

Oracle® Database

Database New Features Guide



19c
F14269-06
October 2019

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

Copyright © 2015, 2019, Oracle and/or its affiliates. All rights reserved.

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 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 installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. 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 and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon 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, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

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

Contents

Preface

Audience	vii
Documentation Accessibility	vii
Related Documents	vii
Conventions	vii

1 Oracle Database Release 19c New Features

Application Development	1-1
Application Express	1-1
Improved Create Application Wizard	1-1
REST Enabled SQL Support	1-2
Social Sign-In Authentication	1-2
Web Source Modules	1-2
Improved Create Page Wizard	1-3
New REST Workshop	1-3
General	1-4
Application Continuity for Java: Declarative Request Demarcation	1-4
Application Continuity for Java: New States Management	1-4
Easy Connect Plus	1-4
Oracle Network Log File Segmentation	1-5
SQL*Net: Auto-Detection of Support for Out-of-Band Breaks	1-5
JSON	1-5
Materialized View Support for Queries Containing JSON_TABLE	1-5
JSON Update Operations	1-6
SQL/JSON Syntax Simplifications	1-6
JSON Object Mapping	1-6
New SQL/JSON Function JSON_SERIALIZE and JSON Data Guide Support for GeoJSON Data	1-6
SQL	1-6
DISTINCT Option for LISTAGG Aggregate	1-7
Availability	1-7
General	1-7

Dynamically Change Oracle Data Guard Broker Fast-Start Failover Target	1-7
Simplified Database Parameter Management in Oracle Data Guard Broker	1-8
Simplified Database Parameter Management in Oracle Data Guard Broker	1-8
Propagate Restore Points from Primary to Standby Site	1-8
Simplified Database Parameter Management in Oracle Data Guard Broker	1-8
Oracle Data Guard Multi-Instance Redo Apply Works with the In-Memory Column Store	1-9
Active Data Guard DML Redirection	1-9
Simplified Database Parameter Management in Oracle Data Guard Broker	1-9
Clear Flashback Logs Periodically for Increased Fast Recovery Area Size Predictability	1-10
New Parameters to Tune Automatic Outage Resolution with Oracle Data Guard	1-10
Finer Granularity Supplemental Logging	1-10
Sharding	1-11
Support for Multi-Shard Query Coordinators on Shard Catalog Standby Databases	1-11
Generation of Unique Sequence Numbers Across Shards	1-11
Support for Multiple PDB Shards in the Same CDB	1-11
Multiple Table Family Support for System-Managed Sharding	1-12
Propagation of Parameter Settings Across Shards	1-12
Big Data and Data Warehousing	1-12
General	1-12
Automatic Indexing	1-12
SQL Diagnostics and Repair Enhancements	1-13
Bitmap Based Count Distinct SQL Function	1-13
Big Data and Performance Enhancements for In-Memory External Tables	1-13
Automatic SQL Plan Management	1-13
Real-Time Statistics	1-14
High-Frequency Automatic Optimizer Statistics Collection	1-14
Hybrid Partitioned Tables	1-14
Database Overall	1-15
Automated Installation, Configuration and Patching	1-15
Ability to Create a Duplicate of an Oracle Database Using DBCA in Silent Mode	1-15
Ability to Relocate a PDB to Another CDB Using DBCA in Silent Mode	1-15
Ability to Create a PDB by Cloning a Remote PDB Using DBCA in Silent Mode	1-15
Simplified Image-Based Oracle Database Client Installation	1-16
Root Scripts Automation Support for Oracle Database Installation	1-16
Support for Dry-Run Validation of Oracle Clusterware Upgrade	1-16
AutoUpgrade and Database Utilities	1-16
AutoUpgrade for Oracle Database	1-17

Oracle Data Pump Ability to Exclude ENCRYPTION Clause on Import	1-17
Oracle Data Pump Allows Tablespaces to Stay Read-Only During TTS Import	1-17
Oracle Data Pump Support for Resource Usage Limitations	1-18
Oracle Data Pump Test Mode for Transportable Tablespaces	1-18
Oracle Data Pump Prevents Inadvertent Use of Protected Roles	1-18
Oracle Data Pump Loads Partitioned Table Data One Operation	1-18
Oracle Data Pump Allows Wildcards for Dump File in Object Store	1-19
Oracle Data Pump Import Supports More Object Store Credentials	1-19
Diagnosability	1-19
General	1-19
Oracle Trace File Analyzer Support for Using an External SMTP Server for Notifications	1-20
Oracle Cluster Health Advisor Integration into Oracle Trace File Analyzer	1-20
Oracle Trace File Analyzer REST API Support	1-20
Oracle Trace File Analyzer Search Extended to Support Metadata Searches	1-21
Oracle ORAchk and Oracle EXAchk REST Support	1-21
Oracle ORAchk and Oracle EXAchk Support for Encrypting Collection Files	1-21
Oracle ORAchk and Oracle EXAchk Support for Remote Node Connections Without Requiring Passwordless SSH	1-22
Oracle ORAchk and Oracle EXAchk Now Show Only the Most Critical Checks by Default	1-22
Oracle Trace File Analyzer Supports New Service Request Data Collections	1-22
Performance	1-23
General	1-23
SQL Quarantine	1-23
Database In-Memory Wait on Populate	1-23
Resource Manager Automatically Enabled for Database In-Memory	1-23
Memoptimized Rowstore Fast Ingest	1-24
Automatic Database Diagnostic Monitor (ADDM) Support for Pluggable Databases (PDBs)	1-24
Resource Manager Automatically Enabled for Database In-Memory	1-24
High-Frequency SQL Plan Management Evolve Advisor Task	1-24
Workload Capture and Replay in a PDB	1-25
RAC and Grid	1-25
General	1-25
Parity Protected Files	1-25
Secure Cluster Communication	1-26
Automated PDB Relocation	1-26
Zero-Downtime Oracle Grid Infrastructure Patching	1-26
Automated Transaction Draining for Oracle Grid Infrastructure Upgrades	1-26
Oracle Restart Patching and Upgrading	1-27

Colocation Tag for Client Routing	1-27
Optional Install for the Grid Infrastructure Management Repository	1-27
Resupport of Direct File Placement for OCR and Voting Disks	1-28
Dynamic Services Fallback Option	1-28
Security	1-28
General	1-28
New ALTER SYSTEM Clause FLUSH PASSWORDFILE_METADATA_CACHE	1-29
Transparent Online Conversion Support for Auto-Renaming in Non-Oracle- Managed Files Mode	1-29
Support for Additional Algorithms for Offline Tablespace Encryption	1-29
Key Management of Encrypted Oracle-Managed Tablespaces in Transparent Data Encryption	1-30
Support for Host Name-Based Partial DN Matching for Host Certificates	1-30
New PDB_GUID Audit Record Field for SYSLOG and the Windows Event Viewer	1-30
New EVENT_TIMESTAMP_UTC Column in the UNIFIED_AUDIT_TRAIL View	1-31
Passwords Removed from Oracle Database Accounts	1-31
Signature-Based Security for LOB Locators	1-31
Unified Auditing Top-Level Statements	1-32
Privilege Analysis Now Available in Oracle Database Enterprise Edition	1-32
Support for Oracle Native Encryption and SSL Authentication for Different Users Concurrently	1-32
Ability to Grant or Revoke Administrative Privileges to and from Schema- Only Accounts	1-33
Automatic Support for Both SASL and Non-SASL Active Directory Connections	1-33
Database Vault Operations Control for Infrastructure Database Administrators	1-33
Database Vault Command Rule Support for Unified Audit Policies	1-33

Preface

This document describes new features implemented in Oracle Database 19c.

- [Audience](#)
- [Documentation Accessibility](#)
- [Related Documents](#)
- [Conventions](#)

Audience

Read Oracle Database New Features Guide if you want to learn about features, options, and enhancements that are new in Oracle Database 19c.

Documentation Accessibility

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

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Related Documents

For more information, see the following documents in the Oracle Database 19c documentation set:

- *Oracle Database Error Messages*
- *Oracle Database Administrator's Guide*
- *Oracle Database Concepts*
- *Oracle Database Reference*

Conventions

The following text conventions are used in this document:

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

1

Oracle Database Release 19c New Features

This chapter contains descriptions of all of the features that are new to Oracle Database Release 19c.

- [Application Development](#)
- [Availability](#)
- [Big Data and Data Warehousing](#)
- [Database Overall](#)
- [Diagnosability](#)
- [Performance](#)
- [RAC and Grid](#)
- [Security](#)

Application Development

- [Application Express](#)
- [General](#)
- [JSON](#)
- [SQL](#)

Application Express

- [Improved Create Application Wizard](#)
- [REST Enabled SQL Support](#)
- [Social Sign-In Authentication](#)
- [Web Source Modules](#)
- [Improved Create Page Wizard](#)
- [New REST Workshop](#)

Improved Create Application Wizard

The updated Create Application Wizard features a new low-code approach to creating applications and simpler, modernized wizards for creating applications.

