

Oracle® Database

Oracle Database New Features



Release 23c

F48428-16

October 2023

ORACLE®

Copyright © 2022, 2023, Oracle and/or its affiliates.

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

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

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

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

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

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

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

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

Contents

1 Introduction

2 Application Development

JSON	2-1
JSON-Relational Duality	2-1
JSON Schema	2-1
XML and JSON Search Index Enhancements	2-1
Changes for JSON Search Index and Data Guide	2-2
DBMS_AQ Support for JSON Arrays	2-2
Enhancement to JSON_TRANSFORM	2-2
JSON Type Support for External Tables	2-3
JSON/JSON_VALUE will Convert PL/SQL Aggregate Type to/from JSON	2-3
JSON_ARRAY Constructor by Query	2-3
JSON_EXPRESSION_CHECK Parameter	2-3
New JSON Data Dictionary Views	2-4
ORDERED in JSON_SERIALIZE	2-4
Precheckable Constraints using JSON SCHEMA	2-4
Predicates for JSON_VALUE and JSON_QUERY	2-4
Tools to Migrate JSON Text Storage to JSON Type Storages	2-5
SQL	2-5
Application Usage Annotations	2-5
Direct Joins for UPDATE and DELETE Statements	2-5
IF [NOT] EXISTS Syntax Support	2-5
New Database Role for Application Developers	2-6
Aggregation over INTERVAL Data Types	2-6
Application Usage Domains	2-6
Automatic PL/SQL to SQL Transpiler	2-6
Client Describe Call Support for Tag Options	2-7
DEFAULT ON NULL for UPDATE Statements	2-7
DESCRIBE Now Supports Column Annotations	2-7
Error Improvement	2-7
Extended CASE Controls	2-8

GROUP BY Column Alias or Position	2-8
Improved TNS Error Messages	2-8
ODP.NET: Annotations	2-8
ODP.NET: SQL Domain Name and Schema	2-9
Oracle C++ Call Interface (OCCI) Support for SQL BOOLEAN Data Type	2-9
Oracle Client Driver Support for SQL BOOLEAN Data Type	2-9
SELECT Without FROM Clause	2-9
SQL BOOLEAN Data Type	2-9
SQL Domain Metadata Support in OCCI	2-10
SQL UPDATE RETURN Clause Enhancements	2-10
SQL*Plus Support for SQL BOOLEAN Data Type	2-10
Table Value Constructor	2-10
Unicode 15.0 Support	2-11
Graph	2-11
Native Representation of Graphs in Oracle Database	2-11
Support for the ISO/IEC SQL Property Graph Queries (SQL/PGQ) Standard	2-11
Property Graph: Native Representation of Graphs in Oracle Database	2-12
Property Graph: Support for the ISO/IEC SQL Property Graph Queries (SQL/PGQ) Standard	2-12
Property Graph: Use JSON Collections as a Graph Data Source	2-12
Property Graph: Use Native Representation of Graphs in Oracle Database with Graph Tools	2-12
Microservices	2-13
Kafka APIs for TxEventQ	2-13
ODP.NET: Advanced Queuing and Transactional Event Queues	2-13
Prometheus/Grafana for Oracle	2-13
Python and REST Drivers for Transactional Event Queues (TxEventQ)	2-14
Saga APIs using Oracle Saga Framework	2-14
Transactional Event Queues (TxEventQ) Propagation	2-14
General	2-15
Efficient Table DDL Change Notification	2-15
Enhanced Inter-Session Communication with DBMS_PIPE	2-15
GB18030-2022 Support	2-15
JDBC RSI Support for Data Load Mode	2-15
ODP.NET: Asynchronous Programming	2-16
Oracle Call Interface (OCI) Support for String Indexed PL/SQL Associative Arrays	2-16
Session Exit on Invalidation	2-16
SQL*Plus ARGUMENT Command	2-16
SQL*Plus OERR Command and Improved HELP Syntax	2-17
SQL*Plus PING Command and Command Line Option	2-17
SQL*Plus SET ERRORDetails Command	2-17
Unicode IVS (Ideographic Variation Sequence) Support	2-17

Java	2-18
Java in the Database: JDK 11 Support Including Modules	2-18
Java in the Database: Web Services Callout Enhancement	2-18
Java in the Database: HTTP and TCP Access While Disabling Other OS Calls	2-18
JavaScript	2-18
Multilingual Engine JavaScript Modules and Environments	2-18
Multilingual Engine Module Calls	2-19
Multilingual Engine Post-Execution Debugging	2-19
Multilingual Engine JavaScript SODA API	2-19
Multilingual Engine JavaScript Support for JSON Data Type	2-19
Application Connectivity	2-20
Reset Database Session State	2-20
Implicit Connection Pooling for Database Resident Connection Pooling (DRCP)	2-20
Implicit Connection Pooling for Oracle Connection Manager in Traffic Director Mode (CMAN-TDM)	2-20
Improved Oracle Connection Manager in Traffic Director Mode (CMAN-TDM) Pool Configuration Settings for Autonomous Database	2-21
JDBC Enhancements to Transparent Application Continuity	2-21
JDBC Extensions for Apps Configuration Providers	2-21
JDBC Support for Kerberos Authentication using JAAS Configuration	2-21
JDBC Support for Kerberos Authentication using User and Password Properties	2-22
JDBC Support for OAuth 2.0 Including OCI IAM and Azure AD	2-22
Multiple Named Pools for Database Resident Connection Pooling (DRCP)	2-22
ODP.NET: Application Continuity	2-23
ODP.NET: OpenTelemetry	2-23
ODP.NET: Pipelining	2-23
ODP.NET: Sharding Split Partitionset	2-23
Oracle Call Interface (OCI) Pipelined Operations	2-24
Oracle Call Interface (OCI) Session Pool Statistics	2-24
Oracle Connection Manager in Traffic Director Mode (CMAN-TDM) Support for Direct Path Applications	2-24
Oracle Connection Manager in Traffic Director Mode (CMAN-TDM) Usage Statistics	2-24
Resumable Cursors	2-25
Shut Down Connection Draining for Database Resident Connection Pooling (DRCP)	2-25
UCP Support for XA Transactions with Sharded Databases	2-25
Database Drivers API Enhancements	2-26
JDBC Support for Database Annotation	2-26
JDBC Support for Pipelined Database Operations	2-26
JDBC Support for Self-Driven Diagnosability	2-26
JDBC Support for SQL BOOLEAN Data Type	2-27
ODBC Support for SQL BOOLEAN Data Type	2-27
Oracle Call Interface (OCI) Support for SQL BOOLEAN Data Type	2-27

Precompiler Support for SQL BOOLEAN Data Type	2-27
UCP Asynchronous Extension	2-28
UCP Support for Self-Driven Diagnosability	2-28

3 Data Analytics

General	3-1
Hybrid Partitioned Tables with Interval and Auto-List Partitioning	3-1
Data Quality Operators in Oracle Database	3-1
Automatic Data Clustering	3-1
Extended Support and Faster Performance for JSON Materialized Views	3-2
Oracle SQL Access to Kafka	3-2
SQL	3-2
Text Indexes with Automatic Maintenance	3-2
Transportable Binary XML	3-3
Concurrent Materialized View Refresh for on-commit	3-3
Enhanced Automatic Indexing	3-3
Enhanced Automatic Materialized Views	3-4
Enhanced Automatic SQL Plan Management	3-4
Enhanced LOB Support for Distributed and Sharded Environments	3-4
Enhanced Parallel Processing Resources Management	3-5
Increased Maximum Size of Inline LOBs of 8000 Bytes	3-5
Materialized View Support for ANSI Joins	3-5
Read-Only Value LOBs	3-5
Semi-Join Materialized Views	3-6
Ubiquitous Search With DBMS_SEARCH Packages	3-6
In-Memory	3-7
Automatic In-Memory Enhancements for Improving Column Store Performance	3-7
Automatic In-Memory Sizing for Autonomous Databases	3-7
In-Memory RAC-Level Global Dictionary	3-7
Vectorized Query Processing: Multi-Level Joins and Aggregations	3-8
Machine Learning - Enhancements	3-8
Automated Time Series Model Search	3-8
Explicit Semantic Analysis Support for Dense Projection with Embeddings	3-8
GLM Link Functions	3-9
Improved Data Prep for High Cardinality Categorical Features	3-9
Lineage: Data Query Persisted with Model	3-9
Multiple Time Series	3-9
Outlier Detection using Expectation Maximization (EM) Clustering	3-10
Partitioned Model Performance Improvement	3-10
XGBoost Support for Constraints and for Survival Analysis	3-10

Spatial	3-11
Spatial: 3D Models and Analytics	3-11
Spatial: REST APIs for GeoRaster	3-11
Workspace Manager: Improved Security when using Oracle Workspace Manager	3-11

4 Data Warehousing/Big Data

General	4-1
Enhanced Partitioning Metadata	4-1
Extended Language Support in Oracle Text	4-1
External Table Partition Values in the File Path	4-1
Logical Partition Change Tracking for Materialized View Refresh and Staleness Tracking	4-2
Staging Tables	4-2

5 Cloud Migration

General	5-1
Classic Queue to Transactional Event Queues (TxEventQ) Online Migration Tool	5-1
Cloud Premigration Advisor Tool for Source Database Migrations to Other Oracle Databases	5-1

6 Cloud Operations

Manageability	6-1
Hybrid Read-Only Mode for Pluggable Databases	6-1
Real-Time SQL Monitoring Enhancements	6-1
Control PDB Open Order	6-1
ODP.NET: Centralized Config Providers	6-2
Optimized Performance for Parallel File System Operations	6-2
Read-Only Users and Sessions	6-2
Continuous Availability	6-3
Enhanced Upgrade of Time Zone Data	6-3
Optimized Read-Write Operations for Database Processes	6-3
Support for the Coexistence of DGPDB and GoldenGate Capture	6-3
General	6-4
Adaptive Result Cache Object Exclusion	6-4
Diagnose and Repair SQL Exceptions Automatically at Compile-Time	6-4
Read-Only Tablespace on Object Storage	6-4
SYSDATE and SYSTIMESTAMP Reflect PDB Time Zone	6-4
Unified Memory	6-5

7 High Availability

Data Guard	7-1
Oracle Data Guard Redo Decryption for Hybrid Disaster Recovery Configurations	7-1
General	7-1
Flashback Time Travel Enhancements	7-1
Minimized Stall and More Accurate Lag Detection in Fast-Start Failover Maximum Performance	7-2
Transaction Guard Support during DBMS_ROLLING	7-2

8 Security

SQL Firewall	8-1
Oracle SQL Firewall Included in Oracle Database	8-1
Encryption	8-1
Transport Layer Security (TLS) 1.3 Now Supported in Oracle Database	8-1
New Parameter to Control the Behavior of TLS Certificate DN Match	8-2
Simplified Transport Layer Security Configuration	8-2
Ability to Configure Transport Layer Security Connections Without Client Wallets	8-2
New sqlnet.ora Parameter to Prevent the Use of Deprecated Cipher Suites	8-2
AES-XTS Encryption Mode Support for TDE Tablespace Encryption	8-3
Changes for TDE Encryption Algorithms and Modes	8-3
Improved and More Secure Local Auto-Login Wallets	8-3
Audit	8-4
Audit Object Actions at the Column Level for Tables and Views	8-4
Control Authorizations for Unified Auditing and Traditional Auditing	8-4
Authentication	8-4
Microsoft Azure Active Directory Integration with Additional Oracle Database Environments	8-4
ODP.NET: Azure Active Directory Single Sign-On	8-5
Increased Oracle Database Password Length	8-5
JDBC-Thin Support for Longer Passwords	8-5
Oracle Data Pump Export and Import Support for Longer Encryption Passwords	8-5
Oracle Call Interface (OCI) and Oracle C++ Call Interface (OCCI) Password Length Increase	8-6
Updated Kerberos Library and Other Improvements	8-6
Enhancements to RADIUS Configuration	8-6
UTL_HTTP Support for SHA-256 and Other Digest Authentication Standards	8-6
XDB HTTP SHA512 Digest Authentication	8-7
Authorization	8-7
Schema Privileges to Simplify Access Control	8-7
Oracle Label Security Triggers Are Now Part of the New LBAC_TRIGGER Schema	8-7

Oracle Data Dictionary Protection Extended to Non-SYS Oracle Schemas with Separation of Duties Protection	8-7
GoldenGate Capture and Apply User Roles	8-8
Autonomous Database	8-8
Identity and Access Management Integration with Oracle Autonomous Cloud Databases	8-8
ODP.NET: Oracle Identity and Access Management	8-8
Oracle Client Increased Database Password Length	8-9
Other	8-9
Secure Distributed Transaction Recovery Background Process (RECO)	8-9
IP Rate Limit in CMAN	8-9
OCI Attributes for Microsoft Azure Active Directory Integration with Additional Oracle Database Environments	8-9
ODP.NET Transport Layer Security (TLS) 1.3	8-10

9 OLTP and Core Database

Availability	9-1
Directory-Based Sharding Method	9-1
Oracle Globally Distributed Database Raft Replication	9-1
Automatic Data Move on Sharding Key Update	9-1
Automatic Transaction Quarantine	9-2
Automatic Transaction Rollback	9-2
Creating Immutable Backups Using RMAN	9-2
Fine-Grained Refresh Rate Control For Duplicated Tables	9-2
Global Partitioned Index Support on Subpartitions	9-3
JDBC Support for Split Partition Set	9-3
Managing Flashback Database Logs Outside the Fast Recovery Area	9-3
Move Data Chunks Between Shardspaces	9-3
New Duplicated Table Type - Synchronous Duplicated Table	9-4
New Partition Set Operations for Composite Sharding	9-4
Oracle Globally Distributed Database Coordinated Backup and Restore Enhancements	9-4
Parallel Cross-Shard DML Support	9-5
PL/SQL Function Cross-Shard Query Support	9-5
Pre-Deployment Diagnostic for Oracle Globally Distributed Database	9-5
RMAN Backup Encryption Algorithm Now Defaults to AES256	9-6
RMAN Operational, Diagnostics, and Upgrade Enhancements	9-6
Simplified Database Migration Across Platforms Using RMAN	9-6
Support for Oracle Database Version Specific RMAN SBT Library	9-6
Blockchain	9-7
Blockchain Table User Chains	9-7
Blockchain Table Row Versions	9-7

Add and Drop User Columns in Blockchain and Immutable Tables	9-7
Blockchain Table Countersignature	9-8
Blockchain Table Delegate Signer	9-8
New Special Privilege Required to Set Long Idle Retention Times for Blockchain and Immutable Tables	9-8
Database Architecture	9-9
Lock-Free Reservations	9-9
Wide Tables	9-9
Consolidated Service Backgrounds for Oracle Instance	9-9
Improve Performance and Disk Utilization for Hybrid Columnar Compression	9-9
Unrestricted Direct Loads	9-10
General	9-10
Unrestricted Parallel DMLs	9-10
ACFS Auto Resize Variable Threshold	9-10
ACFS Cross Version Replication	9-11
ACFS Encryption Migration from OCR to OKV	9-11
ACFS Replication Switchover	9-11
ACFS Snapshots RMAN Sparse Backup and Restore	9-11
ACFS Sparse Backup and Restore of Snapshots	9-11
ACFSutil plogconfig Log Files Wrapping Info	9-12
CEIL and FLOOR for DATE, TIMESTAMP, and INTERVAL Data Types	9-12
Centralized Config Providers	9-12
Oracle Data Pump Filters GoldenGate ACDR Columns from Tables	9-12
PDB Snapshot Carousel ACFS Support	9-13
SQL*Loader Supports SODA (Simple Oracle Document Access)	9-13
Manageability and Performance	9-13
Advanced LOW IOT Compression	9-13
Automatic SecureFiles Shrink for Autonomous Database	9-13
Automatic Storage Compression	9-14
Enhanced Query History Tracking and Reporting	9-14
Fast Ingest (Memoptimize for Write) Enhancements	9-14
Improved Performance of LOB Writes	9-15
Improved System Monitor (SMON) Process Scalability	9-15
Pluggable Database Support in Oracle Data Guard Environments	9-15
Refreshable PDBs in DBCA	9-16

10 Diagnosability

General	10-1
Cluster Health Monitor Improved Diagnosability	10-1
Enhanced Cluster Health Advisor Support for Oracle Pluggable Databases	10-1
Reduce Time to Resolve	10-1

Add Verified SQL Plan Baseline	10-1
CMAN Diagnostics and Logging Enhancements	10-2
DBMS_DICTIONARY_CHECK PL/SQL Package	10-2
Estimate the Space Saved with Deduplication	10-2
Extent-Based Scrubbing	10-3
High Availability Diagnosability Using the DBMS_SCHEDULER Package	10-3
In-Memory Advisor	10-3
Oracle Call Interface (OCI) APIs to Enable Client-Side Tracing	10-3
Rename LOB Segment	10-4
Selective In-Memory Columns	10-4

11 Installation, Upgrade, and Patching

Upgrade	11-1
AutoUpgrade Release Update (RU) Upgrades	11-1
AutoUpgrade Sets Parallelism Based on System Resources	11-1
AutoUpgrade Supports Upgrades with Keystore Access to Databases Using TDE	11-1
AutoUpgrade Unplug-Plugin Upgrades to Different Systems	11-2
REST APIs for AutoUpgrade	11-2

1

Introduction

Oracle Database 23c is the next long term support release of Oracle Database. Oracle Database 23c, **code named “App Simple,”** accelerates Oracle's mission to make it simple to develop and run all data-driven applications.

It's the sum of all the features from the Oracle Database 21c innovation release plus over 300 new features and enhancements. Key focus areas include JSON, graph, microservices, and developer productivity.



Note:

For information about desupported features, see Oracle Database Changes, Desupports, and Deprecations.

JSON Relational Duality

Data can be transparently accessed and updated as either JSON documents or relational tables.

Developers benefit from the strengths of both, which are simpler and more powerful than Object Relational Mapping (ORM).

See [JSON-Relational Duality](#).

Operational Property Graphs in SQL

Developers can now build real-time graph analysis applications against operational data directly in the Oracle Database, utilizing its industry leading security, high availability and performance capabilities.

See [Support for the ISO/IEC SQL Property Graph Queries \(SQL/PGQ\) Standard](#).

Microservice Support

Alongside Oracle's already comprehensive support for microservices, new functionality makes it simpler to implement cross-service transactions.

See [Microservices](#).

Lock-Free Column Value Reservations

Lock-free column value reservations allow applications to reserve part of a value in a column without locking the row; for example, reserve part of a bank account balance or reserve an item in inventory without locking out all other operations on the bank account or item.

See [Lock-Free Reservations](#).

Kafka APIs for TxEventQ

Kafka applications can now run directly against the Oracle Database with minimal code changes, leveraging high performance Transaction Event Queues (TxEventQ).

