

Oracle® Database

Database New Features Guide



19c
F14269-01
April 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
Social Sign-In Authentication	1-1
REST Enabled SQL Support	1-2
Improved Create Application Wizard	1-2
Improved Create Page Wizard	1-3
Web Source Modules	1-3
New REST Workshop	1-4
General	1-4
Application Continuity for Java: New States Management	1-4
Application Continuity for Java: Declarative Request Demarcation	1-4
Oracle Network Log File Segmentation	1-5
SQL*Net: Auto-Detection of Support for Out-of-Band Breaks	1-5
Java	1-5
Java Library for Reactive Streams Ingestion	1-5
JSON	1-6
Materialized View Support for Queries containing JSON_TABLE	1-6
SQL/JSON Syntax Simplifications	1-6
New SQL/JSON Function JSON_SERIALIZE and JSON Data Guide Support for GeoJSON Data	1-6
JSON Update Operations	1-7
JSON-Object Mapping	1-7
SQL	1-7
DISTINCT option for LISTAGG aggregate	1-7
Availability	1-7

General	1-8
Simplified Database Parameter Management in a Broker Configuration	1-8
Dynamically change Fast-Start Failover (FSFO) target	1-8
Observe only mode for Broker's Fast_Start Failover (FSFO)	1-9
Flashback Standby database when Primary database is flashed back	1-9
Propagate Restore Points from Primary to Standby site	1-9
Oracle Data Guard Multi-Instance Redo Apply works with the In-Memory Column Store	1-10
Active Data Guard DML Redirection	1-10
PDB Recovery catalog	1-10
Clear Flashback logs periodically for increased FRA size predictability	1-10
New Parameters for tuning automatic outage resolution with Data Guard	1-11
Finer granularity Supplemental Logging	1-11
Sharding	1-11
Propagation of Parameter Settings Across Shards	1-12
Support for Multiple PDB Shards in the Same CDB	1-12
Multiple Table Family Support for System-Managed Sharding	1-12
Support for Multi-Shard Query Coordinators on Shard Catalog Standby Databases	1-12
Generation of Unique Sequence Numbers Across Shards	1-13
Big Data and Data Warehousing	1-13
General	1-13
SQL Diagnostics and Repair Enhancements	1-13
Automatic Indexing	1-14
Bitmap based count distinct SQL Function	1-14
Big Data and Performance Enhancements for In-Memory External Tables	1-14
Automatic Resolution of SQL Plan Regressions	1-15
Real-Time Statistics	1-15
High-Frequency Automatic Optimizer Statistics Collection	1-15
Hybrid Partitioned Tables	1-15
Database Overall	1-16
Automated install, config, and patch	1-16
Ability to Create a Duplicate of an Oracle Database Using DBCA in Silent Mode	1-16
Ability to Create a PDB by Cloning a Remote PDB Using DBCA in Silent Mode	1-16
Ability to Relocate a PDB to Another CDB Using DBCA in Silent Mode	1-16
Simplified Image Based Oracle Database Client Installation	1-17
Root Scripts Automation Support for Oracle Database Installation	1-17
Support for Dry-Run Validation of Oracle Clusterware Upgrade	1-17
Automated upgrade, migration and utilities	1-17
Oracle Data Pump Ability to Exclude ENCRYPTION Clause on Import	1-18

Oracle Data Pump Allows Tablespaces to Stay Read-Only During TTS Import	1-18
Oracle Data Pump Test Mode for Transportable Tablespaces	1-18
Oracle Data Pump Support for Resource Usage Limitations	1-18
General	1-19
Oracle Data Pump Prevents Inadvertent Use of Protected Roles	1-19
Oracle Data Pump Allows Wildcards for Dumpfile 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 REST API Support	1-20
Oracle Trace File Analyzer Search Extended to Support Metadata Searches	1-20
Oracle Trace File Analyzer Supports New Service Request Data Collections	1-20
Oracle ORAchk and Oracle EXAchk Support for Encrypting Collection Files	1-21
Oracle ORAchk and Oracle EXAchk REST Support	1-21
Oracle Cluster Health Advisor Integration into Oracle Trace File Analyzer	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 Support for Using an External SMTP Server for Notifications	1-22
Performance	1-23
General	1-23
SQL Quarantine	1-23
Resource Manager Automatically Enabled for Database In-Memory	1-23
Database In-Memory Wait on Populate	1-23
Memoptimized Rowstore - Fast Ingest	1-24
Automatic Database Diagnostic Monitor (ADDM) Support for Pluggable Databases (PDBs)	1-24
Real-Time SQL Monitoring for Developers	1-24
Workload Capture and Replay in a PDB	1-24
RAC and Grid	1-25
General	1-25
Parity Protected Files	1-25
Automated PDB Relocation	1-25
Automated Transaction Draining for Oracle Grid Infrastructure Upgrades	1-25
Oracle Restart Patching and Upgrading	1-26
Zero-Downtime Oracle Grid Infrastructure Patching	1-26
Security	1-26
General	1-26