The Create Application Wizard now supports the ability to create advanced pages such as Dashboards and Master-Detail. The wizard also supports adding common frameworks or "Features" when creating an application such as access control, activity reporting, or theme selection. In addition, the revamped wizard supports the ability to

customize user interface options such as Theme Style, the application icon, and page icons.

You also have the ability to refine a previous wizard definition by going back into the Create Application Wizard and retrieving the definition from a previous wizard (or blueprint) and then update the definitions and regenerate another application.

Related Topics

- *Oracle® Application Express App Builder User's Guide*

REST Enabled SQL Support

You can easily create REST enabled SQL references by defining a name, the endpoint URL, and authentication information within shared components.

Oracle Application Express passes the SQL or PL/SQL query to ORDS over REST returning a self-describing JavaScript Object Notation (JSON) response. The JSON object contains result set metadata, the result data, and pagination details. REST enabled SQL references are used as the basis for all report types, such as interactive reports and classic reports, but not interactive grid regions. References can also be used with calendars, Oracle JET Data Visualization components (JET charts), trees, and PL/SQL processes.

Each SQL statement defines Oracle Database links and work over SQL*Net (or over the internet in cloud environments), and must open a session within the remote database for each SQL or PL/SQL executed. By contrast, REST enabled SQL references are defined at the Oracle Application Express workspace-level, and work with JSON over HTTP and HTTPS, which makes them easy to use in cloud environments or over the internet. References can also scale significantly better since ORDS utilizes a connection pool on the remote database.

Related Topics

- *Oracle® Application Express App Builder User's Guide*

Social Sign-In Authentication

Social Sign-In authentication scheme supports authentication with Google, Facebook, and other social networks that support OpenIDConnect / OAuth2 standards.

Social Sign-In authentication is primarily useful for internet-facing applications where an unknown number of users from social networks may use the application, or if the company has standardized on systems that perform user credential verification for authentication such as Oracle Identity Cloud Service, an internal OpenIDConnect, or OAuth2 system.

Related Topics

- *Oracle® Application Express App Builder User's Guide*

Web Source Modules

Web Source Modules provide a declarative method to define references to external Representational State Transfer (REST) APIs and generic JSON data feeds.

In earlier Oracle Application Express releases, you could define Simple Object Access Protocol (SOAP) and REST web services and then utilize them within limited Oracle

Application Express components. Defining such services was manual, time consuming, and error prone. The new Web Source Modules are highly declarative because they use discovery to understand and define the incoming structure of the web service.

Web Source Modules store additional metadata about how to parse response data and map it as a virtual table with rows and columns. A module can contain one or many Web Source Operations, which are the references to a concrete external web service. Modules also include post-processing SQL, which modifies the data being processed by the Oracle Application Express component. You can use post-processing SQL to apply functions, aggregations, or join with local tables.

Web Source Modules are the basis for all report types, such as interactive reports and classic reports, but not interactive grid regions. You can also use these modules with calendars, Oracle JET Data Visualization components (JET charts), trees, and PL/SQL processes.

Related Topics

- *Oracle® Application Express App Builder User's Guide*

Improved Create Page Wizard

The updated Create Page Wizard features new page types, common frameworks or "Features" for existing applications, and support for email and job reporting.

The new Side by Side Master Detail page features a left panel for searching the master record and a right panel that displays the master record using a value pair report, and up to four detail reports using classic reports. The new Dashboard page enables you to select from different chart layouts that are based on sample data. The generated charts can very easily be updated in Page Designer post-generation.

The wizard also enables you to include new common frameworks or "Features" for an application, such as access control, activity reporting, theme selection, and more (providing the application is utilizing the Universal Theme). The wizard also supports Email Reporting, Job Reporting (providing jobs are defined in the default schema), and the ability to create an Administration page to manage features.

Related Topics

- *Oracle® Application Express App Builder User's Guide*

New REST Workshop

If Oracle Application Express uses Oracle REST Data Services (ORDS) 17.4 or later, the new REST Workshop utilizes the ORDS repository.

Prior to Oracle Database release 18c, version 18.1, RESTful services definitions created within Oracle Application Express were stored within the metadata tables of the core Oracle Application Express schema. Utilizing the ORDS repository for REST services makes it easier to manage RESTful services in a single place using a multitude of tools, including Oracle Application Express, SQL Developer, SQL*Plus, and Oracle SQL Developer Command-Line (SQLcl).

While Oracle Application Express-based REST services continue to work, you cannot create new or edit existing Oracle Application Express based RESTful services. Oracle recommends migrating all RESTful services to the Oracle REST Data Services (ORDS) repository.

Related Topics

- [Oracle® Application Express SQL Workshop Guide](#)

General

- [Application Continuity for Java: Declarative Request Demarcation](#)
- [Application Continuity for Java: New States Management](#)
- [Easy Connect Plus](#)
- [Oracle Network Log File Segmentation](#)
- [SQL*Net: Auto-Detection of Support for Out-of-Band Breaks](#)

Application Continuity for Java: Declarative Request Demarcation

When Application Continuity for Java is configured in `AUTO` mode (that is service `FAILOVER_TYPE=AUTO`), the Java Database Connectivity (JDBC) driver injects a `beginRequest` call at runtime after the creation of a JDBC connection with the replay data source.

This feature ensures zero downtime for Java applications and third-party connection pools without the need to make code changes.

Related Topics

- [Oracle® Database JDBC Developer's Guide](#)

Application Continuity for Java: New States Management

This feature introduces new session states including `AL8KW_ERR_OVLAP`, `AL8KW_EDITION`, `AL8KW_SQL_TXLP`, and `AL8KW_ROW_ARCHIVAL`. These session states are saved during normal activity and restored at failover when `FAILOVER_RESTORE` is set and `FAILOVER` equals `AUTO`.

This feature enhances transparency in Application Continuity for Java.

Related Topics

- [Oracle® Database JDBC Developer's Guide](#)

Easy Connect Plus

The Easy Connect syntax that applications use to connect to Oracle Database has improved functionality. The new version is called Easy Connect Plus.

Easy Connect Plus simplifies Oracle Database application configuration and deployment for common use cases. With Easy Connect Plus, you no longer need to configure Oracle Net parameter files such as `tnsnames.ora` and `sqlnet.ora`. Easy Connect Plus also no longer requires you to set the `TNS_ADMIN` environment variable.

Related Topics

- [Oracle® Database Net Services Administrator's Guide](#)

Oracle Network Log File Segmentation

This feature allows you to configure the maximum size and number of text log files for Oracle Network components, such as Oracle Net Listener, Connection Manager (CMAN), and global services manager.

This feature allows better management of log files, particularly in Cloud environments.

Related Topics

- [Oracle® Database Net Services Administrator's Guide](#)

SQL*Net: Auto-Detection of Support for Out-of-Band Breaks

This feature automatically probes the network path between the client and the server in order to determine the status of out-of-band support, and automatically enable or disable it.

Out-of-band breaks were enabled by default for UNIX platforms in past releases. However, this configuration causes numerous problems when network devices on the path between the client and the server do not allow out-of-band data to pass through. This data may either be dropped or inlined leading to server-side problems such as Transparent Network Substrate (TNS) errors or data corruption. These problems are often very hard to diagnose. The solution is to turn off usage of out-of-band data manually by setting a `sqlnet.ora` parameter.

Related Topics

- [Oracle® Database Net Services Administrator's Guide](#)

JSON

- [Materialized View Support for Queries Containing JSON_TABLE](#)
- [JSON Update Operations](#)
- [SQL/JSON Syntax Simplifications](#)
- [JSON Object Mapping](#)
- [New SQL/JSON Function JSON_SERIALIZE and JSON Data Guide Support for GeoJSON Data](#)

Materialized View Support for Queries Containing JSON_TABLE

Queries with `JSON_EXISTS`, `JSON_VALUE`, and other functions can now utilize a materialized view created over a query that uses `JSON_TABLE` function.

This feature is particularly useful when the JavaScript Object Notation (JSON) documents in a column contain arrays. This type of materialized view provides fast performance for accessing data within those JSON arrays.

Related Topics

- [Oracle® Database JSON Developer's Guide](#)

JSON Update Operations

You can now update JavaScript Object Notation (JSON) documents using new SQL function `JSON_MERGEPATCH`, applying one or more changes to multiple documents with a single statement.

This feature improves the flexibility of JSON update operations.

Related Topics

- [Oracle® Database JSON Developer's Guide](#)

SQL/JSON Syntax Simplifications

You can now use simpler syntax for field projection, SQL/JSON path expressions, and SQL/JSON generation function `JSON_OBJECT`.

The SQL interface for JavaScript Object Notation (JSON) processing is easier to use for certain operations.

Related Topics

- [Oracle® Database JSON Developer's Guide](#)

JSON Object Mapping

You can now map JavaScript Object Notation (JSON) data to and from SQL object types and collection types.

This feature makes it easier for programs that use SQL objects and collections to interact with JSON-based applications.