See [Kafka APIs for TxEventQ](#).

JavaScript Stored Procedures

Developers can now create stored procedures using JavaScript in the database. This functionality also allows developers to leverage the huge number of JavaScript libraries.

See [JavaScript](#).

Priority Transactions

Low priority transactions that block high priority transactions can be automatically aborted. This feature reduces the administrative burden on the DBA while maintaining high transaction throughput.

See [Automatic Transaction Rollback](#).

Application Usage Domains

Application Usage Domains allow developers to declare the intended usage of data (columns) in a centralized and light-weight manner. For example, declare a column holds an email, URL, password, currency, etc. Apps can use Application Usage Domains to automatically generate code or verify values.

See [Application Usage Domains](#).

Many Datatype and SQL Enhancements

The following are among the many datatype and SQL enhancements:

- [SQL BOOLEAN Data Type](#)
- [Direct Joins for UPDATE and DELETE Statements](#)
- [Unicode 15.0 Support](#)
- [SELECT Without FROM Clause](#)
- [GROUP BY Column Alias or Position](#)

Up to 4096 Columns per Table

Database tables now support up to 4096 columns. This feature simplifies the development of applications needing large numbers of attributes, such as ML and IoT.

See [Wide Tables](#).

Improved Machine Learning Algorithms

New improvements to Oracle In-Database Machine Learning algorithms make it simpler to categorize text and data while offering better performance and flexibility.

See [Machine Learning - Enhancements](#).

Sharding Enhancements

New functionality makes it simpler to create and manage shard replicas. New sharding models also improve the distribution of data for shard keys with few unique values.

See [Oracle Globally Distributed Database Raft Replication](#).

Schema Privileges

System privileges can now be granted at the schema level. This feature simplifies the privilege management process and as a result, makes it easy to secure databases.

See [Schema Privileges to Simplify Access Control](#).

Developer Role

A new role allows administrators to quickly assign developers only the privileges they need to design, build, and deploy applications for the Oracle Database.

See [New Database Role for Application Developers](#).

SQL Firewall

Included in Oracle Database, SQL Firewall provides real-time protection against common database attacks by monitoring and blocking unauthorized SQL and SQL injection attacks, no matter the SQL execution path.

See [Oracle SQL Firewall Included in Oracle Database](#).

Azure AD OAuth2 Integration

New functionality enables single sign-on to Oracle Database service instances or on-premises Oracle Databases from Microsoft Azure Cloud.

See [JDBC Support for OAuth 2.0 Including OCI IAM and Azure AD](#).

2

Application Development

JSON

JSON-Relational Duality

JSON Relational Duality Views are fully updatable JSON views over relational data. Data is still stored in relational tables in a highly efficient normalized format but can be accessed by applications in the form of JSON documents.

Duality views provide you with game-changing flexibility and simplicity by overcoming the historical challenges developers have faced when building applications using relational or document models.

Related Resources

[View Documentation](#)

JSON Schema

JSON Schema-based validation is allowed with the SQL condition `IS JSON` and with a PL/SQL utility function. A JSON schema is a JSON document that specifies allowed properties (field names) and the corresponding allowed data types, and whether they are optional or mandatory.

By default, JSON data is schemaless, providing flexibility. However, you may want to ensure that your JSON data contains particular mandatory fixed structures and typing, besides other optional and flexible components, which can be done via JSON Schema validation.

Related Resources

[View Documentation](#)

XML and JSON Search Index Enhancements

The Oracle Text XML search index syntax and JSON search index syntax are now consistent. Additionally, the performance of JSON and XML search indexes has been improved.

Using the same syntax for XML or JSON search indexes and better performance increase productivity.

Related Resources

[View Documentation](#)

Changes for JSON Search Index and Data Guide

JSON search index and JSON data guide are enhanced in these ways. The first two represent changes in the default behavior.

1. When creating a JSON search index, by default a data guide is not created.
2. By default, `DBMS_JSON` procedures `create_view`, `get_view_sql`, and `add_virtual_columns` resolve name conflicts; that is, the default value of parameter `resolveNameConflicts` is `TRUE`, not `FALSE`. This means that if a resulting field name exists in the same data guide then it is suffixed with a new sequence number, to make it unique.
3. Function `json_dataguide` is enhanced to detect ISO 8601 date-time string values, using flag option `DBMS_JSON.detect_datetime`.

When this option is present, field values that are strings in the ISO 8601 date and time formats supported by Oracle are represented in a data guide with the value of field `type` not as `string` but as `timestamp`.

The default changes improve usability and performance for JSON data guides.

Related Resources

[View Documentation](#)

DBMS_AQ Support for JSON Arrays

You can use a JSON data type array as the payload for Advanced Queuing (AQ) message-passing functions, which process an array of messages as a single operation. This applies to the AQ interfaces for C (Oracle Call Interface), PL/SQL, and Java (JDBC).

Advanced Queuing can directly use JSON data for its bulk message passing. With JSON being an increasingly popular format for data exchange, this functionality provides more flexible application development and improves developer productivity.

Related Resources

[View Documentation](#)

Enhancement to JSON_TRANSFORM

`JSON_TRANSFORM` is extended to support right-hand-side path expressions, nested paths, and arithmetic operations. A `SORT` operator is supported which allows sorting the elements in an array.

`JSON_TRANSFORM` is the main SQL operator for modifying JSON data, both for on-disk updates and transient changes in the `SELECT` clause of a query. This enhancement increases update capabilities, such as arithmetic calculations and operations on nested arrays and raises developer productivity.

Related Resources

[View Documentation](#)

JSON Type Support for External Tables

Support for access and direct-loading of JSON-type columns is provided for external tables. JSON data type is supported as a column type in the external table definition. Newline-delimited and JSON-array file options are supported, which facilitates importing JSON data from an external table.

This feature makes it easier to load data into a JSON-type columns.

Related Resources

[View Documentation](#)

JSON/JSON_VALUE will Convert PL/SQL Aggregate Type to/from JSON

The PL/SQL JSON constructor is enhanced to accept an instance of a corresponding PL/SQL aggregate type, returning a JSON object or array type populated with the aggregate type data.

The PL/SQL JSON_VALUE operator is enhanced so that its returning clause can accept a type name that defines the type of the instance that the operator is to return.

JSON constructor support for aggregate data types streamlines data interchange between PL/SQL applications and languages that support JSON.

Related Resources

[View Documentation](#)

JSON_ARRAY Constructor by Query

A subquery can be used as an argument to SQL/JSON function JSON_ARRAY to define the array elements. This functionality is part of the SQL/JSON standard.

This feature increases your developer productivity and higher interoperability with other SQL/JSON standard-compliant solutions.

Related Resources

[View Documentation](#)

JSON_EXPRESSION_CHECK Parameter

A new parameter `JSON_EXPRESSION_CHECK` allows to enable/disable a JSON query check. The values are `on` and `off`. The default is `off`. For now, this parameter is limited to JSON-relational duality views. An error is raised if a JSON path expression on a duality view does not match to an underlying column, for example if the path expression has a typo. The error is raised during query compilations.

This simplifies working with JSON-relational duality views, as incorrect JSON path expressions do not need to be debugged at runtime but instead are flagged at query compilation time (by raising an error).

Related Resources

[View Documentation](#)

New JSON Data Dictionary Views

New dictionary views `*_JSON_INDEXES` and `*_TABLE_VIRTUAL_COLUMNS` have been added.

These new views provide better insight into the database objects that have been created to work with JSON data.

Related Resources

[View Documentation](#)

ORDERED in JSON_SERIALIZE

The SQL function `JSON_SERIALIZE` has an optional keyword `ORDERED`, which reorders the key-value pairs alphabetically (ascending only). It can be combined with optional keywords `PRETTY` and `ASCII`.

Ordering the result of serialization makes it easier for both tools and humans to compare values.

Related Resources

[View Documentation](#)

Precheckable Constraints using JSON SCHEMA

To avoid sending invalid data to the database, an application can often precheck (validate) it. PL/SQL function `DBMS_JSON_SCHEMA.describe` provides JSON schemas that apps can use to perform validation equivalent to that performed by database column-level check constraints, and it records constraints that have no equivalent JSON schema.

Applications can also check which columns are precheckable with a JSON schema by consulting static dictionary views `ALL_CONSTRAINTS`, `DBA_CONSTRAINTS`, and `USER_CONSTRAINTS`.

When you create or alter a table you can use keyword `PRECHECK` to determine whether column check constraints can be prechecked outside the database. If no equivalent JSON schema exists for a given `PRECHECK` column check constraint then an error is raised.

Early detection of invalid data makes applications more resilient and reduces potential system downtime. All applications have access to the same information about whether data for a given column is precheckable, and if so what JSON schema validates it.

Related Resources

[View Documentation](#)

Predicates for JSON_VALUE and JSON_QUERY

JSON path expressions with predicates can be used in `JSON_VALUE` and `JSON_QUERY`. The functionality is part of the SQL/JSON standard.

Applying JSON path expressions more widely for querying JSON data boosts your developer's productivity and simplifies code development.

Related Resources

[View Documentation](#)

Tools to Migrate JSON Text Storage to JSON Type Storages

The new PL/SQL procedure, `DBMS_JSON.json_type_convertible_check`, checks whether existing data stored as JSON text can be migrated to JSON data type. There are several alternative ways to migrate the data after this check succeeds.

Leveraging the binary JSON data type format provides the best performance for processing JSON data. Providing a simple and easy way to ensure existing data can be transformed successfully to binary JSON format helps you to adopt the preferred storage format for JSON data.

Related Resources

[View Documentation](#)

SQL

Application Usage Annotations

Application usage annotations enable you to store and retrieve metadata about database objects. These are name-value pairs or simply a name. These are free-form text fields applications can use to customize business logic or user interfaces.

Application usage annotations help you use database objects in the same way across all applications. This simplifies development and improves data quality.

Related Resources

[View Documentation](#)

Direct Joins for UPDATE and DELETE Statements

Join the target table in `UPDATE` and `DELETE` statements to other tables using the `FROM` clause. These other tables can limit the rows changed or be the source of new values.

Direct joins make it easier to write SQL to change and delete data.

Related Resources

[View Documentation](#)

IF [NOT] EXISTS Syntax Support

DDL object creation, modification, and deletion now support the `IF EXISTS` and `IF NOT EXISTS` syntax modifiers. This enables you to control whether an error should be raised if a given object exists or does not exist.

The `IF [NOT] EXISTS` syntax can simplify error handling in scripts and by applications.

Related Resources[View Documentation](#)

New Database Role for Application Developers

The `DB_DEVELOPER_ROLE` role provides an application developer with all the necessary privileges to design, implement, debug, and deploy applications on Oracle databases.

By using this role, administrators no longer have to guess which privileges may be necessary for application development.

Related Resources[View Documentation](#)

Aggregation over INTERVAL Data Types

You can pass `INTERVAL` data types to the `SUM` and `AVG` aggregate and analytic functions.

This enhancement makes it easier for developers to calculate totals and averages over `INTERVAL` values.

Related Resources[View Documentation](#)

Application Usage Domains

An application usage domain is a dictionary object that belongs to a schema and encapsulates a set of optional properties and constraints for common values, such as credit card numbers or email addresses. After you define an application usage domain, you can define table columns to be associated with that domain, thereby explicitly applying the domain's optional properties and constraints to those columns.

With application usage domains, you can define how you intend to use data centrally. They make it easier to ensure you handle values consistently across applications and improve data quality.

Related Resources[View Documentation](#)

Automatic PL/SQL to SQL Transpiler

PL/SQL functions within SQL statements are automatically converted (transpiled) into SQL expressions whenever possible.

Transpiling PL/SQL functions into SQL statements can speed up overall execution time.

Related Resources[View Documentation](#)

Client Describe Call Support for Tag Options

Annotations enable you to store and retrieve metadata about database objects. These are either name-value pairs or only a name. These are free-form text fields that applications can use to customize business logic or user interfaces.

Annotations help you to use database objects in the same way, across all applications. This simplifies development and improves data quality.

Related Resources

[View Documentation](#)

DEFAULT ON NULL for UPDATE Statements

You can define columns as `DEFAULT ON NULL` for update operations, which was previously only possible for insert operations. Columns specified as `DEFAULT ON NULL` are automatically updated to the specific default value when an update operation tries to update a value to `NULL`.

This feature simplifies application development and removes your need for complex application code or database triggers to achieve the desired behavior. Development productivity is increased and code becomes less error-prone.

Related Resources

[View Documentation](#)

DESCRIBE Now Supports Column Annotations

The SQL*Plus `DESCRIBE` command can now display annotation information for columns that have associated annotations available.

Annotations help you to use database objects in the same way across all applications. This simplifies development and improves data quality.

Related Resources

[View Documentation](#)

Error Improvement

The Oracle Call Interface (OCI) `OCIError()` function has been enhanced to optionally include an Oracle URL with error messages. The URL page has additional information about the Oracle error.

This feature allows users to more easily access information about the cause of the error and the actions that can be taken.

Related Resources

[View Documentation](#)

Extended CASE Controls

The `CASE` statement is extended in PL/SQL to be consistent with the updated definitions of `CASE` expressions and `CASE` statements in the SQL:2003 Standard [ISO03a, ISO03b].

Dangling predicates allow tests other than equality to be performed in simple `CASE` operations. Multiple choices in `WHEN` clauses allow `CASE` operations to be written with less duplicated code.

Related Resources

[View Documentation](#)

GROUP BY Column Alias or Position

You can now use column alias or `SELECT` item position in `GROUP BY`, `GROUP BY CUBE`, `GROUP BY ROLLUP`, and `GROUP BY GROUPING SETS` clauses. Additionally, the `HAVING` clause supports column aliases.

These enhancements make it easier to write `GROUP BY` and `HAVING` clauses. It can make SQL queries much more readable and maintainable while providing better SQL code portability.

Related Resources

[View Documentation](#)

Improved TNS Error Messages

This feature enhances common TNS error messages by providing more information.

Having a better description of errors improves diagnosability.

Related Resources

[View Documentation](#)

ODP.NET: Annotations

Oracle Database annotations enable storing and retrieving metadata about database objects. These are either name-value pairs or only a name. These are free-form text fields that applications can use to customize business logic or user interfaces. All ODP.NET provider types support Oracle Database annotations.

Annotations help you to use database objects in the same way, across all applications. This simplifies development and improves data quality.

Related Resources

[View Documentation](#)

ODP.NET: SQL Domain Name and Schema

A SQL domain is a dictionary object that belongs to a schema and encapsulates a set of optional properties and constraints for common values, such as credit card numbers or email addresses. After you define a SQL domain, you can define table columns to be associated with that domain, thereby explicitly applying the domain's optional properties and constraints to those columns.

With SQL domains, you can define how you intend to use data centrally. They make it easier to ensure you handle values consistently across applications and improve data quality.

All ODP.NET provider types support using SQL domains, as well as retrieving SQL domain schema and name information.

Related Resources

[View Documentation](#)

Oracle C++ Call Interface (OCCI) Support for SQL BOOLEAN Data Type

Oracle C++ Call Interface (OCCI) now supports querying and binding of the new SQL `BOOLEAN` data type.

Using the SQL `BOOLEAN` data type enables applications to represent state more clearly.

Related Resources

[View Documentation](#)

Oracle Client Driver Support for SQL BOOLEAN Data Type

Oracle client drivers support fetching and binding the new `BOOLEAN` database column.

Applications can use the native database `BOOLEAN` column data type with a native driver `BOOLEAN` data type. This enhancement makes working with `BOOLEAN` data types easier for developers.

Related Resources

[View Documentation](#)

SELECT Without FROM Clause

You can now run `SELECT` expression-only queries without a `FROM` clause.

This new feature improves SQL code portability and ease of use for developers.

Related Resources

[View Documentation](#)

SQL BOOLEAN Data Type

Oracle Database now supports the ISO SQL standard-compliant `BOOLEAN` data type. This enables you to store `TRUE` and `FALSE` values in tables and use `BOOLEAN` expressions in SQL statements.

The `BOOLEAN` data type standardizes the storage of `Yes` and `No` values and makes it easier to migrate to Oracle Database.

Related Resources

[View Documentation](#)

SQL Domain Metadata Support in OCCI

Provide access to the SQL domain metadata (domain name and domain schema) for the database columns described in OCCI applications.

Database adds SQL domains to columns and the column metadata need to expose the same in all the data access drivers.

Related Resources

[View Documentation](#)

SQL UPDATE RETURN Clause Enhancements

The `RETURNING INTO` clause for `INSERT`, `UPDATE`, and `DELETE` statements are enhanced to report old and new values affected by the respective statement. This allows developers to use the same logic for each of these DML types to obtain values pre- and post-statement execution. Old and new values are valid only for `UPDATE` statements. `INSERT` statements do not report old values and `DELETE` statements do not report new values.

The ability to obtain old and new values affected by `INSERT`, `UPDATE`, and `DELETE` statements, as part of the SQL command's execution, offers developers a uniform approach to reading these values and reduces the amount of work the database must perform.

Related Resources

[View Documentation](#)

SQL*Plus Support for SQL `BOOLEAN` Data Type

SQL*Plus supports the new SQL `BOOLEAN` data type in SQL statements and the `DESCRIBE` command. Enhancements to the `COLUMN` and `VARIABLE` command syntax have also been made.

SQL*Plus scripts can take advantage of the new SQL `BOOLEAN` data type for easy development.

Related Resources

[View Documentation](#)

Table Value Constructor

The database's SQL engine now supports a `VALUES` clause for many types of statements. This new clause allows for materializing rows of data on the fly by specifying them using the new syntax without relying on existing tables. Oracle supports the `VALUES` clause for the `SELECT`, `INSERT`, and `MERGE` statements.

The introduction of the new `VALUES` clause allows developers to write less code for ad-hoc SQL commands, leading to better readability with less effort.

Related Resources

[View Documentation](#)

Unicode 15.0 Support

The National Language Support (NLS) data files for `AL32UTF8` and `AL16UTF16` character sets are updated to match version 15.0 of the Unicode Standard character database.

This enhancement enables Oracle Database to conform to the latest version of the Unicode Standard.

Related Resources

[View Documentation](#)

Graph

Native Representation of Graphs in Oracle Database

Oracle Database now has native support for property graph data structures and graph queries.

Property graphs provide an intuitive way to find direct or indirect dependencies in data elements and extract insights from these relationships. The enterprise-grade manageability, security features, and performance features of Oracle Database are extended to property graphs. Developers can easily build graph applications using existing tools, languages, and development frameworks. They can use graphs in conjunction with transactional data, JSON, Spatial, and other data types.