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

2 Oracle Database 19c Release 19.3 New Features

AutoUpgrade and Database Utilities	2-1
Database Utilities	2-1
Oracle Data Pump loads partitioned table data in one operation	2-1
Upgrade	2-1
AutoUpgrade for Oracle Database	2-1
RAC and Grid	2-2
General	2-2
Colocation Tag for Client Routing	2-2
Optional Install for the Grid Infrastructure Management Repository	2-2
Resupport of Direct File Placement for OCR and Voting Disks	2-3
Dynamic Services Fallback Option	2-3

Preface

This document describes new features implemented in Oracle Database 18c.

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

Audience

Oracle Database New Features Guide is addressed to people familiar with previous releases of Oracle Database who would like to become familiar with features, options, and enhancements that are new in this release of the database.

Documentation Accessibility

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

Access to Oracle Support

Oracle 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 18c 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](#)
- [Java](#)
- [JSON](#)
- [SQL](#)

Application Express

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

Social Sign-In Authentication

Social Sign-In preconfigured authentication scheme supports authentication with Google, Facebook, and other social network that supports OpenIDConnect or OAuth2 standards.

Social Sign-In authentication is primarily useful for the following use cases:

- Your application is internet facing and you expect an unknown number of users from social networks to use your application.
- Your company has standardized on one of these providers, Oracle Identity Cloud Service, an internal OpenIDConnect or OAuth2 system for authentication.

User credential verification is performed by these systems. Be aware that anyone who is registered at this provider can use your application, unless you use authorization schemes for protection.

Related Topics

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

REST Enabled SQL Support

As opposed to creating Database Links, Oracle Application Express installations that utilize Oracle REST Data Services (ORDS) 17.3 or later, can now execute any SQL through a REST endpoint. Developers 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, and a self-describing JSON response is returned. The JSON object contains result set metadata, the result data, and pagination details.

REST Enabled SQL references can be 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, JET Charts, Trees, and PL/SQL Processes.

Oracle Database Links are defined within each SQL statement 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 as ORDS utilizes a connection pool on the remote database.

Related Topics

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

Improved Create Application Wizard

The completely revamped Create Application Wizard provides the following:

- An all new low-code method of creating applications.
- A new and improved user experience for creating applications.
- Simpler and modernized wizards for creating pages.
- The ability to create more advanced pages such as Dashboards and Master-Detail.
- Support for adding common frameworks or features when creating an application such as access control, activity reporting, or theme selection.
- The ability to customize user interface options such as Theme Style, the application icon, and page icons.

The primary benefits of utilizing this new wizard are the ability to quickly create new Oracle Application Express applications with zero coding, and more advanced pages.

Another key benefit is the ability to refine a previous wizard definition. Developers can go back into the Create Application Wizard, retrieve the definition from a previous wizard (blueprint), update the definitions and regenerate another application.

Related Topics

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

Improved Create Page Wizard

The revamped Create Page Wizard includes the following improvements:

- **New page type: Side by Side Master Detail** - The left panel allows searching on the master record. The right panel displays the master record using a value pair report, and up to four detail reports using classic reports.
- **New page type: Dashboard** - Select from different chart layouts that are based on sample data. The generated charts can easily be updated in Page Designer post-generation.
- The ability to add common frameworks or features to an existing application, such as access control, activity reporting, theme selection, and more (providing the application is utilizing the Universal Theme).
- The ability to add Email Reporting and Job Reporting (providing jobs are defined in the default schema).
- The ability to create an Administration page for features, or select an existing page.

The improvements made to the Create Page Wizard are designed to easily deliver new page types, and incorporate powerful features into existing Oracle Application Express applications.

Related Topics

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

Web Source Modules

Oracle Application Express introduces a new data source type called Web Source Modules, a declarative method to define references to external REST APIs and generic JSON data feeds. 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.