Related Topics

- [Oracle® Database JSON Developer's Guide](#)

New SQL/JSON Function `JSON_SERIALIZE` and JSON Data Guide Support for GeoJSON Data

You can use new SQL/JSON function `JSON_SERIALIZE` to serialize JavaScript Object Notation (JSON) data to text. Aggregate function `JSON_DATAGUIDE` can now detect GeoJSON geographic data.

You can use `JSON_SERIALIZE` to extract JSON values as text for printing or display. You can use SQL function `JSON_DATAGUIDE` to create a view that projects such data as `SDO_GEOMETRY` data.

Related Topics

- [Oracle® Database JSON Developer's Guide](#)

SQL

- [DISTINCT Option for LISTAGG Aggregate](#)

DISTINCT Option for LISTAGG Aggregate

The `LISTAGG` aggregate function now supports duplicate elimination by using the new `DISTINCT` keyword.

The `LISTAGG` aggregate function orders the rows for each group in a query according to the `ORDER BY` expression and then concatenates the values into a single string. You can remove duplicate values from the specified expression before concatenation into a single string using the new `DISTINCT` keyword. This removes the need to create complex query processing to find the distinct values before using the aggregate `LISTAGG` function. Use the `DISTINCT` option to remove duplicate values within the `LISTAGG` function.

The result is simpler, faster, more efficient SQL.

Related Topics

- [Oracle® Database Data Warehousing Guide](#)

Availability

- [General](#)
- [Sharding](#)

General

- [Dynamically Change Oracle Data Guard Broker Fast-Start Failover Target](#)
- [Simplified Database Parameter Management in Oracle Data Guard Broker](#)
- [Simplified Database Parameter Management in Oracle Data Guard Broker](#)
- [Propagate Restore Points from Primary to Standby Site](#)
- [Simplified Database Parameter Management in Oracle Data Guard Broker](#)
- [Oracle Data Guard Multi-Instance Redo Apply Works with the In-Memory Column Store](#)
- [Active Data Guard DML Redirection](#)
- [Simplified Database Parameter Management in Oracle Data Guard Broker](#)
- [Clear Flashback Logs Periodically for Increased Fast Recovery Area Size Predictability](#)
- [New Parameters to Tune Automatic Outage Resolution with Oracle Data Guard](#)
- [Finer Granularity Supplemental Logging](#)

Dynamically Change Oracle Data Guard Broker Fast-Start Failover Target

The fast-start failover target standby database can be changed dynamically, to another standby database in the target list, without disabling fast-start failover.

In earlier releases of Oracle Database, you had to disable fast-start failover to move to a new target standby database. This exposes the broker configuration to a period where automatic failover cannot be used at all. You use the `SET FAST_START`

FAILOVER TARGET command to dynamically change the fast-start failover target standby database.

Related Topics

- *Oracle® Data Guard Broker*

Simplified Database Parameter Management in Oracle Data Guard Broker

The management of database parameters in an Oracle Data Guard broker configuration is simplified by allowing all parameter management through SQL*Plus. Inconsistencies between a database's Data Guard parameter settings and the Data Guard Broker's property settings are eliminated.

You can now manage all Oracle Data Guard-related parameter settings using the SQL*Plus ALTER SYSTEM command or in the Data Guard broker command-line interface (DGMGRL) with the new EDIT DATABASE ... SET PARAMETER command. Parameter changes made in DGMGRL are immediately executed on the target database.

Simplified Database Parameter Management in Oracle Data Guard Broker

The management of database parameters in an Oracle Data Guard broker configuration is simplified by allowing all parameter management through SQL*Plus. Inconsistencies between a database's Data Guard parameter settings and the Data Guard Broker's property settings are eliminated.

You can now manage all Oracle Data Guard-related parameter settings using the SQL*Plus ALTER SYSTEM command or in the Data Guard broker command-line interface (DGMGRL) with the new EDIT DATABASE ... SET PARAMETER command. Parameter changes made in DGMGRL are immediately executed on the target database.

Propagate Restore Points from Primary to Standby Site

Restore points created on the primary database are propagated to the standby sites, so that they are available even after a failover operation.

Normal restore points or guaranteed restore points are defined at the primary site to enable fast point-in-time recovery in the event of logical corruptions. These restore points are stored in the control file. In the event of a failover, the standby database becomes the primary database. However, the restore point information is lost. Propagating restore points from the primary to the standby simplifies the complexity of the restore and recovery process after a failover because the standby database is updated with the restore points created on the primary database.

Related Topics

- *Oracle® Data Guard Concepts and Administration*

Simplified Database Parameter Management in Oracle Data Guard Broker

The management of database parameters in an Oracle Data Guard broker configuration is simplified by allowing all parameter management through SQL*Plus. Inconsistencies between a database's Data Guard parameter settings and the Data Guard Broker's property settings are eliminated.

You can now manage all Oracle Data Guard-related parameter settings using the SQL*Plus ALTER SYSTEM command or in the Data Guard broker command-line interface (DGMGRL) with the new EDIT DATABASE ... SET PARAMETER command. Parameter changes made in DGMGRL are immediately executed on the target database.

Oracle Data Guard Multi-Instance Redo Apply Works with the In-Memory Column Store

The In-Memory Column Store and Data Guard Multi-Instance Redo Apply can now be enabled at the same time on an Active Data Guard standby. Previously the two features were mutually exclusive.

You can now use the fastest redo apply technology (Data Guard Multi-Instance Redo Apply) and the fastest analytical query technology (In-Memory Column Store) on the same Active Data Guard standby to gain the best of both features. Multi-Instance Redo Apply uses information in the In-Memory Column Store on the Active Data Guard standby to increase apply speed where possible.

Related Topics

- *Oracle® Data Guard Concepts and Administration*

Active Data Guard DML Redirection

Incidental Data Manipulation Language (DML) operations can be run on Active Data Guard standby databases. This allows more applications to benefit from using an Active Data Guard standby database when some writes are required.

DML redirection helps in load balancing between the primary and standby databases. When incidental DML is issued on an Active Data Guard standby database, the update is passed to the primary database where it is executed. The resulting redo of the transaction updates the standby database after which control is returned to the application.

Related Topics

- *Oracle® Data Guard Concepts and Administration*

Simplified Database Parameter Management in Oracle Data Guard Broker

The management of database parameters in an Oracle Data Guard broker configuration is simplified by allowing all parameter management through SQL*Plus. Inconsistencies between a database's Data Guard parameter settings and the Data Guard Broker's property settings are eliminated.

You can now manage all Oracle Data Guard-related parameter settings using the SQL*Plus ALTER SYSTEM command or in the Data Guard broker command-line interface (DGMGRL) with the new EDIT DATABASE ... SET PARAMETER command. Parameter changes made in DGMGRL are immediately executed on the target database.

Related Topics

- *Oracle® Data Guard Broker*

Related Topics

- *Oracle® Data Guard Broker*

Related Topics

- *Oracle® Data Guard Concepts and Administration*

Related Topics

- *Oracle® Database Backup and Recovery User's Guide*

Clear Flashback Logs Periodically for Increased Fast Recovery Area Size Predictability

Fast recovery area management and database health are improved by automatically deleting flashback logs that are beyond the retention period.

The fast recovery area is critical for databases because it stores backups, online redo logs, archived redo logs, and flashback logs. Because many databases can all use the fast recovery area, multiple databases are impacted when the fast recovery area becomes full. This feature makes flashback space usage become predictable from a storage management perspective, since flashback uses no more space than is required by retention. It also allows you to control cumulative space pressure by adjusting the flashback retention.

Related Topics

- *Oracle® Database Backup and Recovery User's Guide*

New Parameters to Tune Automatic Outage Resolution with Oracle Data Guard

Oracle Data Guard automatic outage resolution can be tuned to fit your specific needs.

Oracle Data Guard has several processes, on the primary database and standby databases, which communicate with each other over the network to manage redo transport and archiving. In certain failure situations, network hangs, disconnects, and disk I/O issues, these processes can hang potentially causing delays in redo transport and gap resolution. Oracle Data Guard has an internal mechanism to detect these hung processes and terminate them thus allowing the normal outage resolution to occur. You can now use two new parameters, `DATA_GUARD_MAX_IO_TIME` and `DATA_GUARD_MAX_LONGIO_TIME`, to tune wait times for a specific Oracle Data Guard configuration based on the user network and disk I/O behavior.

Related Topics

- *Oracle® Database Reference*

Finer Granularity Supplemental Logging

Fine-grained supplemental logging provides a way for partial database replication users to disable supplemental logging for uninteresting tables so that even when supplemental logging is enabled in a database or schema level, there is no supplemental logging overhead for uninteresting tables.

Use of this feature significantly reduces the overhead in terms of resource usage and redo generation when only some of the tables in the database require supplemental logging, such as in an Oracle GoldenGate partial replication configuration. Supplemental logging was designed and implemented for logical standby databases or

for full database replication requirements. This adds unnecessary overhead in environments where only a subset of tables is being replicated.