Related Resources

[View Documentation](#)

Support for the ISO/IEC SQL Property Graph Queries (SQL/PGQ) Standard

The ISO SQL standard has been extended to include comprehensive support for property graph queries and creating property graphs in SQL. Oracle is among the first commercial software products to support this standard.

Developers can easily build graph applications with SQL using existing SQL development tools and frameworks. Support of the ISO SQL standard allows for greater code portability and reduces the risk of application lock-in.

Related Resources

[View Documentation](#)

Property Graph: Native Representation of Graphs in Oracle Database

Oracle Database now has native support for property graph data structures and graph queries.

Property graphs provide an intuitive way to find direct or indirect dependencies in data elements and extract insights from these relationships. The enterprise-grade manageability, security features, and performance features of Oracle Database are extended to property graphs. Developers can easily build graph applications using existing tools, languages, and development frameworks. They can use graphs in conjunction with transactional data, JSON, Spatial, and other data types.

Related Resources

[View Documentation](#)

Property Graph: Support for the ISO/IEC SQL Property Graph Queries (SQL/PGQ) Standard

The ISO SQL standard has been extended to include comprehensive support for property graph queries and creating property graphs in SQL. Oracle is among the first commercial software products to support this standard.

Developers can easily build graph applications with SQL using existing SQL development tools and frameworks. Support of the ISO SQL standard allows for greater code portability and reduces the risk of application lock-in.

Related Resources

[View Documentation](#)

Property Graph: Use JSON Collections as a Graph Data Source

SQL/PGQ queries can be executed on graphs represented as a JSON column (SQL/PGQ is the ISO standard for property graphs).

Developers can store a graph as a schema-less object in the database. Vertices and edges in a graph can have varying number and types of properties.

Related Resources

[View Documentation](#)

Property Graph: Use Native Representation of Graphs in Oracle Database with Graph Tools

Developers can visualize graphs and graph query results that use the native property graph object in Oracle Database.

Developers can query, analyze, and visualize graphs created by SQL DDL statements using built-in advanced tools in Oracle Database. Existing applications can use this native representation of graphs without changing tools and UI.

Related Resources

[View Documentation](#)

Microservices

Kafka APIs for TxEventQ

Transactional Event Queues (TxEventQ) now support the KafkaProducer and KafkaConsumer classes from Apache Kafka.

Oracle Database can now be used as a source or target for applications using the Kafka APIs.

Related Resources

[View Documentation](#)

ODP.NET: Advanced Queuing and Transactional Event Queues

ODP.NET Core and managed ODP.NET now support Advanced Queuing (AQ) and Transactional Event Queues (TxEventQ) application programming interfaces (APIs) that can be used in modern applications, such as microservices. TxEventQ's highly optimized and partitioned implementation leverages the functions of Oracle database so that producers and consumers can exchange messages at high throughput, by storing messages persistently, and propagate messages between queues on different databases. TxEventQ are a high performance partitioned implementation with multiple event streams per queue, while AQ is a disk-based implementation for simpler workflow use cases.

ODP.NET developers can leverage the same APIs no matter if they use TxEventQ or AQ. The APIs provide access to a robust and feature-rich message queuing systems integrated with Oracle database. It can be used with web, mobile, IoT, and other data-driven and event-driven applications to stream events or communicate with each other as part of a workflow.

Related Resources

[View Documentation](#)

Prometheus/Grafana for Oracle

Prometheus/Grafana for Oracle will provide database metrics for developers running in a Kubernetes/Docker (K8S) environment. Database metrics are stored in Prometheus, a time-series database and metrics tailored for developers are displayed using Grafana dashboards. A database metrics exporter aids the metrics exports from database views into Prometheus time series database.

Developers of modern applications like microservices use observability at the app tier and often overlook the data tier. Data-driven applications don't get a full picture of the execution and performance of the application. Traditional database metrics are seen through AWR reports and Enterprise Manager, which are more targeted to the DBAs and less to the developers. For developers and architects, Prometheus and Grafana have become the tools for configuring metrics dashboards, setting alerts and taking remedial action. Developers can now tie in the app-tier metrics, Kubernetes container metrics, and Oracle database metrics on behalf of the application together in a single dashboard. In addition to metrics, logs and tracing is also enabled to truly get unified observability on a single pane of glass.

Related Resources

[View Documentation](#)

Python and REST Drivers for Transactional Event Queues (TxEventQ)

Database 23c introduces support in new languages for Transactional Event Queues (TxEventQ). TxEventQ can now be used in Python and with REST APIs (REST APIs implemented to be like Kafka's Confluent REST APIs). TxEventQ already has support for PL/SQL, C/C++, and Java using JMS or JDBC.

This feature increases developer productivity by allowing REST APIs applications to take advantage of Transactional Event Queues (TxEventQ) to handle application and data events. With increasing popularity of Python for Machine Learning applications, TxEventQ in the Oracle Database can now be part of the Machine Learning application data and events infrastructure.

Related Resources

[View Documentation](#)

Saga APIs using Oracle Saga Framework

Oracle Saga APIs are implemented in the database and provide a framework to implement transactional semantics for microservices built with the Oracle Database. The orchestrator Saga framework provides a way to maintain atomic data consistency across microservices.

Sagas are concurrent and execute local transactions in each participant database making it more efficient than distributed ACID transactions, thereby simplifying application code and increasing developer productivity.

Related Resources

[View Documentation](#)

Transactional Event Queues (TxEventQ) Propagation

Queues are used widely to send and receive events and messages between applications, increasingly being built as microservices. Transactional Event Queues (TxEventQ) are queues built into the Oracle Database. Queue propagation allows multiple databases to act as producers and consumers of events and messages. Producer applications can send events in queues in one database, set up queue propagation to a remote database, and Consumer applications can consume events in queues in the remote database.

Queue propagation is used to consolidate critical events and data from remote locations to a central location for consolidated processing. Propagation is used to operate Transactional Event Queues (TxEventQ) as a reliable and secure Event Mesh, with multiple queues across multiple databases participating to send events and messages across the enterprise reliably and to remote subscribers with permissions. TxEventQ supports exactly-once messaging which makes it simpler to build and test applications.

Related Resources

[View Documentation](#)

General

Efficient Table DDL Change Notification

Applications can now be notified when DDLs occur on tables.

Applications that need or want to be aware of table metadata can be notified of DDL changes rather than having to continuously poll for them.

Related Resources

[View Documentation](#)

Enhanced Inter-Session Communication with DBMS_PIPE

DBMS_PIPE, an in-database messaging framework for inter-session communication, got enhanced to support a broader set of applications and use cases. DBMS_PIPE now can share messages across multiple database sessions with concurrent reads, provides more flexibility in managing messages overall, and supports persistence and inter-instance and inter-database communication through object store buffering.

Providing a more comprehensive in-database inter-session messaging and communication enables more applications to take advantage of DBMS_PIPE, improving the reliability and scalability of applications. It also increases developer productivity by eliminating the need for more complex application architectures.

Related Resources

[View Documentation](#)

GB18030-2022 Support

The implementation of the Oracle client character set ZHS32GB18030 is updated to support the latest GB18030-2022 standard.

This feature enables Oracle Database to conform to the latest edition of the GB18030 standard, which is required for all software products sold in China.

Related Resources

[View Documentation](#)

JDBC RSI Support for Data Load Mode

In RSI stream mode, connection and prepared statement instances are created for every batch. With the new data load mode, the instances are created once and saved in the thread's local context.

This feature brings faster data ingestion into the Oracle Database.

Related Resources

[View Documentation](#)

ODP.NET: Asynchronous Programming

ODP.NET supports the .NET Task Asynchronous Programming (TAP) model with the core and managed drivers.

With support for TAP and the `async` and `await` keywords, ODP.NET data access operations are more responsive and easier to develop for asynchronicity.

Related Resources

[View Documentation](#)

Oracle Call Interface (OCI) Support for String Indexed PL/SQL Associative Arrays

PL/SQL string indexed associative arrays are now supported by Oracle Call Interface (OCI). Applications can natively pass these associative arrays between the database and the client application allowing for creating, binding, and manipulating of this collection type.

This feature allows for more straightforward and less error-prone code development.

Related Resources

[View Documentation](#)

Session Exit on Invalidation

Set `SESSION_EXIT_ON_PACKAGE_STATE_ERROR` to `true` to force a hard session exit when a session's state has been invalidated.

Exiting sessions after state invalidation avoids errors that can occur when applications mishandle invalid state.

Related Resources

[View Documentation](#)

SQL*Plus ARGUMENT Command

A new `ARGUMENT` command lets users of batch scripts control how SQL*Plus treats script argument variables for which the users have not explicitly set values. With this command, users are now able to control when to prompt for input or use a default value for each unset script argument.

This feature gives SQL script processing more resiliency and flexibility, allowing script actions to be customized by users if they want to alter parameter values.

Related Resources

[View Documentation](#)

SQL*Plus OERR Command and Improved HELP Syntax

A new `OERR` command in SQL*Plus allows users to see Oracle error message Cause and Action text within SQL*Plus for a user-supplied error number. The existing `HELP` command has also been enhanced to show the same text.

This feature allows developers to immediately get more information about error messages.

Related Resources

[View Documentation](#)

SQL*Plus PING Command and Command Line Option

A new SQL*Plus `PING` command and equivalent command line option can be used to show the round-trip time from SQL*Plus to either the network listener or to the database.

The network listener check is equivalent to the traditional `tnsping` command line utility that administrators use to check basic network connectivity. The option to check the database round-trip time is commonly used as a liveness check to ensure that the database itself is reachable.

This feature gives users of SQL*Plus power to verify basic connectivity, which is useful in many troubleshooting or post-install scenarios.

Related Resources

[View Documentation](#)

SQL*Plus SET ERRORDetails Command

A new `SET ERRORDetails` command lets users decide whether additional information should be displayed when Oracle errors are generated in failure scenarios. Additional information that can be displayed is the error help URL and the message Cause and Action text.

This feature improves the developer experience by allowing faster troubleshooting.

Related Resources

[View Documentation](#)

Unicode IVS (Ideographic Variation Sequence) Support

The new `UCA1210_JAPANESE_IVS` collation allows the processing of Unicode Ideographic Variation Sequence (IVS) in Japanese text. The SQL functions `LENGTHC()`, `SUBSTRC()`, `INSTRC()`, and `LIKEC()` are also enhanced to count IVSs as single complete characters.

This feature enables application developers to build applications supporting Unicode IVS. It is an important requirement for markets, such as Japan, where processing of data including names such as person names, place names, and historic texts often need to support ideographic characters represented in Unicode IVS.

Related Resources

[View Documentation](#)

Java

Java in the Database: JDK 11 Support Including Modules

In this release, the Oracle JVM infrastructure has been re-architected to support JDK 11 capabilities including the Java module system.

This feature fosters productivity through the design or reuse and execution of code and libraries based on Java 11, inside the database.

Related Resources

[View Documentation](#)

Java in the Database: Web Services Callout Enhancement

This feature furnishes an enhanced implementation of the Web Services Call-Out Utility. Java, PL/SQL, and SQL can now perform a more efficient Web Services Callout.

This feature fosters extensibility and productivity by allowing the Oracle database to invoke external Web Services.

Related Resources

[View Documentation](#)

Java in the Database: HTTP and TCP Access While Disabling Other OS Calls

Oracle JVM now offers a more flexible Lockdown Profile configuration for on-premises and cloud database services (for example, the Autonomous Database). HTTP and TCP callouts can now be enabled separately from other OS calls to allow deployments that depend on HTTP and TCP access.

This feature couples extensibility (for example, making HTTP and TCP callouts) with enhanced security for Java code running in the database.

Related Resources

[View Documentation](#)

JavaScript

Multilingual Engine JavaScript Modules and Environments

Multilingual Engine (MLE) Modules and Environments allow JavaScript code to persist and be managed in the database. Call specifications provide a means to call JavaScript functions from an MLE module anywhere you can call PL/SQL functions.

The introduction of JavaScript Modules and Environments as schema objects in Oracle Database allows developers to follow established and well-known workflows

used in client-side JavaScript development. Complex projects can be broken down into smaller, more manageable pieces worked on independently by team members.

Related Resources

[View Documentation](#)

Multilingual Engine Module Calls

Multilingual Engine (MLE) Module Calls allow developers to invoke JavaScript functions stored in modules from SQL and PL/SQL. Call Specifications written in PL/SQL link JavaScript to PL/SQL code units.

Thanks to Module Calls, developers can use JavaScript functions anywhere PL/SQL functions are called.

Related Resources

[View Documentation](#)

Multilingual Engine Post-Execution Debugging

Oracle Multilingual Engine (MLE) allows developers to debug their JavaScript code by conveniently and efficiently collecting runtime state while the program is being processed, a method referred to as post-execution debugging. After the code has finished running, the collected data can be used to analyze program behavior, discover, and fix bugs.

Post-execution debugging offers a convenient way to extract runtime state information from a JavaScript code unit at runtime without having to change the observed code.

Related Resources

[View Documentation](#)

Multilingual Engine JavaScript SODA API

Simple Oracle Document Access (SODA) is a set of NoSQL-style APIs that let you create and store collections of documents (in particular JSON) in Oracle Database, retrieve them, and query them, without needing to know SQL or how the documents are stored in the database. With the introduction of MLE, JavaScript support for SODA documents exists for client-side and server-side development.

Supporting the Simple Oracle Document Access (SODA) API in JavaScript gives developers a choice between using JSON in a relational or No-SQL way, simplifying the development process and improving the portability of code.

Related Resources

[View Documentation](#)

Multilingual Engine JavaScript Support for JSON Data Type

Support for JavaScript Object Notation (JSON) is an integral part of Oracle database. Oracle supports JSON natively with relational database features, including transactions, indexing, declarative querying, and views. A rich set of SQL functions is available to manipulate JSON in a relational model. Oracle Multilingual Engine (MLE) fully supports JSON: both dynamic MLE as well as MLE Module Calls support interactions with the JSON data type.

JSON and JavaScript objects are closely related, forming a natural match in such a way that makes working with JSON very easy with JavaScript code.

Related Resources

[View Documentation](#)

Application Connectivity

Reset Database Session State

The reset database session state feature clears the session state set by the application when the request ends. The `RESET_STATE` database service attribute cleans up dirty sessions so that the applications cannot see the state of these sessions. This feature applies to all applications that connect to the database using database services.

This feature uses the `RESET_STATE` attribute on the database service to direct the database to clean the session state at the end of each request so that developers do not have to clean the session state manually. By using this feature, you ensure that there are no data leaks from a previous session.

Related Resources

[View Documentation](#)

Implicit Connection Pooling for Database Resident Connection Pooling (DRCP)

This feature enables the automatic assignment of DRCP servers to and from an application connection at runtime when the application starts and finishes database operations, even if the application does not explicitly close the connection.

This feature can provide better scalability and efficient usage of database resources for applications that do not use application connection pooling.

Related Resources

[View Documentation](#)

Implicit Connection Pooling for Oracle Connection Manager in Traffic Director Mode (CMAN-TDM)

Client applications that do not use an application connection pool can take advantage of CMAN-TDM Proxy Resident Connection Pooling (PRCP) without making any application changes.

The new feature enables the automatic assignment of PRCP servers to and from an application connection at runtime when the application starts and finishes database operations even if the application does not explicitly close the connection.

This feature can reduce the size of PRCP pools required. It provides better scalability and efficient usage of resources for applications that do not use Oracle Session Pooling or Universal Connection Pooling (UCP).

Related Resources[View Documentation](#)

Improved Oracle Connection Manager in Traffic Director Mode (CMAN-TDM) Pool Configuration Settings for Autonomous Database

Oracle Connection Manager in Traffic Director Mode (CMAN-TDM) has new Proxy Resident Connection Pooling (PRCP) configuration settings for use with Autonomous Database. Per-PDB PRCP pools can be enabled, allowing you to consolidate connection pools for each PDB and share these sessions across multiple services that belong to the same PDB. The maximum PRCP pool size can be dynamically configured based on the new `cman.ora` parameter `TDM_PERPDB_PRCP_CONNFACOR` and the Oracle Compute Unit (OCPU) count allocated to each PDB.

The per-PDB PRCP mode provides efficient usage of database resources by reducing the number of pools in a CMAN-TDM gateway. Pool sizing can also now be more autonomous, reducing the need for manual re-configuration.

Related Resources[View Documentation](#)

JDBC Enhancements to Transparent Application Continuity

This feature allows, when the RDBMS server supports it, templates (for example, stable restorable attributes), which are cross-session (one template might be used by multiple sessions). This feature also brings the ability to avoid the combinatorial explosion of templates by quarantining session states that are different in most sessions and therefore cannot be shared.

This feature simplifies high availability by moving most application continuity configurations to the server-side. Java applications inherit transparently (that is, no code required) the latest server-side enhancements.

Related Resources[View Documentation](#)

JDBC Extensions for Apps Configuration Providers

JDBC instrumentalization for securely pulling Java Apps configuration from central stores such as Azure Config Store or OCI Object store or any JSON file accessible from generic web servers.

This feature simplifies Java application configuration in multi-Cloud environments.

Related Resources[View Documentation](#)

JDBC Support for Kerberos Authentication using JAAS Configuration

By default, the Oracle JDBC Thin driver uses the default Kerberos login module, bundled with Oracle JDK (`com.sun.security.auth.module.Krb5LoginModule`). This feature enables

applications that want to override the default behavior to specify a JAAS configuration file through the connection properties.

This feature provides flexibility with Kerberos Authentication configuration.

Related Resources

[View Documentation](#)

JDBC Support for Kerberos Authentication using User and Password Properties

This feature enables the users to configure Kerberos Principal and Password through the User and Password properties. The JDBC Thin driver takes care of initializing the `KerberosLoginModule` on behalf of the applications.

This feature simplifies Kerberos Authentication configuration.

Related Resources

[View Documentation](#)

JDBC Support for OAuth 2.0 Including OCI IAM and Azure AD

The Oracle JDBC driver provides support for OAuth 2.0 authentication for Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Cloud Service or the Azure Active Directory.

This feature simplifies Java application authentication to the Oracle Autonomous Database using OAuth 2.0 in multi-Cloud environments (OCI, Azure) in lieu of traditional credentials mechanisms such as username/password or strong authentication mechanisms, such as Kerberos or Radius.

Related Resources

[View Documentation](#)

Multiple Named Pools for Database Resident Connection Pooling (DRCP)

Database Resident Connection Pooling (DRCP) now supports multiple named pools. New `DBMS_CONNECTION_POOL.ADD_POOL()` and `DBMS_CONNECTION_POOL.REMOVE_POOL()` procedures are added. Oracle Net connection string syntax is enhanced so a pool name can be specified for each connection. Existing procedures can be used to start, stop, or configure the named pools. Existing `GV$` and `V$` views show the appropriate pool name(s) in use.