Web Source Modules can also include post-processing SQL which modifies the data before being processed by the Oracle Application Express component. This SQL can be used to apply functions, aggregations, or join with local tables.

Web Source Modules can be used as the basis for all report types, such as interactive reports and classic reports, but not Interactive Grid regions. These modules can also be used with Calendars, JET Charts, Trees, and PL/SQL Processes.

In earlier releases of Oracle Application Express, it was possible to define 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 as they use discovery to understand and define the incoming structure of the web service.

Related Topics

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

New REST Workshop

Before release 18.1, the definitions of RESTful services created within Application Express were stored within the metadata tables of the core Application Express schema. Now Oracle Application Express can utilize the Oracle REST Data Services (ORDS) repository, provided that Application Express is using ORDS 17.4 or above. Oracle recommends the migration of all RESTful services to the Oracle REST Data Services (ORDS) repository.

Existing Application Express-based REST services can readily be migrated to the new ORDS repository. The Application Express-based REST services continue to work, however, you cannot create new or edit existing Application Express based RESTful services.

By utilizing the ORDS repository for REST services, it is far easier to manage RESTful services in a single place using a multitude of tools, including Application Express, SQL Developer, SQL Plus, and SQLcl.

Related Topics

- *Oracle® Application Express SQL Workshop Guide*

General

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

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`, which 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*

Application Continuity for Java: Declarative Request Demarcation

With the introduction of this feature, when Application Continuity for Java is configured in `AUTO` mode (that is service `FAILOVER_TYPE=AUTO`), the JDBC driver injects a `beginRequest` call at runtime, after the creation of a JDBC connection with the `Replay Datasource`.

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*

Oracle Network Log File Segmentation

The maximum size and number of text log files can be configured for Oracle Network components such as Oracle Net Listener, CMAN and GSM.

This feature prevents issues of ever-increasing log file sizes.

Related Topics

- *Oracle® Database Net Services Administrator's Guide*

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

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 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.

The goal of this feature is to automatically probe 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.

Related Topics

- *Oracle® Database Net Services Administrator's Guide*

Java

- [Java Library for Reactive Streams Ingestion](#)

Java Library for Reactive Streams Ingestion

This feature describes the new Java library for high speed ingestion of data streams with non-blocking back pressure. Java applications that use the provided APIs may continuously receive and ingest data from a large group of clients. The ingestion process is non-blocking and extremely fast through the direct path load into the database tables. Through the Universal Connection Pool (UCP), the ingestion process inherits Oracle RAC and Sharded database support, and furnishes high availability and scalability.

This feature enables implementing high-speed ingestion of streaming data with scalability and high availability.

Related Topics

- *Oracle® Database JDBC Developer's Guide*

JSON

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

Materialized View Support for Queries containing JSON_TABLE

Materialized views query rewriting has been enhanced so that queries with `JSON_EXISTS`, `JSON_VALUE` and other functions can utilize a materialized view created over a query that contains a `JSON_TABLE` function.

This feature is particularly useful when the JSON documents in a table 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](#)

SQL/JSON Syntax Simplifications

Syntax simplifications are offered for SQL/JSON path expressions, SQL/JSON generation with the `json_object` function, and field projection with the SQL/JSON nested clause.

These features make the SQL interface for JSON processing easier to use for certain operations. For example, a query such as `SELECT JSON_OBJECT(*) FROM emp` can be used to construct a JSON representation for a row of table `EMP`. The `JSON NESTED` clause can be used to succinctly map JSON values to table rows: `SELECT * FROM customer NESTED jsonCol COLUMNS(first_name, last_name, address, state, zip)`.

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 the new SQL/JSON function `json_serialize` to serialize JSON data to text or to UTF-encoded BLOB data. The `json_dataguide` SQL aggregate function can now detect GeoJSON geographic data in your documents.

The `json_serialize` function is useful for extracting JSON values as text for printing or display. You can now use `json_dataguide` to create a view that projects such data as SQL data type `SDO_GEOMETRY`.

Related Topics

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

JSON Update Operations

You can now update a JSON document declaratively using the new SQL function `json_mergepatch`. You can apply one or more changes to multiple documents by using a single statement.

This feature improves the flexibility of JSON update operations.

Related Topics

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

JSON-Object Mapping

This feature enables the mapping of 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](#)

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. With the new `DISTINCT` keyword, duplicate values can be removed from the specified expression before concatenation into a single string. This removes the need to create complex query processing to find the distinct values before using the aggregate `LISTAGG` function. With the `DISTINCT` option, the processing to remove duplicate values can be done directly within the `LISTAGG` function.