Sharding

- [Support for Multi-Shard Query Coordinators on Shard Catalog Standby Databases](#)
- [Generation of Unique Sequence Numbers Across Shards](#)
- [Support for Multiple PDB Shards in the Same CDB](#)
- [Multiple Table Family Support for System-Managed Sharding](#)
- [Propagation of Parameter Settings Across Shards](#)

Support for Multi-Shard Query Coordinators on Shard Catalog Standby Databases

You enable a multi-shard query coordinator on the shard catalog's Oracle Active Data Guard standby databases.

Running a multi-shard query coordinator on the shard catalog active standby databases improves the scalability and availability of a multi-shard query workload, whereas before Oracle Database 19c, only the primary shard catalog database could be used as the multi-shard query coordinator.

Related Topics

- [Oracle® Database Using Oracle Sharding](#)

Generation of Unique Sequence Numbers Across Shards

You can generate globally unique sequence numbers across shards for any case in which a sequence object must be a single logical object across all shards of a sharded database.

You can use this functionality to generate globally unique IDs for non-primary key columns with unique constraints. This feature does not require you to manage the global uniqueness of a given non-primary key column in your application.

Related Topics

- [Oracle® Database Using Oracle Sharding](#)

Support for Multiple PDB Shards in the Same CDB

You can use more than one PDB in a CDB for shards or shard catalog databases, with certain restrictions. For example, this feature allows a CDB to contain shard PDBs from different sharded databases, each with its own separate shard catalog database.

When you have multiple PDBs in a CDB, customers and applications that require separate sharded databases can share the same system resources for cost reduction and ease of management.

Related Topics

- [Oracle® Database Using Oracle Sharding](#)

Multiple Table Family Support for System-Managed Sharding

You can create more than one table family in a sharded database, each of which can be sharded with a different sharding key.

Different applications that access different table families can now be hosted on one sharded database. This feature applies to system-managed sharded databases only.

Related Topics

- [Oracle® Database Using Oracle Sharding](#)

Propagation of Parameter Settings Across Shards

You can manage and propagate parameter settings to all of the database shards centrally from the shard catalog.. Before Oracle Database 19c, you had to configure `ALTER SYSTEM` parameter settings individually on each shard in a sharded database.

The ability to automatically propagate parameter settings to all of the shards from a central server is saves time and is less prone to error.

Related Topics

- [Oracle® Database Using Oracle Sharding](#)

Big Data and Data Warehousing

- [General](#)

General

- [Automatic Indexing](#)
- [SQL Diagnostics and Repair Enhancements](#)
- [Bitmap Based Count Distinct SQL Function](#)
- [Big Data and Performance Enhancements for In-Memory External Tables](#)
- [Automatic SQL Plan Management](#)
- [Real-Time Statistics](#)
- [High-Frequency Automatic Optimizer Statistics Collection](#)
- [Hybrid Partitioned Tables](#)

Automatic Indexing

The automatic indexing feature automates index management tasks, such as creating, rebuilding, and dropping indexes in an Oracle Database based on changes in the application workload.

This feature improves database performance by managing indexes automatically in an Oracle Database.

Related Topics

- [Oracle® Database Administrator's Guide](#)

SQL Diagnostics and Repair Enhancements

The SQL diagnostics and repair tools, such as SQL Test Case Builder and SQL Repair Advisor have been enhanced to provide better diagnosis and repair capabilities for managing problematic SQL statements.

These enhancements enable more effective diagnosis and repair of problematic SQL statements.

Related Topics

- *Oracle® Database Administrator's Guide*

Bitmap Based Count Distinct SQL Function

Use new bit vector SQL operators to speed up `COUNT DISTINCT` operations within a SQL query. To compute `COUNT(DISTINCT)` for numeric expressions, you can create a bit vector representation of the expressions and aggregate them before the final bit count. The resulting bit vector can be materialized, such as in a materialized view.

You can construct bit vectors by further grouping on a larger set of `GROUP BY` keys than targeted queries, so that you can use one materialized view to rewrite multiple `GROUP BY` queries with `COUNT(DISTINCT)` expressions by using `ROLLUP`.

In most scenarios, bit vector SQL functions combined with materialized views can provide significant performance improvements for queries with `COUNT(DISTINCT)` operations, which are common in data warehousing environments. The new operators are naturally evaluated in parallel and take advantage of hardware-optimized bitmap operations. By creating materialized views with bit vectors at lower-level aggregation levels, you can reuse the same materialized view to rewrite queries at higher level of aggregation levels by using `ROLLUP`.

Related Topics

- *Oracle® Database Data Warehousing Guide*

Big Data and Performance Enhancements for In-Memory External Tables

In-Memory external tables add support for `ORACLE_HIVE` and `ORACLE_BIGDATA` drivers, parallel query, Oracle Real Application Clusters, Oracle Active Data Guard, and on-demand population.

By using the new Big Data drivers, you avoid the cost and complexity of materializing data before populating it into the In-Memory Column Store (IM column store). You can use the SQL analytical capabilities of Oracle Database and Database In-Memory to analyze both internal and external data. Support for parallel query and full scan population means applications have fewer limitations when accessing data that resides outside the database.

Related Topics

- *Oracle® Database In-Memory Guide*

Automatic SQL Plan Management

Automatic SQL plan management resolves plan regressions without user intervention. For example, if high-load statements are performing suboptimally, then SQL plan

management evolve advisor can locate the statements automatically, and then test and accept the best plans.

SQL plan management searches for SQL statements in the Automatic Workload Repository (AWR). Prioritizing by highest load, it looks for alternative plans in all available sources, adding better-performing plans to the SQL plan baseline. Oracle Database also provides a plan comparison facility and improved hint reporting.

The impact of SQL statement performance regressions is significantly reduced using automation.

Related Topics

- *Oracle® Database SQL Tuning Guide*

Real-Time Statistics

Oracle Database automatically gathers online statistics during conventional data manipulation language (DML) operations.

Statistics can go stale between execution of `DBMS_STATS` statistics gathering jobs. By gathering some statistics automatically during DML operations, the database augments the statistics gathered by `DBMS_STATS`. Fresh statistics enable the optimizer to produce more optimal plans.

Related Topics

- *Oracle® Database SQL Tuning Guide*

High-Frequency Automatic Optimizer Statistics Collection

You can configure a lightweight, high-frequency automatic task that periodically gathers optimizer statistics for stale objects.

Statistics can go stale between executions of `DBMS_STATS` jobs. By gathering statistics more frequently, the optimizer can produce more optimal plans.

Related Topics

- *Oracle® Database SQL Tuning Guide*

Hybrid Partitioned Tables

The hybrid partitioned tables feature extends Oracle partitioning by enabling partitions to reside in both Oracle Database segments and in external files and sources. This feature significantly enhances the functionality of partitioning for Big Data SQL where large portions of a table can reside in external partitions.

Hybrid partitioned tables enable you to integrate internal partitions and external partitions into a single partition table. With this feature, you can also move non-active partitions to external files, such as Oracle Data Pump files, for a cheaper storage solution.

Related Topics

- *Oracle® Database VLDB and Partitioning Guide*

Database Overall

- [Automated Installation, Configuration and Patching](#)
- [AutoUpgrade and Database Utilities](#)

Automated Installation, Configuration and Patching

- [Ability to Create a Duplicate of an Oracle Database Using DBCA in Silent Mode](#)
- [Ability to Relocate a PDB to Another CDB Using DBCA in Silent Mode](#)
- [Ability to Create a PDB by Cloning a Remote PDB Using DBCA in Silent Mode](#)
- [Simplified Image-Based Oracle Database Client Installation](#)
- [Root Scripts Automation Support for Oracle Database Installation](#)
- [Support for Dry-Run Validation of Oracle Clusterware Upgrade](#)

Ability to Create a Duplicate of an Oracle Database Using DBCA in Silent Mode

You can now create a duplicate of an Oracle Database by using the `createDuplicateDB` command of Database Configuration Assistant (DBCA) in silent mode.

This feature enables developers to work on identical copies of an Oracle Database.

Related Topics

- [Oracle® Database Administrator's Guide](#)

Ability to Relocate a PDB to Another CDB Using DBCA in Silent Mode

You can now relocate a pluggable database (PDB) to another multitenant container database (CDB) by using the `relocatePDB` command of Database Configuration Assistant (DBCA) in silent mode.

This feature enables automating the PDB life cycle operation of relocating a PDB using DBCA in silent mode.

Related Topics

- [Oracle® Database Administrator's Guide](#)

Ability to Create a PDB by Cloning a Remote PDB Using DBCA in Silent Mode

You can now create a pluggable database (PDB) by cloning a remote PDB using the `createFromRemotePDB` parameter of the `createPluggableDatabase` command of Database Configuration Assistant (DBCA) in silent mode.

This feature enables automating the PDB life cycle operation of cloning a PDB using DBCA in silent mode.

Related Topics

- [Oracle® Database Administrator's Guide](#)

Simplified Image-Based Oracle Database Client Installation

Starting with Oracle Database 19c, the Oracle Database client software is available as an image file for download and installation. You must extract the image software into a directory where you want your Oracle home to be located, and then run the `runInstaller` script to start the Oracle Database client installation. Oracle Database client installation binaries continue to be available in the traditional format as non-image zip files.