Having multiple pools allows finer control on the DRCP pool usage. It helps prevent situations where some applications dominate the use of a single pool.

Related Resources

[View Documentation](#)

ODP.NET: Application Continuity

ODP.NET Core and managed drivers now support Application Continuity (AC) and Transparent Application Continuity (TAC). AC and TAC mask outages from end users and applications by recovering the in-flight database sessions following recoverable outages, including transactions. The recovery is transparent such that the end user merely experiences a slightly delayed execution, but no perceptible outage nor error.

AC and TAC improve the user experience for both unplanned outages and planned maintenance. They enhance the fault tolerance of systems and .NET applications that use an Oracle database. Developers can use AC and TAC with existing .NET apps without making any code changes.

Related Resources

[View Documentation](#)

ODP.NET: OpenTelemetry

OpenTelemetry is a popular open-source observability framework for instrumenting, generating, collecting, and exporting telemetry data. It provides a common specification and protocol so that multiple services can furnish a unified version of traces, metrics, and logs.

Numerous managed ODP.NET and ODP.NET Core APIs have been instrumented to support OpenTelemetry observability and standards. Developers and operators can customize the ODP.NET metrics collected with OpenTelemetry.

With OpenTelemetry support, monitoring, tracking, and analyzing how ODP.NET operations interact in cloud computing, microservices and distributed systems becomes easier using this industry standard.

Related Resources

[View Documentation](#)

ODP.NET: Pipelining

ODP.NET core and managed drivers support pipelining for its database communication. It allows subsequent database requests to be sent and queued transparently even while ODP.NET awaits a database response.

Pipelining improves overall app performance and allows database resources to be used more effectively. ODP.NET does not need to wait for the database to respond from previous requests before submitting subsequent requests.

Related Resources

[View Documentation](#)

ODP.NET: Sharding Split Partitionset

ODP.NET supports sharding split partitionset events in core and managed drivers. A split partitionset is an operation on the sharding database performed when moving data within a specified super sharding key to a different shardspace.

ODP.NET connection pools receive database events about data in a chunk being split and moved across these partition sets. The provider updates its sharding topology appropriately and transparently so that end users are dispensed the correct connection to use.

Related Resources

[View Documentation](#)

Oracle Call Interface (OCI) Pipelined Operations

Oracle Call Interface (OCI) has been enhanced to support pipelining of operations. Pipelining enables applications to submit multiple database operations without waiting for a response from the server. The application has control over when the responses to the pipelined operations are harvested. This allows applications to continue work without being blocked while the database is generating results.

This feature is used to increase the overall throughput and responsiveness of applications and languages that use OCI. Pipelining reduces the server and client idle times in comparison with the traditional request-response model.

Related Resources

[View Documentation](#)

Oracle Call Interface (OCI) Session Pool Statistics

The Oracle Call Interface (OCI) session pool usage statistics can be viewed.

The statistics help in tuning pool sizes for better performance, and aid in understanding the life cycle of connections.

Related Resources

[View Documentation](#)

Oracle Connection Manager in Traffic Director Mode (CMAN-TDM) Support for Direct Path Applications

Oracle Database's Set Current Schema and Direct Path API features are now supported by Oracle Connection Manager in Traffic Director Mode (CMAN-TDM).

This feature enables more client applications to leverage CMAN-TDM connection multiplexing capabilities.

Related Resources

[View Documentation](#)

Oracle Connection Manager in Traffic Director Mode (CMAN-TDM) Usage Statistics

A new `V$TDM_STATS` view can be used to query usage statistics for CMAN-TDM per-PDB Proxy Resident Connection Pools (PRCP), such as the number of active client

connections in the connection pool, the number of busy and free server connections, the maximum number of connections reached, and more.

Providing usage statistics helps to improve the monitoring and tuning of CMAN-TDM.

Related Resources

[View Documentation](#)

Resumable Cursors

Resumable cursors, those that span transactions, will be replayable with Transparent Application Continuity. Such cursors are common in batch processing (such as loading sets of records) and require special handling to reposition those cursors during replay with Transparent Application Continuity.

Broadened support with Transparent Application Continuity for applications that rely upon resumable cursors, those that span commits. These cursors are very common in repetitive batch operations, looping through sets of records for updates and inserts with a commit for each set of records. Now, TAC will be able to replay the transactions that were interrupted (and not yet committed).

Related Resources

[View Documentation](#)

Shut Down Connection Draining for Database Resident Connection Pooling (DRCP)

A new, optional `DRAINTIME` argument to `DBMS_CONNECTION_POOL.STOP_POOL()` allows active DRCP pools to be closed after a specified connection drain time, or be closed immediately without waiting for connections to be idle.

This feature gives DBAs better control over DRCP usage and configuration.

Related Resources

[View Documentation](#)

UCP Support for XA Transactions with Sharded Databases

This feature allows sharded database connections to participate in eXtended Architecture (XA) transactions managed by WebLogic Server Transaction Manager.

This feature allows reliable XA transactions coupled with the scalability of sharded databases.

Related Resources

[View Documentation](#)

Database Drivers API Enhancements

JDBC Support for Database Annotation

Annotation is a mechanism to store application metadata centrally in the database. Annotations can be specified at creation time (CREATE) or at modification time (ALTER). An individual annotation has a name and an optional value. Both the name and the value are freeform text fields. A schema object can have multiple annotations. JDBC furnishes the `getAnnotations()` method with two signatures (as illustrated below). It returns the annotation associated with the specified table or view. It returns `null` if there is no annotation for the given object.

```
getAnnotations?(java.lang.String objectName, java.lang.String  
domainName, java.lang.String domainOwner) throws java.sql.SQLException
```

```
getAnnotations?(java.lang.String objectName, java.lang.String  
columnName, java.lang.String domainName, java.lang.String domainOwner)  
throws java.sql.SQLException
```

This feature enables sharing metadata across applications and microservices thereby increasing metadata management and productivity.

Related Resources

[View Documentation](#)

JDBC Support for Pipelined Database Operations

In the previous releases, the JDBC driver would not allow another database call to start until the current call had been completed however, with asynchronous and reactive programming, Java applications could perform non-database operations, in the meantime. In this release, the database server and the Oracle JDBC-Thin both support pipelining database operations. Java applications can now asynchronously submit several SQL requests to the server without waiting for the return of the preceding calls.

The combination of Java reactive and asynchronous programming (JDBC Reactive Extension, Reactive Streams (R2DB and Virtual Threads) with database support for pipelining fosters high throughput.

Related Resources

[View Documentation](#)

JDBC Support for Self-Driven Diagnosability

This feature eliminates the need to switch from the non-logging JAR files (for example, `ojdbcXX.jar`) to the debug JAR files (for example, `ojdbcXX_g.jar`) for logging purposes. In addition, it enables logging in the following three ways: logging per connection, logging at the tenant level, or logging globally.

This feature furnishes increased productivity and ease of use for Java applications. It greatly simplifies the debugging of Java applications by removing the need to switch from the production jars to the debug jars.

Related Resources

[View Documentation](#)

JDBC Support for SQL BOOLEAN Data Type

This feature exposes the Oracle RDBMS `BOOLEAN` data type to Java through a new `BOOLEAN` data type in `oracle.jdbc.OracleType Enum`, and `DatabaseMetadata`. This feature also performs the implicit conversion of character and number data types to `BOOLEAN` data types.

Java applications can take advantage of the new JDBC support for the standard JDBC `BOOLEAN` data type. The benefits include: increased portability and the ease of development fostered by the implicit conversion of character and number to `BOOLEAN`.

Related Resources

[View Documentation](#)

ODBC Support for SQL BOOLEAN Data Type

ODBC now supports the new SQL `BOOLEAN` data type.

Using the SQL `BOOLEAN` data type enables applications to represent the state more clearly.

Related Resources

[View Documentation](#)

Oracle Call Interface (OCI) Support for SQL BOOLEAN Data Type

Oracle Call Interface (OCI) now supports querying and binding of the new SQL `BOOLEAN` data type.

Using the SQL `BOOLEAN` data type enables applications to represent state more clearly.

Related Resources

[View Documentation](#)

Precompiler Support for SQL BOOLEAN Data Type

The Pro*C and Pro*COBOL precompilers now support querying and binding of the new SQL `BOOLEAN` data type.

Using the new data type makes it easier to represent boolean state in applications instead of using a character column to indicate Y or N.

Related Resources

[View Documentation](#)

UCP Asynchronous Extension

Universal connection pool (UCP) is extended with asynchronous (reactive) database calls.

This extension furnishes high scalability and throughput to Java applications.

Related Resources

[View Documentation](#)

UCP Support for Self-Driven Diagnosability

The new universal connection pool (UCP) diagnosability feature provides the following capabilities:

- When logging is enabled (it is disabled by default), log records are written into an in-memory ring buffer.
- Tracing is enabled by default. A tracing event dumps the ring buffer into either a data-source-specific buffer or a common buffer.

This feature fosters productivity (for example, real-time debugging) and ease of use for Java applications using the UCP.

Related Resources

[View Documentation](#)

3

Data Analytics

General

Hybrid Partitioned Tables with Interval and Auto-List Partitioning

You can create Hybrid Partitioned Tables using single-level partitioning with interval and automatic list partitioning. This is in addition to existing support for single-level partitioning and range and list partitioning.

These extensions to Hybrid Partitioned Tables in Oracle Database provide a user-friendly partitioning strategy.

Related Resources

[View Documentation](#)

Data Quality Operators in Oracle Database

This release introduces the following two new string matching operators based on approximate or "fuzzy" string matching.

- `PHONIC_ENCODE` converts words or phrases into language-specific codes based on pronunciation.
- `FUZZY_MATCH`, which is language-neutral, gauges the textual similarity between two strings.

The new phonic encoding and fuzzy matching methods enable more sophisticated matching algorithms to be run directly on data in the database rather than only in external applications, providing improved matching performance and efficiency, for example in data de-duplication, linking or enhancement.

Related Resources

[View Documentation](#)

Automatic Data Clustering

Oracle Database automatically and transparently clusters storage-based data in response to the type of queries used by the application workload. This allows the workload to make more efficient use of data access optimizations, such as storage indexes, zone maps, and join zone maps.

This feature significantly improves performance for data warehousing workloads based on zone maps or storage indexes. Once data is clustered, the performance of data-scanning queries improves because larger contiguous areas (or zones) of storage are pruned or skipped when they do not contain the data being matched by a particular query.

Related Resources

[View Documentation](#)

Extended Support and Faster Performance for JSON Materialized Views

Materialized views of JSON tables have been enhanced with the ability to fast refresh more types of Materialized Views of JSON tables as well as Query Rewrite support for these Materialized Views.

The performance for JSON table Materialized Views is significantly improved through better fast refresh capabilities and better query rewrite capabilities for more workloads. You can use JSON table Materialized Views more broadly in your applications, with better performance and less resource utilization.

Related Resources

[View Documentation](#)

Oracle SQL Access to Kafka

Oracle SQL Access to Kafka (DBMS_KAFKA) provides efficient, reliable, and scalable access to data streams from Apache Kafka and OCI Streaming Service. Streaming data can be queried via SQL or loaded into Oracle database tables.

Oracle Database provides efficient, reliable, and scalable integration with Apache Kafka using the DBMS_KAFKA APIs. This API enables Oracle Database to consume data from external data streams without the need for costly, complex direct application connections using proprietary interfaces. Oracle SQL Access to Kafka enables you to use Oracle Databases rich analytic capabilities across all your data.

Related Resources

[View Documentation](#)

SQL

Text Indexes with Automatic Maintenance

You can specify a new automatic maintenance mode for text indexes using the `MAINTENANCE AUTO` index parameter. This method automates the `CTX_DDL.SYNC_INDEX` operation. This is now the default synchronization method for new indexes.

With this method, newly created text indexes do not require you to specify a synchronization interval or manually run a `SYNC_INDEX` operation. A background process automatically performs these tasks without user intervention. This helps in synchronizing a large number of indexes in an optimal manner, and also eliminates the manual or time-based `SYNC` operations. By using a background job rather than the database scheduler, it avoids scheduling conflicts and the risk of running out of available jobs. Overall it makes for simpler, more resilient applications and better utilization of hardware resources.

Related Resources

[View Documentation](#)

Transportable Binary XML

Transportable binary XML (TBX) is a new self-contained XMLType storage method. TBX supports sharding, XML search index, and Exadata pushdown operations, providing better performance and scalability than other XML storage options.

You can migrate existing XMLType storage of a different format to TBX format in any of these ways:

- Insert-as select or create-as-select
- Online Redefinition
- Data Pump

Transportable binary XML (TBX) provides better performance and scalability. With the support of more database architectures, such as sharding or Exadata, and its capability to easily migrate and exchange XML data among different servers, containers, and PDBs, TBX allows your applications to take full advantage of both this new XML storage format on more platforms and architectures.

Related Resources

[View Documentation](#)

Concurrent Materialized View Refresh for on-commit

Materialized view refresh provides concurrent refresh, where multiple sessions can refresh the same on-commit materialized views simultaneously without the need for serialization.

Concurrent refresh broadens the applicability of materialized views for your applications and helps make application development simpler. It provides faster refresh and more up-to-date materialized views.

Related Resources

[View Documentation](#)

Enhanced Automatic Indexing

Indexes incur a maintenance overhead during DML operations. This can work against their improvements to data access performance. The enhancements to Automatic Indexing take a broader view than in previous releases and account for index maintenance costs when deciding which indexes will benefit the workload as a whole. Columns filtered using range predicates are considered for indexes and function-based indexes are now supported. This further increases the scope of Automatic Indexing effectiveness.

Automatic Indexing better assesses the impact of DML operations in your database when choosing automatic indexing. Your performance benefits by determining the overall advantage of an index to your workload.

Related Resources

[View Documentation](#)

Enhanced Automatic Materialized Views

Automatic materialized views have been enhanced to include automatic partitioning. In addition, there is a more accurate internal cost model for automatic materialized view selection, which considers both access benefits and maintenance (refresh) costs, as well as the frequency of execution.

Rewrite capabilities have been broadened, including outer join queries with filter predicates.

Enhancing Automatic Materialized Views with more accurate cost-benefit analysis and broader usability optimizes the management of your materialized view eco system and improves the overall performance of your system.

Related Resources

[View Documentation](#)

Enhanced Automatic SQL Plan Management

Automatic SQL plan management has been enhanced to detect and repair SQL performance regressions more quickly. SQL plan changes are detected at parse-time and, after initial execution, SQL performance is compared with the performance of previous SQL execution plans. If a performance degradation is detected, the plan is repaired accordingly.

With automatic SQL Plan Management, your application service levels improves, and impacts caused by SQL performance (plan) regressions are minimized and addressed transparently and proactively.

Related Resources

[View Documentation](#)

Enhanced LOB Support for Distributed and Sharded Environments

Distributed LOBs are LOBs that are fetched from one server to another and may optionally be returned to the client. Shared LOBs are an extension of distributed LOBs where LOBs are transported between shards or between a shard and the shard coordinator. In previous versions, support for sharded and distributed LOBs were limited to persistent LOBs, and temporary LOBs only where they originate from JSON operations. Now all temporary LOBs (including Value LOBs) and new increased-length inline LOBs are usable as distributed and sharded LOBs.

You can now work with inline LOBs, value LOBs, and all temporary LOBs in distributed and sharded environments.

You experience improved performance, scalability, and garbage collection when you work with temporary LOBs, thus improving your developer productivity and application resilience.

Related Resources

[View Documentation](#)

Enhanced Parallel Processing Resources Management

Parallel processes are released pro-actively before individual statements using parallelism are finished. For example, an uncommitted parallel DML operation or a partially fetched parallel `SELECT` statement with 2 Parallel Server Sets (Producer-Consumer) will release one of the Parallel Server Sets as soon as it has finished working, freeing half of the parallel process for use of other statements.

Releasing parallel processes as early as possible and making them usable for other statements optimizes the utilization of your available resources, improving the overall performance of your systems and applications.

Related Resources

[View Documentation](#)

Increased Maximum Size of Inline LOBs of 8000 Bytes

LOB values are stored either in the table row (inline) or outside of the table row (out-of-line). The maximum size of the inline LOB is increased to 8000 bytes, allowing larger LOB values being stored inside a row. Earlier, the maximum size was 4000.

This provides better input-output performance while processing LOB columns. You can experience the improved performance while running operations, such as full table scans, range scans, and DML.

Related Resources

[View Documentation](#)

Materialized View Support for ANSI Joins

Materialized Views in Oracle Database support full rewrite capabilities for SQL statements using ANSI join syntax and for Materialized View definitions using ANSI join syntax.

Full support of ANSI joins with materialized view rewrite provides a significant performance improvement. Many queries, particularly ones generated by SQL Tools and Reports, often use ANSI join syntax. This enhancement allows such tools to benefit from materialized views for query rewrite regardless of the syntax used by joins.

Related Resources

[View Documentation](#)

Read-Only Value LOBs

Value LOBs, a read-only subset of Temporary LOBs, are valid for a SQL fetch duration and optimize the reading of LOB values in the context of a SQL query. Many applications use LOBs to store medium-sized objects, about a few megabytes in size, and you want to read the LOB value in the context of a SQL query.

Value LOBs provide faster read performance and get automatically freed when the next fetch for a cursor is performed, preventing the accumulation of temporary LOBs and simplifying the LOB management within your application.

Value LOBs provide faster read performance than classical reference LOBs for your workload and don't need specific LOB management in your application. Using Value LOBs improves your application performance and makes implementing applications with LOBs simpler and more manageable.

Related Resources

[View Documentation](#)

Semi-Join Materialized Views

Semi-Join Materialized View Rewrite is a unique form of query rewrite. A single, large unified dimension table in the query is replaced with one or more join-specific materialized views. In a unified dimension data model, where multiple dimension tables are merged into a single large dimension table, semi-join Materialized Views materialize one or more of the joins of such a single, large unified dimension table with the fact table.

This new type of Materialized View significantly improves the runtime and resource consumption for complex analytical operations. Semi-join Materialized Views are especially beneficial when the number of applicable dimension keys derived from the large unified dimension table (through semi-join) is small.

Related Resources

[View Documentation](#)

Ubiquitous Search With DBMS_SEARCH Packages

The new `DBMS_SEARCH` PL/SQL package allows the indexing of multiple schema objects in a single index. You can add a set of tables, external tables, or views as data sources into this index. All the columns in the specified sources are indexed and available for a full-text search.

With a simplified set of `DBMS_SEARCH` APIs, you can create, maintain, and query indexes as follows:

- Create indexes across multiple objects.
- Easily add data sources (specific tables or views) into indexes or remove data sources from them.
- Perform a full-text search within a single data source or across multiple sources using the same index.
- Automatically maintain synchronization of indexes.

This simplifies indexing tasks that were previously performed using the `USER_DATASTORE` procedures, thus enhancing developer productivity.