The result is simpler, faster, more efficient SQL.

Related Topics

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

Availability

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

General

- [Simplified Database Parameter Management in a Broker Configuration](#)
- [Dynamically change Fast-Start Failover \(FSFO\) target](#)
- [Observe only mode for Broker's Fast_Start Failover \(FSFO\)](#)
- [Flashback Standby database when Primary database is flashed back](#)
- [Propagate Restore Points from Primary to Standby site](#)
- [Oracle Data Guard Multi-Instance Redo Apply works with the In-Memory Column Store](#)
- [Active Data Guard DML Redirection](#)
- [PDB Recovery catalog](#)
- [Clear Flashback logs periodically for increased FRA size predictability](#)
- [New Parameters for tuning automatic outage resolution with Data Guard](#)
- [Finer granularity Supplemental Logging](#)

Simplified Database Parameter Management in a Broker Configuration

Users can now manage all Data Guard related parameter settings using the SQL*Plus `ALTER SYSTEM` commands or in DGMGRL with the new `EDIT DATABASE ... SET PARAMETER` command. Parameter changes made in the DGMGRL interface are immediately executed on the target database. In addition this new capability allows the user to modify a parameter on all databases in a Data Guard configuration using the `ALL` qualifier, eliminating the requirement to attach to each database and execute an `ALTER SYSTEM` command or set a Broker property for each database with multiple `EDIT PROPERTY` commands. The `SHOW` command has also been updated to show the current setting of a parameter in the target database.

This feature eliminates any inconsistency between a database's Data Guard parameter settings and the Data Guard Broker's Property settings, simplifying the Database Administrator management of the database parameters by allowing all parameter management through the SQL*Plus interface.

Related Topics

- [Oracle® Data Guard Broker](#)

Dynamically change Fast-Start Failover (FSFO) target

Currently the DBA must disable `Fast_Start Failover` in order to change the FSFO Target standby. With the arrival of multiple FSFO targets in Oracle Database 12.2.0.1, this becomes even more important. This new command allows the user to dynamically change the FSFO Target standby to another standby in the target list without requiring that FSFO be disabled.

Forcing the user to disable FSFO to move to a new target standby exposes the configuration to a period where automatic failover cannot be used at all.

Related Topics

- [Oracle® Data Guard Broker](#)

Observe only mode for Broker's Fast_Start Failover (FSFO)

When a database administrator configures Oracle Data Guard Broker's Fast Start Failover (FSFO) capability, they can now configure it to Observe only creating a test mode to see when a failover or other interaction would have occurred during normal production processing. This allows the user to tune the FSFO properties more precisely and to discover what circumstances in their environment would cause an automatic failover to occur. This makes it easier to justify using automatic failovers to reduce the recovery time for failovers.

This configuration allows users to test an automatic failover configuration without actually causing any impact to the production database. This improves on the existing failover validation that is already present in the Broker, and helps customers feel more at ease about FSFO automatic failover routines.

Related Topics

- *Oracle® Data Guard Broker*

Flashback Standby database when Primary database is flashed back

Flashback Database moves the entire database to an older point in time and opens the database with `RESETLOGS`. In a Data Guard setup, if the primary database is flashed back, the standby site is no longer in sync with the primary. In previous releases, getting the secondary to the same point in time as the primary requires a manual procedure to flash back standby databases. A new parameter is introduced which enables the standby database to be flashed back automatically when Flashback Database is performed on the primary database.

By automatically flashing back the standby database when the primary database is flashed back, time, effort, and human errors are reduced resulting in faster synchronization and reduced recovery time objective (RTO).

Related Topics

- *Oracle® Data Guard Concepts and Administration*

Propagate Restore Points from Primary to Standby site

Normal restore points or guaranteed restore points can be defined at the primary site to enable fast point-in-time recovery in the event of any logical corruption issues. However, this restore point is stored in the control file and is not propagated to the standby database. In the event of a failover, the standby becomes the primary and the restore point information is lost. This feature ensures that the restore point is propagated from the primary to standby sites, so that the restore point is available even after a failover event.

The complexity of the restore and recovery process after a failover is simplified because the standby database is updated with the restore points created on the primary database.