As with Oracle Database and Oracle Grid Infrastructure image file installations, Oracle Database client image installations simplify Oracle Database client installations and ensure best practice deployments.

Related Topics

- [Oracle® Database Client Installation Guide for Linux](#)

Root Scripts Automation Support for Oracle Database Installation

Starting with Oracle Database 19c, the database installer, or setup wizard, provides options to set up permissions to run the root configuration scripts automatically, as required, during a database installation. You continue to have the option to run the root configuration scripts manually.

Setting up permissions for root configuration scripts to run without user intervention can simplify database installation and help avoid inadvertent permission errors.

Related Topics

- [Oracle® Database Installation Guide for Linux](#)

Support for Dry-Run Validation of Oracle Clusterware Upgrade

Starting with Oracle Grid Infrastructure 19c, the Oracle Grid Infrastructure installation wizard (`gridSetup.sh`) enables you to perform a dry-run mode upgrade to check your system's upgrade readiness.

In dry-run upgrade mode, the installation wizard performs all of the system readiness checks that it would perform in an actual upgrade and enables you to verify whether your system is ready for upgrade before you start the upgrade. This mode does not perform an actual upgrade. It helps anticipate potential problems with the system setup and avoid upgrade failures.

Related Topics

- [Oracle® Grid Infrastructure Installation and Upgrade Guide for Linux](#)

AutoUpgrade and Database Utilities

- [AutoUpgrade for Oracle Database](#)
- [Oracle Data Pump Ability to Exclude ENCRYPTION Clause on Import](#)
- [Oracle Data Pump Allows Tablespaces to Stay Read-Only During TTS Import](#)
- [Oracle Data Pump Support for Resource Usage Limitations](#)
- [Oracle Data Pump Test Mode for Transportable Tablespaces](#)

- [Oracle Data Pump Prevents Inadvertent Use of Protected Roles](#)
- [Oracle Data Pump Loads Partitioned Table Data One Operation](#)
- [Oracle Data Pump Allows Wildcards for Dump File in Object Store](#)
- [Oracle Data Pump Import Supports More Object Store Credentials](#)

AutoUpgrade for Oracle Database

AutoUpgrade enables you to upgrade one or many Oracle Database instances at the command-line, using a single command, and a single configuration file.

AutoUpgrade runs pre-upgrade tasks, performs automated fix-ups where needed, processes the database upgrade, and finishes the upgrade by completing post-upgrade tasks. It includes automatic retry and fallback, the option to schedule upgrades for future points in time, and the ability to set, change, or remove initialization parameters as desired.

AutoUpgrade significantly reduces the manual effort associated with database upgrades. It enables you to upgrade multiple databases at the same time. It can even take care of routine upgrades for databases that are not actively managed by skilled database administrators. By reducing the effort needed for upgrades, and by implementing recommended practices automatically during the upgrade process, AutoUpgrade makes the database upgrade process easier to complete, and reduces risk.

Related Topics

- [Oracle® Database Upgrade Guide](#)

Oracle Data Pump Ability to Exclude ENCRYPTION Clause on Import

There is a new transform parameter, `OMIT_ENCRYPTION_CLAUSE`, that causes Data Pump to suppress any encryption clauses associated with objects using encrypted columns.

Better Oracle Cloud migrations are now possible for non-cloud databases that have encrypted columns.

Related Topics

- [Oracle® Database Utilities](#)

Oracle Data Pump Allows Tablespaces to Stay Read-Only During TTS Import

You can now import tablespace files mounted on two different databases as long as the files are set as read-only.

A new option allows you to restore pre-12.2 default behavior, such that tablespace data files are read-only during the transportable tablespace import process. The benefit is that this allows a tablespace data file to be mounted on two databases, so long as it remains read-only. However, using this option requires that the source and target databases have exactly the same daylight savings time (DST) version because `TIMESTAMP WITH TIMEZONE` data is not adjusted upon import. Also, if you specify this parameter, then the database does not automatically rebuild tablespace bitmaps to reclaim space during import. This can make the import process faster at the expense of regaining free space within the tablespace data files.

Related Topics

- *Oracle® Database Utilities*

Oracle Data Pump Support for Resource Usage Limitations

The Oracle Data Pump parameter `MAX_DATAPUMP_JOBS_PER_PDB` is updated, and there is a new parameter, `MAX_DATAPUMP_PARALLEL_PER_JOB`.

`MAX_DATAPUMP_JOBS_PER_JOB` provides more control over the number of jobs that can be started in a multitenant container database environment: Default: 100, Range: 0 to 250, or Auto: 50% of `SESSIONS`. The `MAX_DATAPUMP_PARALLEL_PER_JOB` parameter enables you to obtain more control over the number of parallel workers that you can use for an individual Data Pump job.

These parameters provide you with more control over resource utilization when there are multiple users performing Data Pump jobs in a database environment.

Related Topics

- *Oracle® Database Utilities*

Oracle Data Pump Test Mode for Transportable Tablespaces

You can more easily determine how long an export takes, and discover unforeseen issues not reported by the closure check.

Test mode for Transportable Tablespaces (TTSs) performs a metadata-only export test using TTSs or full transportable export or import. It also removes the requirement for the source database tablespaces to be in read-only mode.

Related Topics

- *Oracle® Database Utilities*

Oracle Data Pump Prevents Inadvertent Use of Protected Roles

Oracle Data Pump prevents inadvertent use of protected roles during export and import with the new command-line parameter `ENABLE_SECURE_ROLES`.

Some Oracle roles require authorization. If you need to use these roles with Oracle Data Pump export and import you must explicitly enable them. The new `ENABLE_SECURE_ROLES` parameter is available for EXPDP and IMPDP clients, and for the Oracle Data Pump PL/SQL API. Starting with Oracle Database 19c, the default is `NO`.

Related Topics

- *Oracle® Database Utilities*

Oracle Data Pump Loads Partitioned Table Data One Operation

Oracle Data Pump can import table data in all partitions of a table as one operation instead of separate operations for each partition.

`GROUP_PARTITION_TABLE_DATA`, a new value for the Import `DATA_OPTIONS` command line parameter, changes Oracle Data Pump default behavior by importing table data in all partitions of a table as one operation. This parameter is useful when

you do not want the default Import behavior that imports each table partition as a separate operation. Import chooses the default. For instance, you can use this parameter when there is a possibility that a table could move to a different partition as part of loading a table. The default is also used when the table was not created by the Import operation.

Related Topics

- [Oracle® Database Utilities](#)

Oracle Data Pump Allows Wildcards for Dump File in Object Store

Oracle Data Pump simplifies importing multiple dump files into Oracle Autonomous Database by allowing wildcards for URL-based dump file names.

When you need to import multiple dump files from the object store service, wildcards in URL-based dump file name can simplify the import command for Oracle Autonomous Database. It can reduce typing and lessen the possibility of a misspelled a dump file name. Do not use a wildcard character in the bucket-name component.

Related Topics

- [Oracle® Database Utilities](#)

Oracle Data Pump Import Supports More Object Store Credentials

Oracle Data Pump import supports object store credentials beyond the `DEFAULT_CREDENTIAL` with a new `CREDENTIAL` parameter for Oracle Autonomous Database.

Oracle Data Pump import is no longer constrained to using the `DEFAULT_CREDENTIAL` in Oracle Autonomous Database. Starting with Oracle Database 19c (and backported to Oracle Database release 18c, version 18.3) the new IMPDP client CLI `CREDENTIAL` parameter accepts any Oracle Cloud Infrastructure (OCI) Object Storage credential created in the Oracle Autonomous Database. Data Pump validates whether the credential exists and the user has access to read the credential. Any errors are returned back to the IMPDP client.

Related Topics

- [Oracle® Database Utilities](#)

Diagnosability

- [General](#)

General

- [Oracle Trace File Analyzer Support for Using an External SMTP Server for Notifications](#)
- [Oracle Cluster Health Advisor Integration into Oracle Trace File Analyzer](#)
- [Oracle Trace File Analyzer REST API Support](#)
- [Oracle Trace File Analyzer Search Extended to Support Metadata Searches](#)
- [Oracle ORAchk and Oracle EXAchk REST Support](#)

- [Oracle ORAck and Oracle EXAck Support for Encrypting Collection Files](#)
- [Oracle ORAck and Oracle EXAck Support for Remote Node Connections Without Requiring Passwordless SSH](#)
- [Oracle ORAck and Oracle EXAck Now Show Only the Most Critical Checks by Default](#)
- [Oracle Trace File Analyzer Supports New Service Request Data Collections](#)

Oracle Trace File Analyzer Support for Using an External SMTP Server for Notifications

In Oracle Database 19c, you can use an external Simple Mail Transfer Protocol (SMTP) server to receive Oracle Trace File Analyzer notifications.