Related Resources

[View Documentation](#)

In-Memory

Automatic In-Memory Enhancements for Improving Column Store Performance

Automatic In-Memory (AIM) has been enhanced to automatically enable creation and removal of Database In-Memory performance features based on an enhanced workload analysis algorithm. These features include Join Groups, caching of hashed dictionary values for join key columns and In-Memory Optimized Arithmetic.

Automatic In-Memory (AIM) has been enhanced to identify and enable or disable Database In-Memory features that can improve performance. It enables features either selectively or globally, depending on which adds the most benefit. This improves application performance and also conserves space in the In-Memory column store without requiring manual intervention.

Related Resources

[View Documentation](#)

Automatic In-Memory Sizing for Autonomous Databases

The In-Memory column store will now automatically grow and shrink dynamically based on workload. This allows the In-Memory column store to be available on Autonomous Database. Exadata scan performance is further improved for objects that are partially populated.

With Automatic In-Memory sizing, there is no longer a need to manually resize the In-Memory column store to accommodate different database workloads. This reduces the administrative effort of enabling Database In-Memory. Automatic In-Memory sizing also allows the In-Memory column store to be enabled on Autonomous Database (ADB), enabling applications running on ADB to also take advantage of faster analytic query performance.

Related Resources

[View Documentation](#)

In-Memory RAC-Level Global Dictionary

Database In-Memory Join Groups now support Global Dictionaries across RAC nodes. With Join Groups a common dictionary is shared by columns that are joined together. In-Memory RAC-level global dictionaries now synchronize these common dictionaries across nodes within a RAC database.

In a RAC environment, this feature further improves database In-Memory performance for distributed hash joins.

Related Resources

[View Documentation](#)

Vectorized Query Processing: Multi-Level Joins and Aggregations

This feature enhances the In-Memory Deep Vectorization framework by fully exploiting SIMD capabilities to further improve hash join and group by aggregation performance. New optimizations include incorporating multi-level hash join support, full In-Memory group by aggregation support, and support for multi-join key and additional join methods.

This feature adds improvements in the performance of joins and aggregations, which are the foundations for analytic queries. This enables faster real-time analytic performance and requires no application SQL changes. This feature is automatically used when enabled, which is the default.

Related Resources

[View Documentation](#)

Machine Learning - Enhancements

Automated Time Series Model Search

This feature enables the Exponential Smoothing algorithm to select the forecasting model type automatically - as well as related hyperparameters - when you do not specify the `EXSM_MODEL` setting. This can lead to more accurate forecasting models.

This feature automates the Exponential Smoothing algorithm hyperparameter search to produce better forecasting models without manual or exhaustive search. It enables non-expert users to perform time series forecasting without detailed understanding of algorithm hyperparameters while also increasing data scientist productivity.

Related Resources

[View Documentation](#)

Explicit Semantic Analysis Support for Dense Projection with Embeddings

The unstructured text analytics algorithm Explicit Semantic Analysis (ESA) is able to output dense projections with embeddings, which are functionally equivalent to the popular doc2vec (document to vector) representation.

Producing a doc2vec representation is useful as input to other machine learning techniques, for example, classification and regression, to improve their accuracy when used solely with text or in combination with other structured data. Use cases include processing unstructured text from call center representative notes on customers or physician notes on patients along with other customer or patient structured data to improve prediction outcomes.

Related Resources

[View Documentation](#)

GLM Link Functions

The in-database Generalized Linear Model (GLM) algorithm now supports additional link functions for logistic regression: probit, cloglog, and cauchit.

These additional link functions expand the set available to match standard Generalized Linear Model (GLM) implementations. They enable increasing model quality, for example, accuracy, by handling a broader range of target column data distributions and expand the class of data sets handled. Specifically, the probit link function supports binary (for example, yes/no) target variables, such as when predicting win/lose, churn/no-churn, buy/no-buy. The asymmetric link function complementary-log-log (cloglog) supports binary target variables where one outcome is relatively rare, such as when predicting time-to-relapse of medical conditions. The cauchit link function supports handling data with, for example, data recording errors, more robustly.

Related Resources

[View Documentation](#)

Improved Data Prep for High Cardinality Categorical Features

This feature introduces the setting `ODMS_EXPLOSION_MIN_SUPP` to allow more efficient, data-driven encoding for high cardinality categorical columns. You can adjust the threshold (define the minimum support required) for the categorical values in explosion mapping or disable the feature, as needed.

This feature introduces a more efficient, data-driven encoding of high cardinality categorical columns, allowing users to build models without manual data preparation of such columns.

It efficiently addresses large datasets with millions of categorical values by recoding categorical values to include only those with sufficient support, enabling you to overcome memory limitations.

Related Resources

[View Documentation](#)

Lineage: Data Query Persisted with Model

This feature enables users to identify the data query that was used to provide the training data for producing a model. The `BUILD_SOURCE` column in the `ALL/DBA/USER_MINING_MODELS` view enables users to access the data query used to produce the model.

This feature records the query string that is run to specify the build data, within the model's metadata to better support the machine learning lifecycle and MLOps.

Related Resources

[View Documentation](#)

Multiple Time Series

The Multiple Time Series feature of the Exponential Smoothing algorithm enables conveniently constructing Time Series Regression models, which can include multivariate time series inputs and indicator data like holidays and promotion flags. It enables constructing

Time Series Regression models to include multivariate time series inputs and indicator data like holidays and promotion flags.

This feature automates much of what a data scientist would perform manually by generating backcasts and forecasts on one or more input time series, where the target time series also receives confidence bounds. The result is used as input to other ML algorithms, for example, to support time series regression using XGBoost, with multivariate categorical, numeric, and time series variables.

Related Resources

[View Documentation](#)

Outlier Detection using Expectation Maximization (EM) Clustering

The Expectation Maximization algorithm is expanded to support distribution-based anomaly detection. The probability of anomaly is used to classify an object as normal or anomalous. The EM algorithm estimates the probability density of a data record, which is mapped to a probability of an anomaly.

Using Expectation Maximization (EM) for anomaly detection expands the set of algorithms available to support anomaly detection use cases, like fraud detection. Since different algorithms are capable of identifying patterns in data differently, having multiple algorithms available is beneficial when addressing machine learning use cases.

Related Resources

[View Documentation](#)

Partitioned Model Performance Improvement

This feature improves the performance for a high number of partitions (up to 32K component models) in a partitioned model and speeds up the dropping of individual models within a partitioned model.

Machine learning use cases often require building one model per subset of data, e.g., a model per state, region, customer, or piece of equipment. The partitioned models capability already automated the building of such models - providing a single model abstraction for simplified scoring - and this enhancement improves overall performance when using larger number of partitions.

Related Resources

[View Documentation](#)

XGBoost Support for Constraints and for Survival Analysis

The in-database XGBoost algorithm expands the in-database XGBoost algorithm support for the machine learning technique survival analysis. Also supported are feature interaction constraints and monotonic constraints, which allow choosing how variables are allowed to interact.

Survival analysis is an important machine learning technique for multiple industries. This enhancement enables increased model accuracy when predicting, for example, equipment failures and healthcare outcomes. Specifically, this supports data scientists

with the Accelerated Failure Time (AFT) model - one of the most used models in survival analysis – to complement the Cox proportional hazards regression model.

Interaction and monotonic constraints provide for greater control over the features used to achieve better predictive accuracy by leveraging user domain knowledge when specifying interaction terms.

Related Resources

[View Documentation](#)

Spatial

Spatial: 3D Models and Analytics

The point cloud feature of Oracle Database supports change detection through SQL and PL/SQL APIs.

This feature automates discovery of relevant changes between two point clouds, enabling easy inclusion in applications, such as modeling changes in forest canopies, assessing damages to landscape due to fire, flood, landslides, or earthquakes, and measuring progress over time in infrastructure projects.

Related Resources

[View Documentation](#)

Spatial: REST APIs for GeoRaster

Oracle Database includes a comprehensive set of REST APIs for working with GeoRaster data such as satellite imagery.

In addition to existing PL/SQL and Java APIs, developers can use REST APIs to perform GeoRaster query and data manipulation operations. This feature simplifies the development of cloud applications which frequently depend on REST APIs.

Related Resources

[View Documentation](#)

Workspace Manager: Improved Security when using Oracle Workspace Manager

Database users can have workspace manager objects in their own schema instead of in the WMSYS schema.

Oracle Workspace Manager enables collaborative development, what-if analysis from data updates, and maintains a history of changes to the data. Developers can create multiple workspaces and group different versions of table row values in different workspaces. With this enhancement, developers have improved security when using this feature. All the workspace manager objects can be stored and invoked in their own user schema in Oracle Autonomous Database and user-managed databases.

Related Resources

[View Documentation](#)

4

Data Warehousing/Big Data

General

Enhanced Partitioning Metadata

All data dictionary views holding partitioning-related metadata, for example, `[ALL_TAB_PARTITIONS]`, have two additional columns representing the high value (boundary) information of partitions and subpartitions in JSON and CLOB format.

Providing the high value (boundary) partitioning information in JSON and as CLOB allows you to use this information programmatically. This enables simple and automated processing of this information for schema retrieval or lifecycle management operations.

Related Resources

[View Documentation](#)

Extended Language Support in Oracle Text

Language support is extended in Oracle Text, now supporting up to 48 languages. Additionally, there is extended support for all languages. To avoid the extended language support increasing your install footprint on disk, a new mechanism is introduced to control the downloaded languages on demand.

Customers benefit from the improved support for languages and can download the required files for only the languages they support. They can avoid using disk space unnecessarily for unneeded languages.

Related Resources

[View Documentation](#)

External Table Partition Values in the File Path

External tables pointing to data in the object store can consist of a large number of files. These files can be organized across multiple directories, and even directory trees. You can use external table partitioning with folder names as part of the file paths. External table columns can also now return the file name of the source file for each row. The partition values can be derived from the directory name or file name.

External tables and SQL*Loader can load large numbers of data files in object stores and meet the requirements for Hive-generated tables organized across multiple directories, and even directory trees. This feature enables external table partitioning based upon the directory and file name; for example, you can have files for different months or for different states in separate directories.

Related Resources

[View Documentation](#)

Logical Partition Change Tracking for Materialized View Refresh and Staleness Tracking

Logical Partition Change Tracking (LPCT) tracks the staleness of materialized views. LPCT operates at a fine level of logical granularity and gives you the flexibility to align the boundaries of logical partitions with the business rules and with changes applied to tables. It evaluates the staleness of the base tables for individual logical partitions without using a materialized view log or requiring any of the tables used in the materialized view to be partitioned.

With LPCT, materialized view staleness is tracked at the granularity of the logical partitions. This functionality significantly broadens the applicability of query rewrite for your application due to the fine-grained query rewrite. With LPCT, you perform refresh operations targeted at stale logical partitions only, improving the refresh time and avoiding complete re-loading data.

Related Resources

[View Documentation](#)

Staging Tables

Staging tables are heap tables optimized for fast data ingestion and for handling volatile data. Key table attributes are set to defaults for these use cases without any additional user interaction.

Creating staging tables rather than 'normal' tables saves you time and effort so that you do not need to tune your table attributes for fast data ingestion with volatile data content. A staging table is configured by default with optimal configuration settings in order to guarantee the best possible performance and to avoid unnecessary performance debugging and tuning.

Related Resources

[View Documentation](#)

5

Cloud Migration

General

Classic Queue to Transactional Event Queues (TxEventQ) Online Migration Tool

Advanced Queuing (AQ) is a key feature of the Oracle Database used to build application workflows using events. Queues in the Oracle Database come in two flavors - Classic Queues and the new and improved Transactional Event Queues (TxEventQ). TxEventQ are Kafka-like and highly performant due to their persistent in-memory caching implementation.

Oracle Database 23c introduces an online migration tool to migrate from AQ Classic Queues to the new TxEventQ. The database creates a new TxEventQ with aliased name and routes enqueues to it, while it dequeues drain messages from the existing AQ. After the messages are drained, the migration is complete and the TxEventQ takes the name of the AQ queue. The TxEventQ migration interface checks for compatibility and adaptability for migration of an AQ deployment, and can commit or roll back a migration after running some tests on TxEventQ.

Existing AQ customers interested in higher throughput queues and with Kafka compatibility using a Kafka Java Client and Confluent like REST APIs, can migrate easily from AQ to Transactional Event Queues (TxEventQ).

TxEventQ is Oracle's next generation messaging system that offers many benefits and features over Advanced Queuing (AQ), like scalability, performance, key-based partitioning, and Kafka compatibility with a Java client, TxEventQ natively supports JSON payload, which makes event-driven microservices/application writing easier in multiple languages - Java, JavaScript, PL/SQL, Python, etc.

Related Resources

[View Documentation](#)

Cloud Premigration Advisor Tool for Source Database Migrations to Other Oracle Databases

The Cloud Premigration Advisor Tool (CPAT) assists you with database migration, both on-premises and in Oracle Cloud. It assesses the characteristics of the source database for migration and determines whether the source database can be migrated successfully to another Oracle Database. The JSON output from CPAT can be read by tools and applications for further processing and reporting.

CPAT helps you to avoid incompatibility issues in migrations. It helps you save time and effort in planning your migration of on-premises and Oracle Cloud databases to Oracle Autonomous Database.

Related Resources

[View Documentation](#)

6

Cloud Operations

Manageability

Hybrid Read-Only Mode for Pluggable Databases

Administrators can configure pluggable databases (PDBs) to operate in a new mode called hybrid read-only. Hybrid read-only mode enables a PDB to operate as either read-write or read-only, depending on the user who is connected to the PDB. For common users, the PDB will be in read-write mode. For local users, the PDB will be restricted to read-only mode.

Hybrid read-only mode enables you to patch and maintain an application in a safe mode for open PDBs without the risk of local users, including higher privileged ones, interfering with the ongoing maintenance operation of the PDB.

Related Resources

[View Documentation](#)

Real-Time SQL Monitoring Enhancements

Real-time SQL Monitoring works independently and concurrently across multiple PDB containers in an efficient manner. SQL statements, PL/SQL procedures and functions, and DBOPs (Database Operations) are monitored at PDB and CDB levels. You can efficiently query SQL Monitor reports across ad-hoc time ranges, DBIDs (internal database identifiers), and CON_DBIDs (CDB identifiers). This data is also accessible through SQL History Reporting.

Additionally, SQL Monitoring data can be exported along with the Automatic Workload Repository (AWR) and imported into another database or container for longer term storage and analysis.

Real-time SQL Monitoring is now supported per-PDB and CDB levels efficiently by default. As a PDBA persona, you can get a more accurate view of the monitored SQL for your application.

SQL Monitoring data can be transported through the AWR framework to a different container or database for longer term storage and offline analysis.

Related Resources

[View Documentation](#)

Control PDB Open Order

Administrators can define a startup order or priority for each pluggable database (PDB) where the most important PDBs are started first. The priority is applied to PDB opening order and upgrade order as follows:

- Restoring PDB states when opening the CDB
- Setting PDB states when using the `PDB OPEN ALL` statement
- Setting the order for PDB database upgrade operations
- Starting PDBs in an Active Data Guard (ADG) switchover or failover

This feature allows critical PDBs to start and open before less important PDBs, reducing the time for the critical applications to become usable.

Related Resources

[View Documentation](#)

ODP.NET: Centralized Config Providers

Managed ODP.NET and ODP.NET Core securely retrieve application configuration data from a centralized on-premises, Oracle Cloud Infrastructure, or Azure location. This store uses a JSON format and contains data, including connection descriptors and tuning parameters.

Centralized application configuration makes management and cloning simpler. It suits modern architectures, such as microservices and serverless deployments.

Related Resources

[View Documentation](#)

Optimized Performance for Parallel File System Operations

This feature allows you to tune performance in environments that contain many PDBs and require multiple `DBMS_FS` requests to be processed in parallel. You can update the number `OFS_THREADS` to increase the number of `DBMS_FS` requests that are executed in parallel. This increases the number of worker threads executing the make, mount, unmount, and destroy operations on Oracle file systems in the Oracle database. An increase in the value of `OFS_THREADS` results in a significant reduction of time taken to execute parallel file system requests in environments that contain multiple PDBs.

This feature significantly reduces the time required to perform parallel file system requests in consolidation environments containing multiple PDBs.

Related Resources

[View Documentation](#)

Read-Only Users and Sessions

You can control whether a user or session is enabled for read-write operations, irrespective of the privileges of the user that is connected to the database. The `READ_ONLY` session applies to any type of user for any type of container. The `READ_ONLY` user only applies to local users.

Providing the capability to disable and re-enable the read-write capabilities of any user or session without revoking and re-granting privileges provides you with more flexibility to temporarily control the privileges of users or sessions for testing, administration, or application development purposes. It also gives you a simple way to control the read-

write behavior within different parts of an application that are used by the same user or session.

Related Resources

[View Documentation](#)

Continuous Availability

Enhanced Upgrade of Time Zone Data

The process of upgrading timezone data to reflect up-to-date Governmental Daylight Saving Time rules is optimized, taking the actual data content of tables into account. Only tables impacted by a Daylight Saving Time rule change will undergo a data change.

Optimizing the necessary data changes for a Daylight Savings Time rule change significantly improves the overall upgrade of timezone data to the absolute bare minimum to bring a database up to the latest global timezone rules. The implicit analysis and reduction of the data required to change significantly reduces the overall timezone upgrade process and the resources needed.

Related Resources

[View Documentation](#)

Optimized Read-Write Operations for Database Processes

To optimize the read and write operations performed by database processes when you access files managed through OFS or DBFS, specify the new `db_access` mount option for the `dbms_fs.mount_oracle_fs` procedure while mounting the file system.

When you enable `db_access`, both memory consumption and CPU usage reduces. The throughput increases while performing read and write operations by database processes on the files managed by OFS.

Related Resources

[View Documentation](#)

Support for the Coexistence of DGPDB and GoldenGate Capture

This project introduces perPDB DataGuard. When DGPDB is configured on a source/Primary database, there are validations that insure there is no GoldenGate Capture pre-existing on the source. GoldenGate capture sessions will be broken if a DGPDB is allowed and executes a role transition.

This project adds support for coexistence of DGPDB and GoldenGate Capture. Changes/support will be required in the LogMiner, redo transport, and Broker layers.

Related Resources

[View Documentation](#)

General

Adaptive Result Cache Object Exclusion

With adaptive result cache object exclusion, the database decides to blocklist certain objects if using the result cache is not beneficial for these objects, based on statistical evidence such as the number of invalidations, the cost savings of using result caching, and others. You have full control over the objects considered for exclusion to ensure you can continue using result cache for all your objects of interest.

Adaptive exclusion of objects that don't benefit or even have a detriment impact on result caching reduces the overall development and management workload for you. It can improve the database performance out of the box.