Related Topics

- *Oracle® Data Guard Concepts and Administration*

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

The Oracle Database 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 (Multi-instance Redo Apply) and the fastest analytical query technology (In-Memory Column Store) on the same Oracle 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

Active Data Guard DML Redirection allows for incidental DML to be issued on an Active Data Guard standby database. When DML is executed, the update is passed to the Primary database where it is executed and the resulting redo of the transaction will update the standby after which control will be returned to the application. The DML is executed preserving all ACID properties of the transaction.

DML redirection allows more applications to benefit from an Active Data Guard standby database when some writes are required.

Related Topics

- *Oracle® Data Guard Concepts and Administration*

PDB Recovery catalog

Pluggable databases (PDBs) are supported as a target database and a virtual private catalog (VPC) user can be used to more granularly control permissions to perform backup and restore operations at a PDB level. Metadata view is also limited, so a VPC user can view only data for which the user has been granted permission. In previous releases, connections to the recovery catalog when the target database is a PDB was not supported.

Oracle Database 19c provides complete backup and recovery flexibility for container database (CDB) and PDB level backups and restores, including recovery catalog support.

Related Topics

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

Clear Flashback logs periodically for increased FRA size predictability

Customers have many databases that all use the Fast Recovery Area (FRA). They usually subscribe to FRA by using the `recovery_dest_size` initialization parameter. Difficulties arise when flashback logs are not cleared until space pressure requires it. In many cases, the only remedy is to turn off flashback logging and turn it back on. This feature makes flashback space usage become predictable from a storage

management perspective, since flashback uses no more space than is required by retention. This feature also allows users to control cumulative space pressure by adjusting the flashback retention.

The FRA is critical for databases because it stores backups, online redo logs, archived redo logs, and flashback logs. When the FRA becomes full, it affects all the databases. By automatically ensuring that flashback logs do not overutilize the space needed for retention, storage management improves and the health of the database is increased.

Related Topics

- [Oracle® Database Backup and Recovery User's Guide](#)

New Parameters for tuning automatic outage resolution with Data Guard

Oracle Data Guard has several processes on the Primary and Standby databases that handle redo transport and archiving which communicate with each other over the network. 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. Data Guard has an internal mechanism to detect these hung processes and terminate them allowing the normal outage resolution to occur. In Oracle Database 19c, the DBA can tune the amount of wait time for this detection period by using two new parameters, `DATA_GUARD_MAX_IO_TIME` and `DATA_GUARD_MAX_LONGIO_TIME`. These parameters allow the wait times to be tuned for a specific Data Guard configuration based on the user network and Disk I/O behavior.

Users can now tune Oracle Data Guard automatic outage resolution to fit their specific needs.

Related Topics

- [Oracle® Database Reference](#)

Finer granularity Supplemental Logging

Supplemental logging was designed and implemented for Logical Standby or full database replication requirements. This adds unnecessary overhead in environments where only a subset of tables is being replicated. 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 database or schema level, there is no supplemental logging overhead for uninteresting tables.

Use of this feature can significantly reduce the overhead in terms of resource usage and redo generation in case when only some of the tables in the database require supplemental logging, such as in a GoldenGate partial replication configuration.

Sharding

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

Propagation of Parameter Settings Across Shards

Before Oracle Database 19c, database administrators had to configure `ALTER SYSTEM` parameter settings on each shard in a sharded database.

This feature allows administrators to centrally manage and propagate parameter settings from the shard catalog to all of the database shards, improving ease of manageability.

Related Topics

- *Oracle® Database Using Oracle Sharding*

Support for Multiple PDB Shards in the Same CDB

Oracle Sharding with Oracle Database 18c supported one pluggable database (PDB) shard in a CDB. In Oracle Database 19c, Oracle Sharding enables you to 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 and ease of management.

Related Topics

- *Oracle® Database Using Oracle Sharding*

Multiple Table Family Support for System-Managed Sharding

The Oracle Sharding feature for Oracle Database 18c supported only one table family (a set of related tables sharing the same sharding key) for each sharded database. In Oracle Database 19c, Oracle Sharding allows a sharded database to support multiple table families, each of which can be sharded with a different sharding key. Data from different table families reside in the same chunks. This feature applies to system-managed sharded databases only.

Different applications that access different table families can now be hosted on one sharded database.

Related Topics

- *Oracle® Database Using Oracle Sharding*

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

Before Oracle Database 19c, only the primary shard catalog database could be used as the multi-shard query coordinator. In Oracle Database 19c, you can also enable the multi-shard query coordinator on the shard catalog's Oracle Active Data Guard standby databases.