In earlier releases of Oracle Trace File Analyzer, to deliver email notifications of alerts, you had to have monitored hosts configured with local sendmail or SMTP support. With external SMTP server notification support, Oracle Trace File Analyzer deployments can leverage complete notification functionality, helping to minimize downtime, and maximizing availability.

Related Topics

- [Oracle® Autonomous Health Framework User's Guide](#)

Oracle Cluster Health Advisor Integration into Oracle Trace File Analyzer

Oracle Trace File Analyzer now integrates with Oracle Cluster Health Advisor, and consumes the problem events that Oracle Cluster Health Advisor detects.

When Oracle Cluster Health Advisor detects a problem event, Oracle Trace File Analyzer automatically triggers the relevant diagnostic collection and sends an email notification. You can configure email notification through the standard Oracle Trace File Analyzer notification process.

Oracle Cluster Health Advisor provides early warnings for Oracle Real Application Clusters (Oracle RAC) database and cluster node related performance issues. Oracle Trace File Analyzer sends email notifications with root cause analysis and corrective recommendations, which enables you to prevent application performance and availability issues proactively.

Related Topics

- [Oracle® Autonomous Health Framework User's Guide](#)

Oracle Trace File Analyzer REST API Support

Oracle Trace File Analyzer now includes REpresentational State Transfer (REST) support, which enables invocation and query over HTTPS.

Oracle REST Data Services (ORDS) is included within the installation to facilitate REST support. REST supports printing details, starting a diagnostic collection, and downloading collections.

The REST interface enables you to configure remote management, and automate data center operations. Oracle Trace File Analyzer when operating through REST APIs supports easy integration into your operations framework and thus improves diagnostic efficiency and reduces recovery time.

Related Topics

- *Oracle® Autonomous Health Framework User's Guide*

Oracle Trace File Analyzer Search Extended to Support Metadata Searches

Starting in this release, metadata stored in the Oracle Trace File Analyzer index is searchable using `tfactl`.

Oracle Trace File Analyzer searches log and trace file metadata using JavaScript Object Notation (JSON) formatted name-value pairs representing data types and events.

The ability to search log and trace file metadata is essential to minimize downtime and maximize availability and to efficiently diagnose and triage issues, especially the recurring issues across instances and nodes. In earlier releases of Oracle Trace File Analyzer, the search function was limited to log and trace file strings.

Related Topics

- *Oracle® Autonomous Health Framework User's Guide*

Oracle ORAchk and Oracle EXAchk REST Support

Oracle ORAchk and Oracle EXAchk now include REpresentational State Transfer (REST) support, which enables invocation and query over HTTPS.

Oracle REST Data Services (ORDS) is included within the installation to facilitate REST support. The REST interface enables you to configure remote management, and automate data center operations. Oracle ORAchk and Oracle EXAchk, when operating through REST APIs, support easy integration into your operations framework and thus improve diagnostic efficiency and reduce recovery time.

Related Topics

- *Oracle® Autonomous Health Framework User's Guide*

Oracle ORAchk and Oracle EXAchk Support for Encrypting Collection Files

Oracle ORAchk and Oracle EXAchk diagnostic collection files may contain sensitive data. Starting in this release, you can encrypt and decrypt diagnostic collection ZIP files and protect them with a password.

Oracle ORAchk and Oracle EXAchk collections and their reports can include sensitive data. When you email or transfer these reports to repositories, it is critical that only the intended recipients can view the sensitive data. To prevent leaks, you can restrict access to sensitive data by encrypting the diagnostic collections and protecting them with a password. This feature is available only on Linux and Solaris platforms.

Related Topics

- *Oracle® Autonomous Health Framework User's Guide*

Oracle ORAchk and Oracle EXAchk Support for Remote Node Connections Without Requiring Passwordless SSH

Starting in this release, you can configure Oracle ORAchk and Oracle EXAchk to autogenerate the private key files for the remote nodes. Alternatively, you can provide a private key.

You can perform operations remotely to centrally manage many database servers or clusters. In many cases, corporate policies prevent passwordless Secure Shell (SSH) configuration. Using the private key authentication, you can run Oracle ORAchk and Oracle EXAchk remotely in these deployments and improve operational efficiency. In earlier releases of Oracle ORAchk and Oracle EXAchk, remotely running Oracle ORAchk or Oracle EXAchk required configuration of passwordless SSH between the remote nodes.

Related Topics

- *Oracle® Autonomous Health Framework User's Guide*

Oracle ORAchk and Oracle EXAchk Now Show Only the Most Critical Checks by Default

Oracle ORAchk and Oracle EXAchk generate reports and show only the most critical checks by default.

The critical checks are those that have the most severe potential effect. Oracle ORAchk and Oracle EXAchk still run all other checks and include them in the report. You can view the checks by selecting the appropriate option under the *Show checks* with the following status control.

In earlier releases of Oracle ORAchk and Oracle EXAchk, reports contained over one hundred checks and thus made the analysis more time-consuming. With the most critical checks, you can analyze the reports efficiently, and quickly resolve critical problems and prevent downtime or performance issues.

Related Topics

- *Oracle® Autonomous Health Framework User's Guide*

Oracle Trace File Analyzer Supports New Service Request Data Collections

This release adds additional database Service Request Data Collections (SRDCs) that cover more ORA errors and problems in the infrastructure such as Oracle Automatic Storage Management (Oracle ASM), Oracle Automatic Storage Management Cluster File System (Oracle ACFS), listeners, auditing, and Recovery Manager (RMAN).

When operations or Oracle Database issues occur that require Oracle Support Services, it is important that you collect and send all of the data and logs necessary to diagnose and resolve the issue in one compact complete archive. SRDCs simplify the collection of required logs and data for specific issues.

Related Topics

- *Oracle® Autonomous Health Framework User's Guide*

Performance

- [General](#)

General

- [SQL Quarantine](#)
- [Database In-Memory Wait on Populate](#)
- [Resource Manager Automatically Enabled for Database In-Memory](#)
- [Memoptimized Rowstore Fast Ingest](#)
- [Automatic Database Diagnostic Monitor \(ADDM\) Support for Pluggable Databases \(PDBs\)](#)
- [Resource Manager Automatically Enabled for Database In-Memory](#)
- [High-Frequency SQL Plan Management Evolve Advisor Task](#)
- [Workload Capture and Replay in a PDB](#)

SQL Quarantine

SQL statements that are terminated by Oracle Database Resource Manager due to their excessive consumption of CPU and I/O resources are automatically quarantined. The execution plans associated with the terminated SQL statements are quarantined to prevent them from being executed again.

This feature protects an Oracle Database from performance degradation by preventing execution of SQL statements that excessively consume CPU and I/O resources.

Related Topics

- [Oracle® Database Administrator's Guide](#)

Database In-Memory Wait on Populate

The `DBMS_INMEMORY_ADMIN.POPULATE_WAIT` function waits until objects at the specified priority have been populated to the specified percentage.

The new function ensures that the specified In-Memory objects have been populated before allowing application access. For example, a database might contain a number of In-Memory tables with a variety of priority settings. In a restricted session, you can use the `POPULATE_WAIT` function to ensure that every In-Memory table is completely populated. Afterward, you can disable the restricted session so that the application is guaranteed to query only In-Memory representations of the tables.

Related Topics

- [Oracle® Database In-Memory Guide](#)

Resource Manager Automatically Enabled for Database In-Memory

When `INMEMORY_SIZE` is greater than 0, Oracle Database Resource Manager is automatically enabled.

The Resource Manager is required to take advantage of In-Memory Dynamic Scans. Because the Resource Manager is automatically enabled when Database In-Memory is enabled, you receive the benefits of enhanced performance and automatic management for CPU resource allocation.

Memoptimized Rowstore Fast Ingest

The fast ingest functionality of Memoptimized Rowstore enables fast data inserts into an Oracle Database from applications, such as Internet of Things (IoT) applications that ingest small, high volume transactions with a minimal amount of transactional overhead. The insert operations that use fast ingest temporarily buffer the data in the large pool before writing it to disk in bulk in a deferred, asynchronous manner.

Using the rich analytical features of Oracle Database, you can now perform data analysis more effectively by easily integrating data from high-frequency data streaming applications with your existing application data.

Related Topics

- *Oracle® Database Performance Tuning Guide*

Automatic Database Diagnostic Monitor (ADDM) Support for Pluggable Databases (PDBs)

You can now use Automatic Database Diagnostic Monitor (ADDM) analysis for pluggable databases (PDBs) in a multitenant environment.

ADDM analysis at a PDB level enables you to tune a PDB effectively for better performance.

Related Topics

- *Oracle® Database Performance Tuning Guide*

Resource Manager Automatically Enabled for Database In-Memory

When `INMEMORY_SIZE` is greater than 0, Oracle Database Resource Manager is automatically enabled.

The Resource Manager is required to take advantage of In-Memory Dynamic Scans. Because the Resource Manager is automatically enabled when Database In-Memory is enabled, you receive the benefits of enhanced performance and automatic management for CPU resource allocation.