Related Resources

[View Documentation](#)

Diagnose and Repair SQL Exceptions Automatically at Compile-Time

SQL diagnostics can automatically detect and repair many severe compile-time SQL exceptions that would otherwise cause SQL statements to fail.

This feature improves the robustness of your applications and its service levels.

Related Resources

[View Documentation](#)

Read-Only Tablespace on Object Storage

Read-only tablespaces can be moved to and from Object Storage transparently, storing portions of a database on lower-cost storage in the Cloud.

Allowing to move tablespaces to Object Storage enables a data lifecycle management strategy, storing the data on the most cost-effective storage tier based on its business value or access frequency.

Related Resources

[View Documentation](#)

SYSDATE and SYSTIMESTAMP Reflect PDB Time Zone

Oracle Multitenant enables an Oracle Database to consolidate multiple pluggable databases as self-contained databases, improving resource utilization and database management. In addition to the centrally managed `SYSDATE` and `SYSTIMESTAMP` following the OS-level setting of a database host, `SYSDATE` and `SYSTIMESTAMP` can be managed per pluggable database following the individual database time zone setting. All user-visible operations adhere to this setting.

The ability to control the time zone behavior for `SYSDATE` and `SYSTIMESTAMP` on a pluggable database level increases to self-containment of individual databases in a

multitenant environment and enhances your consolidation capabilities of independent databases.

Related Resources

[View Documentation](#)

Unified Memory

Unified Memory is a flexible and simple memory configuration for Oracle Databases that uses a single parameter to control database memory allocations, reducing or eliminating the need for system restart to change memory configurations. Unified Memory is especially useful in multiple workload high density database consolidation environments.

Unified Memory simplifies memory management to run multiple workloads in a highly consolidated environment with minimum disruption.

Related Resources

[View Documentation](#)

7

High Availability

Data Guard

Oracle Data Guard Redo Decryption for Hybrid Disaster Recovery Configurations

Oracle Data Guard now provides the capability to decrypt redo operations in hybrid cloud disaster recovery configurations where the cloud database is encrypted with Transparent Data Encryption (TDE) and the on-premises database is not.

Hybrid disaster recovery (DR) with Data Guard is now more flexible and easy to configure. Hybrid disaster recovery for the Oracle Database allows you to expand outage and data protection to take advantage of the automation and resources of Oracle Cloud Infrastructure (OCI). By enabling the ability to quickly configure disaster recovery in OCI, even in cases where on-premises databases might not already be encrypted with Transparent Data Encryption (TDE), the steps required to configure hybrid disaster recovery environments and prepare on-premises databases for a DR configuration with cloud databases in OCI have been greatly reduced.

Related Resources

[View Documentation](#)

General

Flashback Time Travel Enhancements

Flashback Time Travel can automatically track and archive transactional changes to tables. Flashback Time Travel creates archives of the changes made to the rows of a table and stores the changes in history tables. It also maintains a history of the evolution of the table's schema. By maintaining the history of the transactional changes to a table and its schema, Flashback Time Travel enables you to perform operations, such as Flashback Query (`AS OF` and `VERSIONS`), on the table to view the history of the changes made during transaction time.

Flashback Time Travel helps to meet compliance requirements based on record-stage policies and audit reports by tracking and storing transactional changes to a table, which has also been made more efficient and performant in this release.

Related Resources

[View Documentation](#)

Minimized Stall and More Accurate Lag Detection in Fast-Start Failover Maximum Performance

Oracle Data Guard Fast-Start Failover has two additional properties for improved lag detection and status changes. `FastStartFailoverLagType` sets the lag type that Fast-Start Failover must consider when in Maximum Performance mode (`APPLY` or `TRANSPORT`). `FastStartFailoverLagGraceTime` lets the configuration transition to a pre-emptive `LAGGING` state that the observer can acknowledge before reaching the actual lag limit, so the status can transition immediately to `TARGET OVER LAG LIMIT` without waiting for the observer quorum.

The new properties for the Maximum Performance protection mode further enhance Fast-Start Failover capabilities and reduce the impact on application transactions for status changes requiring the observer quorum.

Related Resources

[View Documentation](#)

Transaction Guard Support during DBMS_ROLLING

Transaction Guard support for `DBMS_ROLLING` ensures continuous application operation during the switchover issued by `DBMS_ROLLING` to Transient Logical Standby. The procedure uses the last commit outcome of transactions part of in-flight sessions during a switchover-related outage (or caused by an error/timeout) to protect the applications from duplicate submissions of the transactions on replay.

Application Continuity supported by Transaction Guard during database upgrades using `DBMS_ROLLING` ensures that commit outcomes are guaranteed across the entire upgrade process.

Related Resources

[View Documentation](#)

8

Security

SQL Firewall

Oracle SQL Firewall Included in Oracle Database

Included in Oracle Database, SQL Firewall inspects all incoming SQL statements and ensures that only explicitly authorized SQL is run. You can use SQL Firewall to control which SQL statements are allowed to be processed by the database. You can restrict connection paths associated with database connections and SQL statements. Unauthorized SQL can be logged and blocked. Because SQL Firewall is embedded in the Oracle database, it cannot be bypassed. All SQL statements are inspected, whether local or network, or encrypted or clear text. It examines top-level SQL, stored procedures and the related database objects.

SQL Firewall provides real-time protection against common database attacks by restricting database access to only authorized SQL statements or connections. It mitigates risks from SQL injection attacks, anomalous access, and credential theft or abuse.

SQL Firewall uses session context data such as IP address, operating system user name, and operating system program name to restrict how a database account can connect to the database. This helps mitigate the risk of stolen or misused application service account credentials. A typical use case for SQL Firewall is for application workloads.

You can use SQL Firewall in both the root and a pluggable database (PDB).

Related Resources

[View Documentation](#)

Encryption

Transport Layer Security (TLS) 1.3 Now Supported in Oracle Database

Transport Layer Security (TLS) version 1.3 is supported in Database 23c. TLS 1.3 is the latest and most secure TLS protocol to protect network connections to and from an Oracle database.

Because TLS 1.3 handles initial session setup more efficiently than prior TLS versions, users moving to TLS 1.3 should see improvements in TLS performance, particularly for applications that frequently connect and reconnect to the database. TLS 1.3 also implements newer, more secure cipher suites that improve confidentiality of data in transit.

Related Resources

[View Documentation](#)

New Parameter to Control the Behavior of TLS Certificate DN Match

You can use the `SSL_ALLOW_WEAK_DN_MATCH` parameter to control how `SSL_SERVER_DN_MATCH` allows the service name for partial distinguished name matching and to only check the database server certificate.

DN matching with both the listener and server certificates provides better security to ensure that the client is connecting to the correct database server.

Related Resources

[View Documentation](#)

Simplified Transport Layer Security Configuration

The Transport Layer Security (TLS) configuration between the database client and server has been simplified with streamlined parameters, performance improvements, and an additional parameter to find a wallet. Older TLS protocols have also been removed.

These changes improve security and make it easier to implement TLS.

Related Resources

[View Documentation](#)

Ability to Configure Transport Layer Security Connections Without Client Wallets

An Oracle Database client is no longer required to provide a wallet to hold well-known CA root certificates if they are available elsewhere in the local system.

Transport Layer Security (TLS) encryption requires either one-way authentication or two-way authentication. In one-way authentication (the default), which is commonly used for HTTPS connections, the server certificate is verified using well-known root CA certificates that are already available in local systems. Starting in this release, you will no longer need to install and configure a wallet to hold a well-known root certificate if it is already available in the local system.

This feature greatly simplifies the Oracle Database client installation and the use of TLS protocol to encrypt Oracle Database client-server communications.

Related Resources

[View Documentation](#)

New sqlnet.ora Parameter to Prevent the Use of Deprecated Cipher Suites

You can block the use of deprecated cipher suites by setting the `SSL_ENABLE_WEAK_CIPHERS` `sqlnet.ora` parameter to `FALSE`.

Removing the ability to use older, less secure cipher suites improves protection for data in-motion between the database.

Related Resources[View Documentation](#)

AES-XTS Encryption Mode Support for TDE Tablespace Encryption

Transparent Database Encryption (TDE) tablespace encryption now supports Advanced Encryption Standard (AES) XTS (XEX-based mode with ciphertext stealing mode) in `CREATE TABLESPACE` statements. Earlier versions of Oracle Database TDE used AES-CFB cipher mode.

AES-XTS provides improved security and better performance, especially on platforms where TDE can take advantage of parallel processing and specialized instructions built into processor hardware.

Related Resources[View Documentation](#)

Changes for TDE Encryption Algorithms and Modes

The default encryption algorithm for both TDE column encryption and TDE tablespace encryption is now AES256. The previous default for TDE column encryption was AES192. For TDE tablespace encryption, the default was AES128.

The decryption libraries for the GOST and SEED algorithms are deprecated. New keys cannot use these algorithms. The encryption libraries for both of these libraries are desupported.

The column encryption mode is now Galois/Counter mode (GCM) instead of cipher block chaining (CBC), and the tablespace keys are now used in tweakable block ciphertext stealing (XTS) operating mode instead of cipher feedback (CFB).

The Oracle Recovery Manager (RMAN) integrity check for column encryption keys now uses SHA512 instead of SHA1.

The keys for Oracle RMAN and column keys are now derived from SHA512/AES for key generation. In previous releases, they used SHA-1/3DES as a pseudo-random function.

These enhancements enable your Oracle Database environment to use the latest, most secure algorithms and encryption modes.

Related Resources[View Documentation](#)

Improved and More Secure Local Auto-Login Wallets

A local auto-login wallet is now more tightly bound to the host where it was created or modified (both bare metal and virtual). The local auto-login process is also more secure, does not require additional deployment requirements, and does not require root access.

Local auto-login wallets are more secure now and support both bare metal and virtual environments.

This enhancement also applies to Transparent Data Encryption (TDE) auto-login keystores.

Related Resources

[View Documentation](#)

Audit

Audit Object Actions at the Column Level for Tables and Views

You can create unified audit policies to audit individual columns in tables and views.

This feature enables you to configure more granular and focused audit policies, and ensures that auditing is selective enough to reduce the creation of unnecessary audit records, and effective enough to let you meet your compliance requirements.

Related Resources

[View Documentation](#)

Control Authorizations for Unified Auditing and Traditional Auditing

You can control how privileged users can grant and revoke the Oracle Database `AUDIT_ADMIN` and `AUDIT_VIEWER` roles by using Oracle Database Vault APIs. Database Vault blocks direct modification of the database audit tables except through the `DBMS_AUDIT_MGMT` PL/SQL package by authorized users. A new mandatory default realm (Oracle Audit Realm) protects the `AUDSYS` schema and audit-related objects in the `SYS` schema.

This new Database Vault realms simplifies auditing database vault, consolidating the privileges required for auditing into one authorization mechanism. In addition to facilitating the granting of audit-related privileges to the user, this enhancement provides greater separation of duties for managing auditing in an Oracle Database Vault environment.

Related Resources

[View Documentation](#)

Authentication

Microsoft Azure Active Directory Integration with Additional Oracle Database Environments

You can log into Oracle Databases using your Microsoft Azure Active Directory (Azure AD) single sign-on OAuth2 access token. This feature has been backported to Oracle Database release 19.16 and later, but not for Oracle Database 21c.

This multi-cloud feature integrates authentication and authorization between Azure AD and Oracle Databases.

Related Resources

[View Documentation](#)

ODP.NET: Azure Active Directory Single Sign-On

ODP.NET can log into Oracle databases using a Microsoft Azure Active Directory (Azure AD) OAuth 2.0 access token. Users can sign-on once with Azure AD, acquire the token, and access their on-premises and cloud-based Oracle databases. This feature is available in ODP.NET Core and managed ODP.NET.

This multicloud capability eases authentication and authorization between Azure AD and Oracle Databases by simplifying user access and management.

Related Resources

[View Documentation](#)

Increased Oracle Database Password Length

Oracle Database now supports passwords up to 1024 bytes in length. In previous releases, the Oracle Database password length and the secure role password length could be up to 30 bytes.

Increasing the password length supports an industry-wide trend for stronger authentication. In cases where passwords must be used, the increased length permits passwords that are more difficult to guess.

Related Resources

[View Documentation](#)

JDBC-Thin Support for Longer Passwords

Passwords for database user authentication can now be as long as 1024 characters.

This feature fosters increased authentication security for Java applications in Cloud and On-premises environments.

Related Resources

[View Documentation](#)

Oracle Data Pump Export and Import Support for Longer Encryption Passwords

Oracle Data Pump can protect export files with encryption passwords of up to 1024 bytes long.

Oracle Data Pump enhances security by supporting encryption passwords of up to 1024 bytes long.

Related Resources

[View Documentation](#)

Oracle Call Interface (OCI) and Oracle C++ Call Interface (OCCI) Password Length Increase

Oracle Call Interface (OCI) and Oracle C++ Call Interface (OCCI) now support passwords for database user authentication up to 1024 bytes long.

This feature allows longer passwords to be used to improve security. It also aids database use with tools that generate long passwords.

Related Resources

[View Documentation](#)

Updated Kerberos Library and Other Improvements

Oracle Database supports MIT Kerberos library version 1.20.1, and provides other improvements, such as support for multiple principals, cross-domain, and Windows Credential Guard.

These Kerberos improvements improve security and allow Kerberos to be used in more Oracle Database environments.

Related Resources

[View Documentation](#)

Enhancements to RADIUS Configuration

RADIUS is frequently used to provide multi-factor authentication (MFA) for Oracle Database. Oracle Database 23c now supports the RFC 6613 and 6614 guidelines for RADIUS and implements TCP over Transport Layer Security (TLS) by default. This enhancement introduces new RADIUS-related `sqlnet.ora` parameters to support the new standards. The enhancement also deprecates several RADIUS-related `sqlnet.ora` parameters that are no longer needed to support the new standards.

This update to RADIUS standards support improves security for customers using RADIUS-based authentication.

Related Resources

[View Documentation](#)

UTL_HTTP Support for SHA-256 and Other Digest Authentication Standards

UTL_HTTP is extended to support both SHA-256 and SHA-512/256 for digest authentication, to ensure forward compatibility.

UTL_HTTP can be seen as an API for client-side HTTP access, much like a standard browser. Support for both SHA-256 and SHA-512/256 for digest authentication enables UTL_HTTP to be at par with other standard browsers.

Related Resources

[View Documentation](#)

XDB HTTP SHA512 Digest Authentication

Oracle XDB HTTP protocol server now supports digest authentication SHA512 authentication, which is a more secure digest algorithm than MD5.

This feature improves security when using Oracle XDB from the web.

Related Resources

[View Documentation](#)

Authorization

Schema Privileges to Simplify Access Control

Oracle Database supports granting privileges on schemas (in addition to the existing object, system, and administrative privileges).

This feature improves security by simplifying authorization for database objects, especially for schemas that frequently add new objects. Instead of granting broad system level (* ANY) privileges that apply to the entire database, privileges can now be granted at the individual schema level.

Related Resources

[View Documentation](#)

Oracle Label Security Triggers Are Now Part of the New LBAC_TRIGGER Schema

A new schema, `LBAC_TRIGGER`, is introduced to own the internal triggers that were previously owned by the `LBACSYS` schema. You can migrate existing `LBACSYS` triggers to this new schema.

Both the `LBACSYS` and `LBAC_TRIGGER` schemas are Oracle-maintained and dictionary-protected.

This feature improves security when using the Oracle Label Security option.

Related Resources

[View Documentation](#)

Oracle Data Dictionary Protection Extended to Non-SYS Oracle Schemas with Separation of Duties Protection

Oracle Database schemas can have data dictionary protection with additional separation of duties protection for `SYSOPER`, `SYSASM`, `SYSBACKUP`, `SYSKM`, `SYSRAC`, and `SYSDG`.

Oracle schemas provide critical functionality for Oracle Database features. By enabling these schemas to have data dictionary protection with additional separation of duties, you can prevent inadvertent and malicious changes within these schemas that could endanger Oracle Database functionality.

Related Resources

[View Documentation](#)

GoldenGate Capture and Apply User Roles

New roles `OGG_CAPTURE`, `OGG_APPLY`, `OGG_APPLY_PROCREP`, `XSTREAM_CAPTURE`, `XSTREAM_APPLY` have been created for granting appropriate capture and apply privileges to the GoldenGate and XStream administrators. These new roles replace the functionality in the procedures of the `DBMS_GOLDENGATE_AUTH` and `DBMS_XSTREAM_AUTH` packages, which are now de-supported.

This feature simplifies administrative tasks.

Related Resources

[View Documentation](#)

Autonomous Database

Identity and Access Management Integration with Oracle Autonomous Cloud Databases

You can now log in to additional Oracle Database Oracle Cloud Infrastructure (OCI) DBaaS platforms by using an Identity and Access Management (IAM) password or a token-based authentication. It's possible to log in to these databases by using these IAM credentials from tools, such as SQL*Plus or SQLcl.

This feature improves security through centralized management of credentials for OCI DBaaS database instances.

Related Resources

[View Documentation](#)

ODP.NET: Oracle Identity and Access Management

ODP.NET supports Oracle Identity and Access Management (IAM) cloud service for unified identity across Oracle cloud services, including Oracle Cloud Database Services. ODP.NET can use the same Oracle IAM credentials for authentication and authorization to the Oracle Cloud and Oracle cloud databases, now with IAM SSO tokens. This feature is available in ODP.NET Core and managed ODP.NET.

This capability allows single sign-on and for identity to be propagated to all services Oracle IAM supports including federated users via Azure Active Directory and Microsoft Active Directory (on-premises). A unified identity makes user management and account management easier for administrators and end users.

Related Resources

[View Documentation](#)

Oracle Client Increased Database Password Length

Starting with this release, Oracle Database and client drivers support passwords up to 1024 bytes in length.

The Oracle Database and client password length has been increased to 1024 bytes, up from 30 bytes, to allow users to set longer passwords if needed. The maximum number of characters is based on the character set used since some characters are larger than one byte.

Related Resources

[View Documentation](#)

Other

Secure Distributed Transaction Recovery Background Process (RECO)

Oracle Database enables queries and DMLs on objects hosted on a different database. When objects are updated on a remote database, the transaction on the source database ends up becoming a distributed transaction. If a distributed transaction fails, a database background process (RECO) periodically tries to re-establish contact, with the yet-to-be-notified subordinates and pushes the final outcome to those remote databases.

Secure Distributed Transaction Recovery Background Process (RECO) provides additional security for the RECO process.

Related Resources

[View Documentation](#)

IP Rate Limit in CMAN

You can use Oracle Connection Manager (CMAN) to limit the number of new connections allowed from an IP address in the specified unit of time. This IP rate limit feature enables you to protect your database against potential denial-of-service (DoS) attacks.