Oracle Database 19c improves the scalability and availability of a multi-shard query workload.

Related Topics

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

Generation of Unique Sequence Numbers Across Shards

Before Oracle Database 19c, if you needed a unique number across shards you had to manage it yourself. In Oracle Database 19c, Oracle Sharding allows you to generate globally unique sequence numbers across shards. You can use this functionality to generate globally unique sequence numbers for non-primary key columns with unique constraints, or any other case in which a Sequence object must be a single logical object across all shards of a sharded database.

Customers often need to generate unique IDs for non-primary key columns, such as `order_id`, when the `customer_id` is the sharding key. In this case and others, this feature lets you generate unique sequence numbers across shards, while not requiring you to manage the global uniqueness of a given non-primary key column in your application.

Related Topics

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

Big Data and Data Warehousing

- [General](#)

General

- [SQL Diagnostics and Repair Enhancements](#)
- [Automatic Indexing](#)
- [Bitmap based count distinct SQL Function](#)
- [Big Data and Performance Enhancements for In-Memory External Tables](#)
- [Automatic Resolution of SQL Plan Regressions](#)
- [Real-Time Statistics](#)
- [High-Frequency Automatic Optimizer Statistics Collection](#)
- [Hybrid Partitioned Tables](#)

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](#)

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*

Bitmap based count distinct SQL Function

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

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

In most scenarios, bitvector 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 bitvectors at lower-level aggregation levels, the same materialized view can be reused 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 Resolution of SQL Plan Regressions

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.

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.

Related Topics

- *Oracle® Database SQL Tuning Guide*

Real-Time Statistics

Oracle Database automatically gathers online statistics during conventional 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 Partition 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 Partition Tables enable you to easily integrate internal partitions and external partitions into a single partition table. With this feature, you can also easily 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 install, config, and patch](#)
- [Automated upgrade, migration and utilities](#)
- [General](#)

Automated install, config, and patch

Focus area for database and GI installation, configuration and patching

- [Ability to Create a Duplicate of an Oracle Database Using DBCA in Silent Mode](#)
- [Ability to Create a PDB by Cloning a Remote PDB Using DBCA in Silent Mode](#)
- [Ability to Relocate a PDB to Another CDB 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 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 Create a PDB by Cloning a Remote PDB Using DBCA in Silent Mode

You can now create a PDB by cloning a remote PDB using the `createFromRemotePDB` parameter of the `createPluggableDatabase` command of 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](#)

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

You can now relocate a PDB to another CDB by using the `relocatePDB` command of 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](#)

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](#)

Automated upgrade, migration and utilities

Focus area to include database upgrade, all types of database migration, and utilities such as SQL*Loader and external tables

- [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 Test Mode for Transportable Tablespaces](#)

- [Oracle Data Pump Support for Resource Usage Limitations](#)

Oracle Data Pump Ability to Exclude ENCRYPTION Clause on Import

A new transform parameter, `OMIT_ENCRYPTION_CLAUSE` is introduced 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 PL/SQL Packages and Types Reference](#)

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

A new option allows the user to restore pre-12.2 default behavior, such that tablespace data files can be read-only during the Transportable Tablespace (TTS) 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 will not be adjusted upon import. Also, if this parameter is specified 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 datafiles.

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

Related Topics

- [Oracle® Database Utilities](#)

Oracle Data Pump Test Mode for Transportable Tablespaces

Test Mode for Transportable Tablespaces performs a metadata-only export test using Transportable Tablespaces or Full Transportable Export/Import. It also removes the requirement for the source database tablespaces to be in read-only mode.

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

Related Topics

- [Oracle® Database Utilities](#)

Oracle Data Pump Support for Resource Usage Limitations

Two new parameters, `MAX_DATAPUMP_JOBS_PER_PDB` and `MAX_DATAPUMP_PARALLEL_PER_JOB` are introduced to give DBAs more control over the number of jobs that can be started in a multitenant container database environment, and over the number of parallel workers that can be used for an individual Data Pump job, respectively.

These parameters give the DBA more control over resource utilization when there are multiple users performing data pump jobs in a database environment.

Related Topics

- [Oracle® Database Utilities](#)

General

- [Oracle Data Pump Prevents Inadvertent Use of Protected Roles](#)
- [Oracle Data Pump Allows Wildcards for Dumpfile in Object Store](#)
- [Oracle Data Pump import supports more object store credentials](#)

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`.