Related Topics

- *Oracle® Database In-Memory Guide*

Related Topics

- *Oracle® Database SQL Tuning Guide*

High-Frequency SQL Plan Management Evolve Advisor Task

You can configure the Automatic SPM Evolve Advisor task to run every hour, outside of the standard maintenance window.

By evolving SQL plan baselines more frequently, the optimizer can correct performance regressions more quickly and enforce more optimal SQL execution plans.

Related Topics

- [Oracle® Database SQL Tuning Guide](#)

Workload Capture and Replay in a PDB

Oracle Real Application Testing was designed to capture and replay multitenant databases at the root multitenant container database (CDB) level. Starting with Oracle Database Release 19c, you can capture and replay the workload from within an individual pluggable database (PDB).

This enhancement enables you to capture and replay workloads at the PDB level. This leads to better testing, less downtime, and more effective and efficient change control.

Related Topics

- [Oracle® Database Testing Guide](#)

RAC and Grid

- [General](#)

General

- [Parity Protected Files](#)
- [Secure Cluster Communication](#)
- [Automated PDB Relocation](#)
- [Zero-Downtime Oracle Grid Infrastructure Patching](#)
- [Automated Transaction Draining for Oracle Grid Infrastructure Upgrades](#)
- [Oracle Restart Patching and Upgrading](#)
- [Colocation Tag for Client Routing](#)
- [Optional Install for the Grid Infrastructure Management Repository](#)
- [Resupport of Direct File Placement for OCR and Voting Disks](#)
- [Dynamic Services Fallback Option](#)

Parity Protected Files

You cannot use parity protection for write-once files in Oracle Database Automatic Storage Management (Oracle ASM). Write-once files are files such as archive logs and backup sets.

A great deal of space is consumed when two or three way Oracle ASM mirroring is used for files associated with database backup operations. Backup files are write-once files, and this feature allows parity protection for protection rather than conventional mirroring. Considerable space savings are the result.

Related Topics

- [Oracle® Automatic Storage Management Administrator's Guide](#)

Secure Cluster Communication

Secure Cluster Communication protects the cluster interconnect from common security threats when used together with Single Network Support. Secure Cluster Communication includes message digest mechanisms, protection against fuzzing, and uses Transport Layer Security (TLS) to provide privacy and data integrity between the cluster members.

The increased security for the cluster interconnect is invoked automatically as part of a new Oracle Grid Infrastructure 19c deployment or an upgrade to Oracle Grid Infrastructure 19c. Database administrators or cluster administrators do not need to make any configuration changes for this feature.

Related Topics

- *Oracle® Clusterware Administration and Deployment Guide*

Automated PDB Relocation

In Oracle Grid Infrastructure, you can use Fleet Patching and Provisioning to automate relocation of a pluggable database (PDB) from one multitenant container database (CDB) to another.

You can patch individual PDBs more quickly and without exposing other PDBs to the changes that a patch would bring.

Related Topics

- *Oracle® Clusterware Administration and Deployment Guide*

Zero-Downtime Oracle Grid Infrastructure Patching

Zero-downtime Oracle Grid Infrastructure Patching enables patching of Oracle Grid Infrastructure without interrupting database operations. Patches are applied out-of-place and in a rolling fashion, with one node being patched at a time, while the database instances on the node remain operational. Zero-Downtime Oracle Grid Infrastructure Patching supports Oracle Real Application Clusters (Oracle RAC) databases on clusters with two or more nodes.

Zero-Downtime Grid Infrastructure Patching significantly increases database availability by allowing you to perform a rolling patch of Oracle Grid Infrastructure without interrupting database operations on the node being patched and without affecting capacity or performance on those database instances.

Related Topics

- *Oracle® Clusterware Administration and Deployment Guide*

Automated Transaction Draining for Oracle Grid Infrastructure Upgrades

Automated Transaction Draining for Oracle Grid Infrastructure Upgrades provides automatic draining of transactions against the database instances, one node at a time, in a rolling fashion, according to the database service configurations. Transaction draining capabilities are an integral part of the database service design and are now automatically integrated into the application of rolling Oracle Grid Infrastructure patches.

Automated and coordinated draining of database transactions during rolling patch applications, using Fleet Patching and Provisioning, reduces the impact of patching operations. Once user transactions are drained, patching operations for a particular node on a cluster are completed. The instance and services are then restarted locally and new connections are established before the patching operation rolls on to the next node in the cluster.

Related Topics

- *Oracle® Clusterware Administration and Deployment Guide*

Oracle Restart Patching and Upgrading

Use Fleet Patching and Provisioning to patch and upgrade Oracle Restart. In previous releases, Oracle Restart environments required you to perform patching and upgrade operations, often involving manual intervention. Fleet Patching and Provisioning automates these procedures.

Using Fleet Patching and Provisioning to patch and upgrade Oracle Restart automates and standardizes the processes that are implemented in Oracle Real Application Clusters (Oracle RAC) database installations. This also reduces operational demands and risks, especially for large numbers of Oracle Restart deployments.

Related Topics

- *Oracle® Clusterware Administration and Deployment Guide*

Colocation Tag for Client Routing

The `COLOCATION_TAG` parameter is an alphanumeric string that you can use with the `CONNECT_DATA` parameter of the Transparent Network Substrate (TNS) connect string. When you set the `COLOCATION_TAG` parameter, it attempts to route clients with the same `COLOCATION_TAG` to the same database instance.

Colocation of sessions on the same instance can help decrease inter-instance communication and thereby increase performance for workload that benefits from being executed in the same instance.

Related Topics

- *Oracle® Database Database Net Services Reference*

Optional Install for the Grid Infrastructure Management Repository

Starting with Oracle Grid Infrastructure 19c, the Grid Infrastructure Management Repository (GIMR) is optional for new installations of Oracle Standalone Cluster. Oracle Domain Services Cluster still requires the installation of a GIMR as a service component.

The data contained in the GIMR is the basis for preventative diagnostics based on applied Machine Learning and helps to increase the availability of Oracle Real Application Clusters (Oracle RAC) databases. Having an optional installation for the GIMR allows for more flexible storage space management and faster deployment, especially during the installation of test and development systems.

Related Topics

- *Oracle® Database Upgrade Guide*

Resupport of Direct File Placement for OCR and Voting Disks

Starting with Oracle Grid Infrastructure 19c, the desupport for direct Oracle Cluster Registry (OCR) and voting disk file placement on shared file systems is rescinded for Oracle Standalone Clusters. For Oracle Domain Services Clusters, the requirement to place OCR and voting files in Oracle Automatic Storage Management (Oracle ASM) on top of files hosted on shared file systems and used as Oracle ASM disks remains.

In Oracle Grid Infrastructure 12c Release 2 (12.2), Oracle announced that it would no longer support the placement of the OCR and voting files for Oracle Grid Infrastructure directly on a shared file system. This desupport is now rescinded. Starting with Oracle Grid Infrastructure 19c (version 19.3), with Oracle Standalone Clusters, you can again place OCR and voting disk files directly on shared file systems.

Related Topics

- [Oracle® Database Upgrade Guide](#)

Dynamic Services Fallback Option

For a dynamic database service that is placed using "preferred" and "available" settings, you can now specify that this service should fall back to a "preferred" instance when it becomes available if the service failed over to an available instance.

The Dynamic Services Fallback Option allows for more control in placing dynamic database services and ensures that a given service is available on a preferred instance as long as possible.

Related Topics

- [Oracle® Real Application Clusters Administration and Deployment Guide](#)

Security

- [General](#)

General

- [New ALTER SYSTEM Clause FLUSH PASSWORDFILE_METADATA_CACHE](#)
- [Transparent Online Conversion Support for Auto-Renaming in Non-Oracle-Managed Files Mode](#)
- [Support for Additional Algorithms for Offline Tablespace Encryption](#)
- [Key Management of Encrypted Oracle-Managed Tablespaces in Transparent Data Encryption](#)
- [Support for Host Name-Based Partial DN Matching for Host Certificates](#)
- [New PDB_GUID Audit Record Field for SYSLOG and the Windows Event Viewer](#)
- [New EVENT_TIMESTAMP_UTC Column in the UNIFIED_AUDIT_TRAIL View](#)
- [Passwords Removed from Oracle Database Accounts](#)
- [Signature-Based Security for LOB Locators](#)
- [Unified Auditing Top-Level Statements](#)

- [Privilege Analysis Now Available in Oracle Database Enterprise Edition](#)
- [Support for Oracle Native Encryption and SSL Authentication for Different Users Concurrently](#)
- [Ability to Grant or Revoke Administrative Privileges to and from Schema-Only Accounts](#)
- [Automatic Support for Both SASL and Non-SASL Active Directory Connections](#)
- [Database Vault Operations Control for Infrastructure Database Administrators](#)
- [Database Vault Command Rule Support for Unified Audit Policies](#)

New ALTER SYSTEM Clause FLUSH PASSWORDFILE_METADATA_CACHE

The ALTER SYSTEM clause FLUSH PASSWORDFILE_METADATA_CACHE refreshes the metadata cache with the latest details of the database password file. Querying the V\$PASSWORDFILE_INFO view retrieves the latest details of the database password file.