Malicious clients can send excessive connection requests to the server node. This can saturate the capacity of CMAN to handle new connections per second, and thus cause DoS attacks on your database. Using this security feature, you can prevent these types of attacks by detecting such clients early and rejecting those connections.

Related Resources

[View Documentation](#)

OCI Attributes for Microsoft Azure Active Directory Integration with Additional Oracle Database Environments

You can log into additional Oracle Database environments using your Microsoft Azure Active Directory (Azure AD) single sign-on OAuth2 access token. The previous release supported Azure AD integration for Oracle Cloud Infrastructure (OCI) Autonomous Database (Shared Infrastructure). This release has expanded Azure AD integration to support on-premises

Oracle Database release 19.16 and later. The project adds the OCI attributes needed to supply the bearer token for connection creation.

This multi-cloud feature integrates authentication and authorization between Azure AD and Oracle Databases in Oracle Cloud Infrastructure and on-premises.

Related Resources

[View Documentation](#)

ODP.NET Transport Layer Security (TLS) 1.3

ODP.NET supports Transport Layer version 1.3 for the core, managed, and unmanaged providers. TLS is used to encrypt data and authenticate connections.

TLS 1.3 offers better security and performance over TLS 1.2.

Related Resources

[View Documentation](#)

9

OLTP and Core Database

Availability

Directory-Based Sharding Method

Directory-based sharding is a type of user-defined sharding in Oracle Globally Distributed Database where the location of data records associated with a sharding key is specified dynamically at the time of insert based on user preferences. The key location information is stored in a directory which can hold a large set of key values in the hundreds of thousands. With directory-based sharding, you have the freedom to move individual key values from one location to another, or make bulk movements to scale up or down, or for data and load balancing.

Directory-based Sharding Method improves the user-defined sharding model and provides linear scalability, complete fault isolation, and global data distribution for the most demanding applications.

Related Resources

[View Documentation](#)

Oracle Globally Distributed Database Raft Replication

Raft replication provides built-in replication for Oracle Globally Distributed Database without requiring configuration of Oracle GoldenGate or Oracle Data Guard. Raft replication is logical replication with consensus-based (RAFT) commit protocol, which enables declarative replication configuration and sub-second failover.

RAFT Replication helps simplify management, improves availability and SLA delivery, as well as optimizes hardware utilization for sharded database environments.

Related Resources

[View Documentation](#)

Automatic Data Move on Sharding Key Update

When you update the sharding key value on a particular row of a sharded table, the data with that key value might be mapped to a different partition or shard than where it currently resides. Oracle Globally Distributed Database now handles moving the data to the new location, whether it is in a different partition on the same shard or on a different shard.

This feature makes data movement between partition or shards seamless when sharding key value update occurs due to various reasons, for example, a move to another country or change in roles.

Related Resources

[View Documentation](#)

Automatic Transaction Quarantine

System MONitor (SMON) is a background process responsible for transaction recovery. Transaction Quarantine can now automatically quarantine the recovery of problematic transactions while keeping the database open, allowing SMON to proceed with the recovery of the other transactions. Alerts and diagnostic information are provided to the DBA or operator so that they can review and resolve the quarantine while other database operations continue unaffected.

The benefit of transaction quarantining is increased fault tolerance and high availability of the database. The database stays up and running and continues processing transactions while the quarantine is being resolved.

Related Resources

[View Documentation](#)

Automatic Transaction Rollback

If a transaction does not commit or rollback for a long time while holding row locks, it can potentially block other high-priority transactions. This feature allows applications to assign priorities to transactions and for administrators to set timeouts for each priority. The database will automatically rollback a lower priority transaction and release the row locks held if it blocks a higher priority transaction beyond the set timeout, allowing the higher priority transaction to proceed.

Automatic Transaction Rollback reduces the administrative burden while also helping to maintain transaction latencies and SLAs on higher priority transactions.

Related Resources

[View Documentation](#)

Creating Immutable Backups Using RMAN

RMAN is now compatible with immutable OCI Object Storage using locked retention rules, which prevents deletion or modification of backups.

To help organizations meet ransomware protection or strict regulatory requirements for record management and retention, RMAN now prevents anyone, even an administrator, from deleting or modifying backups in OCI Object Storage.

Related Resources

[View Documentation](#)

Fine-Grained Refresh Rate Control For Duplicated Tables

Oracle Globally Distributed Database enables refresh rate control for individual duplicated tables. Each duplicated table can have a separate refresh rate which is defined either at its creation or by the `ALTER TABLE` statement.

This feature helps optimize the use of resources by customization of refresh rates for individual duplicated tables.

Related Resources[View Documentation](#)

Global Partitioned Index Support on Subpartitions

Globally Distributed Database allows a global partitioned index on the sharding key when the sharded table is sub-partitioned. You can create a primary key/unique indexes on sharded tables that are composite partitioned without having to include sub-partition keys.

The benefit of this feature is that it removes the restriction on the primary key columns when the sharded table is sub-partitioned, as in the composite sharding method.

Related Resources[View Documentation](#)

JDBC Support for Split Partition Set

This feature enables the Java connection pool (UCP) to receive ONS events about data in a chunk being split and moved across partition sets, and then update the sharding topology appropriately.

This feature furnishes high availability to Java applications using Sharded databases.

Related Resources[View Documentation](#)

Managing Flashback Database Logs Outside the Fast Recovery Area

In previous releases, you could store flashback database logs only in the fast recovery area. Now you can optionally designate a separate location for flashback logging. For example, if you have write-intensive database workloads, then flashback database logging can slow down the database if the fast recovery area is not fast enough. In this scenario, you can now choose to write the flashback logs to faster disks. Using a separate destination also eliminates the manual administration to manage the free space in the fast recovery area.

Managing flashback database logs outside the fast recovery area lowers the operational costs related to space management and guarantees the best performance for workloads that are typically impacted by flashback logging on traditional storage.

Related Resources[View Documentation](#)

Move Data Chunks Between Shardspaces

With the Oracle Globally Distributed Database composite sharding method, data is organized into different shardspaces, allowing you to differentiate subsets of data; however, any automatic chunk moves for load balancing occurred within a shardspace. With this release, Oracle Globally Distributed Database provides you with the ability to move data chunks from an existing shardspace to another shardspace.

This feature makes it possible to move data between existing shardspaces, or move some data to a newly added shardspace. You can arrange sharded data for your new business

needs, such as providing new levels of services or resources for certain customers, or move customers from one class of service to another while maintaining regional data sovereignty.

Related Resources

[View Documentation](#)

New Duplicated Table Type - Synchronous Duplicated Table

Oracle Globally Distributed Database introduces a new kind of duplicated table that is synchronized on the shards 'on-commit' on the shard catalog. The rows in a duplicated table on the shards are synchronized with the rows in the duplicated table on the shard catalog when the active transaction performing DMLs on the duplicated tables in the shard catalog is committed.

This features enables efficient and absolute data consistency and synchronization for duplicated tables, across all shards at all times.

Related Resources

[View Documentation](#)

New Partition Set Operations for Composite Sharding

For Oracle Globally Distributed Database sharded databases using the composite sharding method, two new `ALTER TABLE` operations enhance partition set maintenance. Previously, partition set operations did not support specifying tablespace sets for child and reference-partitioned tables that are affected due to add and split partition set operations. `MOVE PARTITIONSET` lets you move a whole partition set from one tablespace set to another, within the same shardspace. `MODIFY PARTITIONSET` lets you add values to the list of values of a given partition set.

These new operations enhance resharding capability. `MOVE PARTITIONSET` gives you the control to move all subpartitions of a given table to another tablespace set, within a given shardspace. You can also specify separate tablespace sets for LOBs and subpartitions. `MODIFY PARTITIONSET` extends the add list values feature of partitions to partition sets.

Related Resources

[View Documentation](#)

Oracle Globally Distributed Database Coordinated Backup and Restore Enhancements

Coordinated backup and restore functionality in Oracle Globally Distributed Database has been extended to include the following:

- Enhanced error handling and diagnosis for backup jobs
- Improved automation of sharded database restore
- Support for running RMAN commands from GDSCTL
- Support for using different RMAN recovery catalogs for different shards
- Encryption of backup sets

- Support for additional backup destinations: Amazon S3, Oracle Object Storage, and ZDLRA

The benefits of this functionality are:

- Easily diagnose problems in backup jobs
- Backups sets can be encrypted so that the data is secured
- Support for additional destinations other than on-disk storage
- Support for different RMAN catalogs and destinations to abide by data residency requirements

Related Resources

[View Documentation](#)

Parallel Cross-Shard DML Support

The Oracle Globally Distributed Database query coordinator runs cross-shard updates and inserts in parallel on multiple shards.

This feature improves cross-shard DML performance by running updates and inserts in parallel rather than serially.

Related Resources

[View Documentation](#)

PL/SQL Function Cross-Shard Query Support

PL/SQL functions are enhanced with the keyword `SHARD_ENABLE` to allow these functions to be referenced in Oracle Globally Distributed Database cross-shard queries. With the new keyword, the query optimizer takes the initiative to push the execution of the PL/SQL function to the shards.

This feature significantly improves performance for PL/SQL functions in sharded database environments.

Related Resources

[View Documentation](#)

Pre-Deployment Diagnostic for Oracle Globally Distributed Database

While processing `GSDSCTL ADD SHARD`, `ADD GSM`, and `DEPLOY` commands, Oracle Globally Distributed Database runs a series of checks to make sure that there is no potential environmental issue.

This feature proactively avoids common pitfalls to reduce time taken to complete a sharded database deployment.

Related Resources

[View Documentation](#)

RMAN Backup Encryption Algorithm Now Defaults to AES256

RMAN encrypted backups now default to `AES256` encryption algorithm. RMAN will continue to support restore using existing backups created with `AES128` or `AES192` encryption algorithms. You may also choose to create new backups using `AES128` by changing the default `AES256` setting. This default change applies to `BACKUP BACKUPSET` command and the `ALLOCATE CHANNEL` command.

To strengthen the security of encrypted backups from being decrypted by malicious users, RMAN encrypted backups now default to the `AES256` encryption standard.

Related Resources

[View Documentation](#)

RMAN Operational, Diagnostics, and Upgrade Enhancements

RMAN now includes easier standby database registration for Oracle Data Guard, better fault tolerance and optimization for Oracle Real Application Cluster (Oracle RAC), enhanced diagnosability which automatically gathers information to help identify issues, and updates to mitigate bottlenecks and pause sessions during recovery catalog upgrades.

RMAN operations are now easier and more resilient for highly available Oracle environments with less complex backup registration, automatic diagnostic gathering, and fewer failures when performing maintenance activities.

Related Resources

[View Documentation](#)

Simplified Database Migration Across Platforms Using RMAN

Using RMAN to migrate databases across different operating system platforms has been streamlined and includes support for databases encrypted with Transparent Data Encryption (TDE) and multi-section backups. New command options allow existing RMAN backups to be used to transport tablespaces or pluggable databases to a new destination database with minimal downtime.

Migrations using RMAN are now easier, faster, and require fewer steps to execute. The new capabilities enable a simple and straightforward migration process, minimizing downtime for your applications, reducing risk, and increasing productivity.

Related Resources

[View Documentation](#)

Support for Oracle Database Version Specific RMAN SBT Library

The Oracle Home directory now includes the database version compatible libraries (`SBT_LIBRARY`) for Zero Data Loss Recovery Appliance, OCI Object Storage and Amazon S3. You can now configure RMAN to directly access libraries from the Oracle Home directory using an alias. For example, if the backup destination is OCI Object Storage, you only have to specify the alias name `oracle.oci` for the `SBT_LIBRARY` parameter. When RMAN attempts to backup to Object Storage, it uses the specified

alias to access the SBT library used for backup cloud service from the Oracle home directory.

The RMAN storage libraries are now included with the database, eliminating the need to download and install additional software and ensuring that you have all the necessary components to immediately start backing up and restoring from Zero Data Loss Recovery Appliance, OCI Object Storage, or Amazon S3.

Related Resources

[View Documentation](#)

Blockchain

Blockchain Table User Chains

Earlier versions of blockchain tables supported only system chains. A system chain (one of the 32 chains per instance) is randomly chosen by Oracle for every new row inserted into a blockchain table.

A user chain is a chain of rows based on a set of up to three user-defined columns of type `NUMBER`, `CHAR`, `VARCHAR2`, and `RAW`. For example, consider a blockchain table created for tracking banking transactions (withdrawals, deposits, transfers) associated with various accounts. Assume there is a column called `ACCOUNTNO` in the blockchain table for account numbers. Each transaction inserts a new entry into this blockchain table for some account number. A user chain can be associated with every unique value in `ACCOUNTNO`. If there are a total of 100 different account numbers, there can be at most 100 user chains. You can then run a verification procedure only on a chain for a specific `ACCOUNTNO`, providing greater data isolation. This feature allows you to create user chains for rows in blockchain tables based on version columns even if they are split across system chains.

Multiple user chains increase the flexibility of applying blockchain tables and their verification procedures to make it easier to leverage tamper-resistant tables in your applications.

Related Resources

[View Documentation](#)

Blockchain Table Row Versions

The blockchain table row version feature allows you to have multiple historical versions of a row that is maintained within a blockchain table corresponding to a set of user-defined columns. A view `<bctable>_last$` on top of the blockchain table allows you to see just the latest version of a row.

This feature allows you to guarantee row versioning when using tamper-resistant blockchain tables in your application.

Related Resources

[View Documentation](#)

Add and Drop User Columns in Blockchain and Immutable Tables

This feature allows evolution of Blockchain and Immutable Tables, namely it allows columns to be added and dropped while maintaining the current data, including that in dropped columns for continuity of crypto-hash chains.

As applications evolve you may need to modify existing tables by adding or dropping columns. In this release, you can easily add or drop columns in previously created Blockchain or Immutable tables. Any rows prior to a column deletion will maintain the data in these columns in order to preserve the integrity of the crypto-hash chains and allow the verification procedures to work across the entire table.

Related Resources[View Documentation](#)

Blockchain Table Countersignature

You can request a database countersignature at the time of signing a row. In addition to recording the countersignature and its metadata in the row, the countersignature and the `signed_bytes` are returned to the caller. The caller can then save the countersignature and `signed_bytes` in another data store, such as Oracle Blockchain Platform, for non-repudiation purposes.

A countersignature can provide user additional guarantees that data has been securely stored in the blockchain table.

Related Resources[View Documentation](#)

Blockchain Table Delegate Signer

A delegate is an alternate user who's allowed to sign rows inserted by the primary user. This feature allows a delegate to sign rows in an immutable or blockchain table on behalf of another user. A delegate's signature is accepted only if the signature can be verified using the public key in the delegate's certificate, which has been added to the dictionary table.

A delegate signer can be used when users are not able to sign the rows they created and they trust their delegate.

Related Resources[View Documentation](#)

New Special Privilege Required to Set Long Idle Retention Times for Blockchain and Immutable Tables

Blockchain or immutable tables with idle retention set to a sufficiently large value cannot be dropped until the newest row of the table becomes very old. This limits the ability to drop the blockchain/immutable table if necessary to prevent a disk space exhaustion attack. Hence, the operation of setting a table's idle retention to a large value is restricted to privileged users via a grant of a new `TABLE RETENTION` system privilege. The idle retention threshold, which specifies when to require the new privilege `BLOCKCHAIN_TABLE_RETENTION_THRESHOLD`, is configurable.

Ability to create blockchain or immutable tables with long retention times and inserting large amounts of data that can not be deleted could potentially be a vector for a denial of service attack via disk space exhaustion. To reduce this risk, the special privilege has been introduced. Only users granted this privilege can set idle retention above the configurable threshold level.

Related Resources

[View Documentation](#)

Database Architecture

Lock-Free Reservations

Lock-Free Reservations enables concurrent transactions to proceed without being blocked on updates of heavily updated rows. Lock-Free reservations are held on the rows instead of locking them. Lock-Free Reservations verifies if the updates can succeed and defers the updates until the transaction commit time.

Lock-Free Reservations improves the end user experience and concurrency in transactions.

Related Resources

[View Documentation](#)

Wide Tables

The maximum number of columns allowed in a database table or view has been increased to 4096. This feature allows you to build applications that can store attributes in a single table with more than the previous 1000-column limit. Some applications, such as Machine Learning and streaming IoT application workloads, may require the use of de-normalized tables with more than 1000 columns.

You now have the ability to store a larger number of attributes in a single row which for some applications may simplify application design and implementation.

Related Resources

[View Documentation](#)

Consolidated Service Backgrounds for Oracle Instance

We are introducing a new set of service processes which execute database service actions.

Service actions are responsible for maintenance tasks, parallel tasks and brokered tasks, consolidated tasks and many more. These were performed by dedicated processes in the database before. The new background scheduler group processes can execute any of these service actions, thus providing consolidation of background service actions.

Related Resources

[View Documentation](#)

Improve Performance and Disk Utilization for Hybrid Columnar Compression

Enhancements to the compression algorithms for Hybrid Columnar Compression (HCC) include improvements for faster compression and decompression speeds, as well as better compression ratios for newly created HCC compressed tables or for existing HCC compressed tables that are rebuilt. The exact benefits can vary based on the data and the chosen compression level.

This feature improves an application's workload performance while reducing database storage utilization.

Related Resources

[View Documentation](#)

Unrestricted Direct Loads

Prior to this feature, after a direct load and prior to a commit, queries and additional DMLs were not allowed on the same table for the same session or for other database sessions. This enhancement allows the loading session to query and perform DML on the same table that was loaded. Other sessions are also able to concurrently perform direct loads and DML. Rollback to a savepoint is also supported.

This feature removes the restrictions that you may have encountered when loading and querying data. Potentially improving the performance of your applications in areas such as Data Warehousing and complex batch processing.

Related Resources

[View Documentation](#)

General

Unrestricted Parallel DMLs

Oracle Database allows DML statements (`INSERT`, `UPDATE`, `DELETE`, and `MERGE`) to be executed in parallel by breaking the DML statements into mutually exclusive smaller tasks. Executing DML statements in parallel can make DSS queries, batched OLTP jobs, or any larger DML operations faster. However, parallel DML operations had a few transactional limitations.

This includes a limitation that restricted transactions with multiple per-table parallel DMLs. This means that once an object is modified by a parallel DML statement, that object cannot be read or modified by later statements of the same transaction. This enhancement removes this limitation, enabling users to run parallel DMLs, and any combination of statements like queries, serial DML, and parallel DML on the same object, within the same transaction.

For users, this simplifies and speeds up data loading and analytical processing by making full use of Oracle Database's parallel execution and parallel query capabilities.

Related Resources

[View Documentation](#)

ACFS Auto Resize Variable Threshold

ACFS auto resize now allows you to configure the threshold percentage for your file system automatic resize.