Oracle Data Pump Allows Wildcards for Dumpfile in Object Store

Oracle Data Pump simplifies importing multiple dumpfiles into Oracle Autonomous Database by allowing wildcards for URL-based dumpfile names.

When you need to import multiple dumpfiles 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. A wildcard character can not be used in the bucket-name component.

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 release 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.

Diagnosability

- [General](#)

General

- [Oracle Trace File Analyzer REST API Support](#)
- [Oracle Trace File Analyzer Search Extended to Support Metadata Searches](#)

- [Oracle Trace File Analyzer Supports New Service Request Data Collections](#)
- [Oracle ORAck and Oracle EXAck Support for Encrypting Collection Files](#)
- [Oracle ORAck and Oracle EXAck REST Support](#)
- [Oracle Cluster Health Advisor Integration into Oracle Trace File Analyzer](#)
- [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 Support for Using an External SMTP Server for Notifications](#)

Oracle Trace File Analyzer REST API Support

Oracle Trace File Analyzer now includes REST support, allowing invocation and query over HTTPS. To help develop and test this REST support, Oracle REST Data Services (ORDS) are included within the installation. REST supports printing details, starting a diagnostic collection, and downloading collections.

Business that automate their data center operations require remote management and automation. The REST interface is an industry standard that supports remote management. Oracle Trace File Analyzer operation through REST APIs supports easy integration into customer operations frameworks, improving diagnostic efficiency and reducing recovery time.

Related Topics

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

Oracle Trace File Analyzer Search Extended to Support Metadata Searches

In earlier releases, Oracle Trace File Analyzer's search function was limited to log and trace file strings. Starting in this release, all metadata stored in the Oracle Trace File Analyzer index is searchable using `tfactl`. Oracle Trace File Analyzer searches log and trace file metadata using JSON-formatted name-value pairs representing datatypes 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 recurring issues across instances and nodes.

Related Topics

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

Oracle Trace File Analyzer Supports New Service Request Data Collections

Service Request Data Collections (SRDCs) simplify the collection of required logs and data for specific issues. This release adds additional database SRDCs that cover more ORA errors and problems in infrastructure such as Oracle Automatic Storage Management (ASM), Oracle Automatic Storage Management Cluster File System (ACFS), Listeners, auditing, and Recovery Manager (RMAN).

When operations or Oracle Database issues occur that require Oracle Support Services, it is essential to send all the data and logs necessary to diagnose and

resolve the issue in one compact complete archive. Oracle Trace File Analyzer's SRDCs provide this functionality while minimizing the steps required to compile and send the archive efficiently. This improves recovery time while improving the administrator's efficiency.

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. This feature is available only on Linux and Solaris platforms.

Companies are increasingly concerned about the leakage of sensitive data. Oracle ORAchk and Oracle EXAchk collections and their reports can include such data. When these reports are transferred to repositories or emailed, it's critical that such data is viewed only by the intended recipients. To prevent leaks, you can restrict access to sensitive data by encrypting the diagnostic collections and protecting them with a password.

Related Topics

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

Oracle ORAchk and Oracle EXAchk REST Support

Oracle ORAchk and Oracle EXAchk now include REST support, allowing invocation and query over HTTPS. To facilitate this REST support, Oracle REST Data Services (ORDS) is included within the installation.

Business that automate their data center operations require remote management and automation. The REST interface is an industry standard that supports remote management. Oracle ORAchk and Oracle EXAchk operation through REST APIs supports easy integration into customer operations frameworks, improving diagnostic efficiency and reducing recovery time.

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 performance issues. By delivering email notifications with root cause analysis and corrective recommendations through Oracle Trace File Analyzer's daemon, operations and database administrators can proactively prevent application performance and availability issues.

Related Topics

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

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

In earlier releases, remotely running Oracle ORAchk or Oracle EXAchk required configuration of passwordless SSH between the remote nodes. 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.

To centrally manage many database servers or clusters, it is more efficient to perform operations remotely. In many cases, corporate policies prevent passwordless SSH configuration. Using private key authentication, you can run Oracle ORAchk and Oracle EXAchk remotely in these deployments and improve operational efficiency.

Related Topics

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

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

By default, only the most critical checks are shown in the report output. The critical checks are those that Oracle judges to have the most severe potential effect. All other checks are still run and are available in the report. You can view them by selecting the appropriate option under the "Show checks with the following status" control.