This functionality is useful when the database password file name or location is changed, and the metadata cache needs to be refreshed with the details of the updated database password file.

Related Topics

- [Oracle® Database Administrator's Guide](#)

Transparent Online Conversion Support for Auto-Renaming in Non-Oracle-Managed Files Mode

Starting with this release, in a Transparent Data Encryption online conversion in non-Oracle-managed files mode, you are no longer forced to include the FILE_NAME_CONVERT clause in the ADMINISTER KEY MANAGEMENT SQL statement. The file name retains its original name.

This enhancement helps prevent you from having to rename files to the original name, sometimes missing files.

Related Topics

- [Oracle® Database Advanced Security Guide](#)

Support for Additional Algorithms for Offline Tablespace Encryption

In previous releases, only the AES128 encryption algorithm was supported for offline tablespace encryption. This release adds support for the AES192 and AES256 encryption algorithms, as well as for the ARIA, GOST, and 3DES encryption algorithms for offline tablespace encryption.

This enhancement helps in scenarios in which you have concerns about auxiliary space usage required by online tablespace encryption.

Related Topics

- [Oracle® Database Advanced Security Guide](#)

Key Management of Encrypted Oracle-Managed Tablespaces in Transparent Data Encryption

In this release, a closed Transparent Data Encryption (TDE) encryption keystore has no impact on internal operations to Oracle-managed tablespaces.

Internal processes can access a keystore when the keystore is closed, which allows the internal process to continue and successfully complete by using an intermediate key that is derived from the TDE master encryption key, while the TDE keystore is closed or is otherwise unavailable.

Closing the TDE keystore has no effect on queries of an encrypted `SYSTEM`, `SYSAUX`, `TEMP`, and `UNDO` tablespace, unlike queries of a user created tablespace, which continue to return an `ORA-28365 wallet is not open` error when the TDE keystore is closed.

User initiated operations such as decrypt on any encrypted Oracle-managed tablespace still require the TDE keystore to be in the `OPEN` state.

Related Topics

- *Oracle® Database Advanced Security Guide*

Support for Host Name-Based Partial DN Matching for Host Certificates

There is new support for partial distinguished name (DN) matching that adds the ability for the client to further verify the server certificate.

The earlier ability to perform a full DN match with the server certificate during the Secure Sockets Layer (SSL) handshake is still supported. The client supports both full and partial DN matching. If the server DN matching is enabled, then partial DN matching is the default.

Allowing partial and full DN matching for certificate verification enables more flexibility based on how the certificates were created.

Related Topics

- *Oracle® Database Security Guide*

New PDB_GUID Audit Record Field for SYSLOG and the Windows Event Viewer

The audit record fields for `SYSLOG` and the Windows Event Viewer now include a new field, `PDB_GUID`, to identify the pluggable database (PDB) associated with a unified audit trail record.

In a multitenant container database (CDB) deployment, the pluggable database that generated a unified audit trail record must be identified in the audit trail. The new field captures this information starting with this release. The data type is `VARCHAR2`.

Related Topics

- *Oracle® Database Security Guide*

New EVENT_TIMESTAMP_UTC Column in the UNIFIED_AUDIT_TRAIL View

The new `EVENT_TIMESTAMP_UTC` column appears in the `UNIFIED_AUDIT_TRAIL` view. Query the `UNIFIED_AUDIT_TRAIL` view based on the `EVENT_TIMESTAMP_UTC` column in the `WHERE` clause. The new column helps partition pruning, improving the read performance of the `UNIFIED_AUDIT_TRAIL` view.

Related Topics

- *Oracle® Database Security Guide*

Passwords Removed from Oracle Database Accounts

Most of the Oracle Database supplied schema-only accounts now have their passwords removed to prevent users from authenticating to these accounts.

This enhancement does not affect the sample schemas. Sample schemas are still installed with their default passwords.

Administrators can still assign passwords to the default schema-only accounts. Oracle recommends changing the schemas back to a schema-only account afterward.

The benefit of this feature is that administrators no longer have to periodically rotate the passwords for these Oracle Database provided schemas. This feature also reduces the security risk of attackers using default passwords to hack into these accounts.

Related Topics

- *Oracle® Database Security Guide*

Signature-Based Security for LOB Locators

Starting with this release, you can configure signature-based security for large object (LOB) locators.

This feature strengthens the security of Oracle Database LOBs, particularly when using instances of LOB data types (CLOB and BLOB) in distributed environments.

LOB signature keys are in both multitenant pluggable databases (PDBs) or in standalone, non-multitenant databases. You can enable the encryption of the LOB signature key credentials by executing the `ALTER DATABASE DICTIONARY ENCRYPT CREDENTIALS` SQL statement; otherwise, the credentials are stored in obfuscated format. If you choose to store the LOB signature key in encrypted format, then the database or PDB must have an open Transparent Data Encryption (TDE) keystore.

Related Topics

- *Oracle® Database Security Guide*

Unified Auditing Top-Level Statements

The unified auditing top-level statements feature enables you to audit top-level user (direct user) activities in the database without collecting indirect user activity audit data.

You can use this feature to audit only the events generated by top-level users, without the overhead of creating audit records for indirect SQL statements. Top-level statements are SQL statements that users directly issue. These statements are important for both security and compliance. Often SQL statements that run from within PL/SQL procedures or functions are not considered top level, so they may be less relevant for auditing purposes.

Related Topics

- *Oracle® Database Security Guide*

Privilege Analysis Now Available in Oracle Database Enterprise Edition

Privilege analysis is now available as part of Oracle Database Enterprise Edition.

Privilege analysis runs dynamic analysis of users and applications to find privileges and roles that are used and unused. Privilege analysis reduces the work to implement least privilege best practices by showing you exactly what privileges are used and not used by each account. Privilege analysis is highly performant and is designed to work in test, development, and production development databases.

As part of this change, the documentation for privilege analysis has moved from the Oracle Database Vault Administrator's Guide to the Oracle Database Security Guide.

Related Topics

- *Oracle® Database Security Guide*

Support for Oracle Native Encryption and SSL Authentication for Different Users Concurrently

In previous releases, Oracle Database prevented the use of Oracle native encryption (also called Advanced Networking Option (or ANO) encryption) and Secure Sockets Layer (SSL) authentication together.

For example, if you set both the `SQLNET.ENCRYPTION_CLIENT` parameter on the client and the `SQLNET.ENCRYPTION_SERVER` parameter on the server to `REQUIRED`, and a TCP/IP with SSL (TCPS) listener is used, then you receive the `ORA-12696 Double Encryption Turned On, login disallowed` error. Starting with this release, you can set the new `SQLNET.IGNORE_ANO_ENCRYPTION_FOR_TCPS` parameter to `TRUE`. This setting ignores the `SQLNET.ENCRYPTION_CLIENT` or `SQLNET.ENCRYPTION_SERVER` when a TCPS client is used and either of these two parameters are set to `REQUIRED`.

Related Topics

- *Oracle® Database Security Guide*

Ability to Grant or Revoke Administrative Privileges to and from Schema-Only Accounts

You can grant administrative privileges, such as `SYSOPER` and `SYSBACKUP`, to schema-only (passwordless) accounts.

Unused and rarely accessed database user accounts with administrative privileges can now become schema-only accounts. This enhancement prevents administrators from having to manage the passwords of these accounts.

Related Topics

- *Oracle® Database Security Guide*

Automatic Support for Both SASL and Non-SASL Active Directory Connections

Starting with this release, support is available for both Simple Authentication and Security Layer (SASL) and Transport Layer Security (TLS) binds for Microsoft Active Directory connections.

For centrally managed users, the Oracle Database initially tries to connect to Active Directory using SASL bind. If the Active Directory server rejects the SASL bind connection, then the Oracle Database automatically attempts the connection again without SASL bind but still secured with TLS.

The Active Directory administrator is responsible for configuring the connection parameters for Active Directory server, but does not need to configure the database to match this new Active Directory connection enhancement. The database automatically adjusts from using SASL to not using SASL bind.

Related Topics

- *Oracle® Database Security Guide*

Database Vault Operations Control for Infrastructure Database Administrators

In a multitenant database, you can now use Oracle Database Vault to block common users (infrastructure database administrators, for example) from accessing local data in pluggable databases (PDBs).

This enhancement prevents common users from accessing local data that resides on a PDB. It enables you to store sensitive data for your business applications and to allow operations to manage the database infrastructure without having to access sensitive customer data.

Related Topics

- *Oracle® Database Vault Administrator's Guide*

Database Vault Command Rule Support for Unified Audit Policies

You can now create Oracle Database Vault command rules for unified audit policies.

You can use command rules to differentiate which users can enable and disable unified audit policies under specific conditions. For example, the `HR` auditor can enable or disable the `HR` Unified Audit policy, while the Finance auditor can do the same for the Finance Unified Audit policy without being able to modify the other policy.

Related Topics

- *Oracle® Database Vault Administrator's Guide*