A more flexible threshold is now available for your file systems auto resize. Previously, the threshold was fixed to 10%. Now, you can customize to your specific use case needs.

Related Resources[View Documentation](#)

ACFS Cross Version Replication

ACFS replication now allows for primary clusters to replicate to standby cluster on a previous or older release.

This feature will provide flexibility in replication configurations, providing ample time for upgrading and lifecycle maintenance.

Related Resources[View Documentation](#)

ACFS Encryption Migration from OCR to OKV

ACFS Encryption now allows you to migrate from OCR to OKV.

This feature allows for a centralized point for key management using Oracle Key Vault.

Related Resources[View Documentation](#)

ACFS Replication Switchover

A new command, `acfsutil repl switchover`, provides a coordinated failover. However, if ACFS cannot establish contact the replication primary site, the command will fail.

Enhanced flexibility in ACFS replication management is now available with the addition of this new command.

Related Resources[View Documentation](#)

ACFS Snapshots RMAN Sparse Backup and Restore

You can now back up and restore PDB snapshot copies on ACFS.

Backing up and restoring PDB snapshot copies on ACFS, provides the space-efficient storage that is inherent of ACFS Snapshots.

Related Resources[View Documentation](#)

ACFS Sparse Backup and Restore of Snapshots

The `acfsutil snap duplicate` command can now generate a backup of an entire ACFS file systems and its snapshots, while preserving its sparseness.

You can now apply a full backup to another location while retaining the original sparseness. You can now replicate an entire ACFS file system and its snapshot tree with this new functionality.

Related Resources[View Documentation](#)

ACFSutil plogconfig Log Files Wrapping Info

ACFSutil plogconfig offers you a way to manage persistent logging configuration settings. `acfsutil plogconfig -q` will now offer you additional information on whether the logs have wrapped or not. You can also get this information with `acfsutil plogconfig -w`, which will offer only this information and not all the comprehensive information offered by `acfsutil plogconfig -q`.

Further information regarding persistent logging is now available, hence enhancing the experience in the realm of diagnosability.

Related Resources[View Documentation](#)

CEIL and FLOOR for DATE, TIMESTAMP, and INTERVAL Data Types

You can now pass `DATE`, `TIMESTAMP`, and `INTERVAL` values to the `CEIL` and `FLOOR` functions. These functions include an optional second argument to specify a rounding unit. You can also pass `INTERVAL` values to `ROUND` and `TRUNC` functions.

These functions make it easy to find the upper and lower bounds for date and time values for a specified unit.

Related Resources[View Documentation](#)

Centralized Config Providers

Database clients can securely pull application configuration data from Azure or OCI Cloud. The store can contain data such as application connection descriptors and tuning parameters. The configuration can also be specified in a JSON file loaded over HTTPS or from a local file system.

Central configuration makes application management and scaling easier. It fits well with architectures, such as microservices and serverless deployments.

Related Resources[View Documentation](#)

Oracle Data Pump Filters GoldenGate ACDR Columns from Tables

The ACDR feature of Oracle GoldenGate adds hidden columns to tables to resolve conflicts when the same row is updated by different databases using active replication. GoldenGate can also create a "tombstone table," which records interesting column values for deleted rows. Oracle Data Pump can exclude the hidden columns and the tombstone tables by setting a new import transform parameter, `OMIT_ACDR_METADATA`.

Oracle Data Pump enhances migration flexibility. It can migrate data from an Oracle GoldenGate ACDR (automatic conflict detection and resolution) environment to a non-ACDR environment by excluding the GoldenGate ACDR metadata during import.

Related Resources

[View Documentation](#)

PDB Snapshot Carousel ACFS Support

Oracle ACFS now supports PDB Snapshot Carousel, which allows you to maintain a library of PDB Snapshots.

Oracle Database files stored on Oracle ACFS file systems can now leverage PDB Snapshot Carousel in conjunction with ACFS snapshot technology.

Related Resources

[View Documentation](#)

SQL*Loader Supports SODA (Simple Oracle Document Access)

SQL*Loader now supports Simple Oracle Document Access (SODA). You can insert, append, and replace external documents into SODA collections in Oracle Database applications by using the SQL*Loader utility in both control file and express modes.

SQL*Loader support for Simple Oracle Document Access (SODA) makes it easier and faster to load schema-less JSON or XML-based application data into Oracle Database.

Related Resources

[View Documentation](#)

Manageability and Performance

Advanced LOW IOT Compression

An index-organized table (IOT) is a table stored in a variation of a B-tree index structure where rows are ordered by primary key. IOTs are useful because they provide fast random access by primary key without duplicating primary key columns in two structures – a heap table and an index. In earlier releases, IOTs only supported Oracle's prefix key compression, which required additional analysis and had the possibility of negative compression (where the overhead of compression outweighed the compression benefits).

Advanced LOW IOT Compression allows you to reduce the overall storage for Oracle Databases.

Related Resources

[View Documentation](#)

Automatic SecureFiles Shrink for Autonomous Database

Automatic SecureFiles Shrink for Autonomous Database automatically selects SecureFiles LOB segments based on a set of criteria and executes the free space shrink operation in the background for the selected segments. With Automatic SecureFiles Shrink and Autonomous

Database, the shrink operation happens transparently in small and gradual steps over time while allowing DDL and DML statements to execute concurrently. In the manual method, you must decide on which LOB segments to shrink using tools like Segment Advisor and use a DDL statement to execute the shrink operation. The manual method may not be feasible for very large LOB segments because it is time-consuming.

Automatic SecureFiles Shrink for Autonomous Database simplifies administrator duties and saves time due to the automation of this process.

Related Resources

[View Documentation](#)

Automatic Storage Compression

Organizations use Hybrid Columnar Compression for space saving and fast analytics performance. However, the compression and decompression overhead of Hybrid Columnar Compression can affect direct load performance. To improve direct load performance, Automatic Storage Compression enables Oracle Database to direct load data into an uncompressed format initially, and then gradually move rows into Hybrid Columnar Compression format in the background.

Automatic Storage Compression improves direct load performance, while keeping the advantages of Hybrid Columnar Compression, including space savings and fast analytics performance.

Related Resources

[View Documentation](#)

Enhanced Query History Tracking and Reporting

Enhanced Query History Tracking and Reporting lets you track and report on a more complete history of user-issued queries than is available in previous releases. This feature provides you with greater capability to track user-initiated queries within a session. It includes non-parallel queries with less than five seconds of execution time, which are not tracked with Real-time SQL Monitoring unless tracking is forced by a hint. Each user can access and report on their own current session history. SYS users and DBAs can view and get query history reports for all current user sessions and can also turn this functionality on or off. Reporting is configurable, with options for selecting the reporting scope and detail level.

Enhanced Query History Tracking and Reporting allows application developers and development operations (DevOps) personas to get detailed insight into the queries that execute on your databases. This insight allows you to better manage and optimize your applications.

Related Resources

[View Documentation](#)

Fast Ingest (Memoptimize for Write) Enhancements

This feature adds enhancements to Memoptimize Rowstore Fast Ingest with support for partitioning, compressed tables, fast flush using direct writes, and direct In-Memory column store population support. These enhancements make the Fast Ingest feature easier to incorporate in more situations where fast data ingest is required.

This feature helps Oracle Database provide better support for applications requiring fast data ingest capabilities. Data can be ingested and then processed all in the same database. This reduces the need for special loading environments, and thus reduces complexity and data redundancy.

Related Resources

[View Documentation](#)

Improved Performance of LOB Writes

You can experience improved read and write performance for LOBs due to the following enhancements:

- Multiple LOBs in a single transaction are buffered simultaneously. This improves performance when you use switch between LOBs while writing within a single transaction.
- Various enhancements, such as acceleration of compressed LOB append and compression unit caching, improve the performance of reads and writes to compressed LOBs.
- The input-output buffer is resized based on the input data for large writes to LOBs with the NOCACHE option. This improves the performance for large direct writes, such as writes to file systems on DBFS and OFS.

This feature adds a host of improvements to accelerate SecureFiles writes for JSON document-based applications, for write calls issued by a database file system, and also for LOB workloads where the underlying data is compressed for storage savings.

Related Resources

[View Documentation](#)

Improved System Monitor (SMON) Process Scalability

Queries can require large amounts of temporary space and some temporary space operations run in critical background processes, like the System Monitor (SMON) process. SMON is responsible for cleaning up temporary segments that are no longer in use. SMON checks regularly to see whether it is needed, and other processes can call SMON. Temporary space management can affect SMON's scalability for other critical actions. This new enhancement instead uses the Space Management Coordinator (SMCO) process so that the responsibility of managing temporary space is offloaded from SMON, thereby improving its scalability.

This feature improves the overall scalability of the SMON process, particularly in a multitenant Oracle RAC cluster.

Related Resources

[View Documentation](#)

Pluggable Database Support in Oracle Data Guard Environments

In release 23c, there is now a pluggable database configuration in a Data Guard environment using Database Configuration Assistant (DBCA).

A command line based silent mode option is available for configuring pluggable databases (PDBs) in Data Guard environments.

Related Resources

[View Documentation](#)

Refreshable PDBs in DBCA

Database Configuration Assistant (DBCA) allows you to clone a remote Pluggable database (PDB) as a refreshable PDB. When a PDB is created as refreshable, the changes of the source PDB will periodically propagate to the refreshable PDB. The refreshable PDB can be configured to refresh manually or automatically during creation.

A DBCA-based graphical user interface or scripted silent mode for cloning a remote refreshable PDB reduces many commands needed to create a remote refreshable PDB clone ensuring a faster and more reliable cloning of PDBs.

Related Resources

[View Documentation](#)

10

Diagnosability

General

Cluster Health Monitor Improved Diagnosability

Cluster Health Monitor introduces a new diagnostic ability to listen for critical component events that could indicate pending or actual failure and report these with recommended corrective actions. In some cases, these actions may be executed autonomously. Such events and actions could then be captured and admins notified through components such as Trace File Analyzer.

Improving the robustness and reliability of the Oracle Database hosting infrastructure is a critical business requirement for enterprises. This improved ability to detect and correct at first failure and self-heal autonomously delivers value by improving business continuity.

Related Resources

[View Documentation](#)

Enhanced Cluster Health Advisor Support for Oracle Pluggable Databases

Cluster Health Advisor's (CHA) diagnostics capability is extended to support 4K pluggable databases (PDBs) from 256. This is critical for Oracle Autonomous Database deployments. CHA's problem detection and root cause analysis improves accuracy by considering database events such as reconfiguration. This improves detection, analysis, and targeted preventative actions for problems, such as instance evictions.

By adding this support to Cluster Health Advisor, performance and availability are kept in line with the deployment size. The business continuity of critical applications is preserved with improved prognostics and targeted preventive actions.

Related Resources

[View Documentation](#)

Reduce Time to Resolve

Add Verified SQL Plan Baseline

The SQL plan management API (DBMS_SPM) includes a new procedure called `ADD_VERIFIED_SQL_PLAN_BASELINE`. It searches the cursor cache, AWR, and automatic SQL tuning set to establish which execution plan is best for a specified SQL statement. It creates an accepted SQL plan baseline for the best plan.

This feature provides improved performance management.

Related Resources

[View Documentation](#)

CMAN Diagnostics and Logging Enhancements

Using the command line interface, you can now monitor statistics for all database service registration operations that the Oracle Connection Manager (CMAN) listener performs. You can also view additional diagnostic details about service registration events in the CMAN and listener log files.

You can view the following details:

- Statistics about registration operations (such as register, update, or unregister) for services, endpoints, handlers, and other information per instance level or for all instances. You can query these details either periodically or for a particular period.
- Additional diagnostic details about service updates, blocked service handlers, and target database instances for client connections in the CMAN and listener log files.

This feature enables you to evaluate statistics about service registration operations at both global and instance levels, analyze their traffic, and diagnose registration issues.

Related Resources

[View Documentation](#)

DBMS_DICTIONARY_CHECK PL/SQL Package

DBMS_DICTIONARY_CHECK is a read-only and light weight PL/SQL package procedure that helps you identify database dictionary inconsistencies that are manifested in unexpected entries in the RDBMS dictionary tables or invalid references between dictionary tables. Database dictionary inconsistencies can cause process failures and, in some cases, instances crash. DBMS_DICTIONARY_CHECK assists you in identifying such inconsistencies and provides a guided remediation to resolve the problem and avoid such database failures.

This feature improves database availability thus reducing the management and maintenance time for environments utilizing this package.

Related Resources

[View Documentation](#)

Estimate the Space Saved with Deduplication

Before you enable deduplication, you can estimate the space that you can save by enabling advanced LOB deduplication for existing LOBS.

This enables you to take an informed decision to enable deduplication. Advanced LOB Deduplication enables Oracle Database to automatically detect duplicate LOB data within a LOB column or partition, and conserve space by storing only one copy of the data.

Related Resources

[View Documentation](#)

Extent-Based Scrubbing

Automatic Storage Management (ASM) extent-based scrubbing changes the granularity level on which ASM scrubs data from a file and disk group level to the extent level.

Compared to scrubbing the whole file, scrubbing specific extent sets significantly reduces the scrubbing turn-around time, improves the data availability, and minimizes the performance impact.

Related Resources

[View Documentation](#)

High Availability Diagnosability Using the DBMS_SCHEDULER Package

The Scheduler In-Memory Tracing feature is aimed at designing and implementing tools for the collection and temporary in-memory storage of scheduler trace messages generated during process execution.

It is critical to successfully restart jobs when they are interrupted by forced shut downs, like a forced patching cycle. With the addition of High Availability (HA) diagnostics in the `DBMS_SCHEDULER` package, you will be able to add real-time in-memory diagnostics during forced shut downs, and address any issues that result from these diagnostics.

This feature provides benefits like easier collection of trace messages generated since the initial failure, reduction in user interaction to collect traces, and significant reduction in multiple requests of problem reproduction.

Related Resources

[View Documentation](#)

In-Memory Advisor

The In-Memory Advisor is now part of Oracle Database and has two components: (1) an eligibility test that identifies databases that are not good candidates for Database In-Memory and (2) an advisor with enhanced analysis capability to better identify workloads that will benefit from Database In-Memory.

The In-Memory Advisor makes it easier and faster to identify databases that can take advantage of Database In-Memory. The In-Memory Advisor is now built into the database in place of having to install a separate standalone utility. An eligibility test provides the ability to quickly eliminate workloads that will not benefit from Database In-Memory, saving time and effort. An enhanced analysis capability that makes identification of analytic workloads that will benefit from Database In-Memory simpler and more accurate. Together, these two components make it much simpler to decide where and when to use Database In-Memory.

Related Resources

[View Documentation](#)

Oracle Call Interface (OCI) APIs to Enable Client-Side Tracing

New Oracle Call Interface (OCI) APIs allow applications to enable and disable client-side OCI diagnostic tracing dynamically without the need to update configuration files or set environment variables.

This feature allows developers to improve OCI application problem troubleshooting and reduce issue resolution time.

Related Resources

[View Documentation](#)

Rename LOB Segment

To rename an existing LOB segment users perform an operation such as `ALTER TABLE ... MOVE`, which could perform slowly since the operation physically moves the LOB data as part of the renaming.

This enhancement improves the performance of renaming a LOB segment, at the table, partition and subpartition level by eliminating the physical movement of the LOB data.

Related Resources

[View Documentation](#)

Selective In-Memory Columns

With Selective In-Memory columns, it is now easier to add or exclude columns for in-memory. An `ALL` sub-clause has been added so that all columns can be either enabled or disabled from in-memory. This reduces the need to have very long strings of included or excluded columns.

With Selective In-Memory columns, the ability to specify `ALL` columns to be enabled or disabled from in-memory reduces the need for very long strings of columns, which reduces the chance for errors and makes configuring in-memory enabled tables easier.

Related Resources

[View Documentation](#)

Installation, Upgrade, and Patching

Upgrade

AutoUpgrade Release Update (RU) Upgrades

AutoUpgrade supports the option of using AutoUpgrade to perform out-of-place Oracle home Release Update patching.

For an out-of-place patch of Oracle Database using AutoUpgrade, AutoUpgrade moves the source database that you want to patch to a new Oracle Database Oracle home, and then patches the database binaries in that target Oracle home with the Release Update that you select. With this option, you can use AutoUpgrade at any time that you want to move the database to a new Oracle home, either as part of a planned upgrade or as part of a patch plan. In a patch operation, AutoUpgrade performs the patch using the following workflow:

1. AutoUpgrade recognizes that the source database and the target Oracle Database are the same base release.
2. AutoUpgrade skips the upgrade steps.
3. AutoUpgrade patches the target database using the Release Update.

Related Resources

[View Documentation](#)

AutoUpgrade Sets Parallelism Based on System Resources

AutoUpgrade automatically evaluates system resources and makes an intelligent decision as to how many upgrade jobs can run simultaneously.

AutoUpgrade uses the `CPU_COUNT` value and system process parameters to determine available system resources, and calibrates both the number of upgrades that can run at a time and the number of parallel threads for each upgrade. Upgrades that exceed a safe threshold are put in a queue so that they can be run as system resources become available.

Related Resources

[View Documentation](#)

AutoUpgrade Supports Upgrades with Keystore Access to Databases Using TDE

AutoUpgrade enhances support for databases that use transparent data encryption (TDE) by enabling keystore generation.

AutoUpgrade now enables you to provide passwords to an external key manager generated and maintained by AutoUpgrade. With this configuration, AutoUpgrade supports unmanned

or automated operations of TDE-enabled databases. AutoUpgrade can open the source database keystore without prompting for the keystore password, and enroll the target database into the TDE external keystore for key management, so that the target database can start automatically.

Related Resources

[View Documentation](#)

AutoUpgrade Unplug-Plugin Upgrades to Different Systems

You can now use the Oracle Database AutoUpgrade Unplug/Plug method to unplug a PDB from one system and plug into a different system and upgrade.

In earlier releases, AutoUpgrade supported unplug/plugin/upgrades on the same server, but it was not possible to unplug a PDB from one server, plug it into a different system, and then upgrade the PDB. With this feature, you can now migrate and upgrade the PDB in a single operation, including migrations to the cloud.

Related Resources

[View Documentation](#)

REST APIs for AutoUpgrade

To facilitate safe and secure remote use of the AutoUpgrade for Oracle Database upgrades, AutoUpgrade now provides REST APIs (ORDS and OCI).

The Oracle REST Data Services (ORDS) database API is a database management and monitoring REST API embedded into Oracle REST Data Services. The Oracle Cloud Infrastructure (OCI) REST API is enabled by configuring the REST Adapter connection to use the OCI Signature Version 1 security policy. You can now use these features to run AutoUpgrade upgrades remotely over SSH.

Related Resources

[View Documentation](#)