In earlier versions, Oracle EXAchk and Oracle ORAchk reports included over a hundred checks and made analysis more time-consuming. When only the most critical checks are included, analyzing the reports becomes more efficient, and you can more quickly resolve critical problems to prevent downtime or performance issues.

Related Topics

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

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

In earlier releases, Oracle Trace File Analyzer required monitored hosts to have local sendmail or SMTP support in order to deliver email notifications of alerts. Starting in this release, you can configure an external SMTP server to receive these notifications from Oracle Trace File Analyzer and alert the administrators.

Oracle Trace File Analyzer's ability to alert administrators immediately when it detects an anomaly or issue is essential to maintain availability and rapidly recover from problems. By extending email notification support to Oracle Trace File Analyzer deployments that cannot send email locally, these deployments can minimize downtime and maximize availability.

Related Topics

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

Performance

- [General](#)

General

- [SQL Quarantine](#)
- [Resource Manager Automatically Enabled for Database In-Memory](#)
- [Database In-Memory Wait on Populate](#)
- [Memoptimized Rowstore - Fast Ingest](#)
- [Automatic Database Diagnostic Monitor \(ADDM\) Support for Pluggable Databases \(PDBs\)](#)
- [Real-Time SQL Monitoring for Developers](#)
- [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 can be 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](#)

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, customers receive the benefits of enhanced performance and automatic management for CPU resource allocation.

Related Topics

- [Oracle® Database In-Memory 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*

Memoptimized Rowstore - Fast Ingest

The Memoptimized Rowstore's fast ingest functionality 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 Oracle Database's rich analytical features, 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 ADDM analysis for 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*

Real-Time SQL Monitoring for Developers

Database users without `SELECT_CATALOG_ROLE` can generate and view SQL Monitor reports for their own SQL statements, including the execution plans and performance metrics.

A primary responsibility of database developers is to write and tune SQL statements. Access to SQL Monitor reports enables developers to perform these duties without requiring database administrator privileges.

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 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](#)
- [Automated PDB Relocation](#)
- [Automated Transaction Draining for Oracle Grid Infrastructure Upgrades](#)
- [Oracle Restart Patching and Upgrading](#)
- [Zero-Downtime Oracle Grid Infrastructure Patching](#)

Parity Protected Files

The `REDUNDANCY` file type property specifies the redundancy for a file group. The `PARITY` value specifies single parity for redundancy. The parity setting is intended for write-once files, such as archive logs and backup sets.

A great deal of space is consumed when two or three way 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](#)

Automated PDB Relocation

In Oracle Grid Infrastructure, you can use Fleet Patching and Provisioning to automate relocation of a PDB from one CDB to another.

This allows individual PDBs to be patched more quickly and without exposing other PDBs to the changes that a patch would bring.

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 can be 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 the user 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 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](#)

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 customers to perform a rolling patch of Oracle Grid Infrastructure without interrupting database operations on the node being patched and without impacting capacity or performance on those database instances.

Related Topics

- [Oracle® Clusterware 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](#)
- [Key Management of Encrypted Oracle-Managed Tablespaces in Transparent Data Encryption](#)
- [Support for Additional Algorithms for Offline Tablespace Encryption](#)

- Support for Host Name-Based Partial DN Matching for Host Certificates
- 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
- Unified Auditing Top-Level Statements
- Passwords Removed from Oracle Database Accounts
- Signature-Based Security for LOB Locators
- New EVENT_TIMESTAMP_UTC Column in the UNIFIED_AUDIT_TRAIL View
- New PDB_GUID Audit Record Field for SYSLOG and the Windows Event Viewer
- 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. The latest details of the database password file can be retrieved by querying the V\$PASSWORDFILE_INFO view.

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*

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

In this release, a closed 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 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*

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

This new support for partial DN matching 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*

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 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

Administrative privileges such as `SYSOPER` and `SYSBACKUP` can now be granted 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, both Simple Authentication and Security Layer (SASL) and Transport Layer Security (TLS) binds are supported 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*

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 can be 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*

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 instances of LOB data types (CLOB and BLOB) are used in distributed environments.

LOB signature keys can be in both multitenant 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 TDE keystore.

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*

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 associated with a unified audit trail record.

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

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 DBAs, 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*

2

Oracle Database 19c Release 19.3 New Features

This chapter contains descriptions of all of the features that are new to Oracle Database 19c Release 19.3.

- [AutoUpgrade and Database Utilities](#)
- [RAC and Grid](#)

AutoUpgrade and Database Utilities

- [Database Utilities](#)
- [Upgrade](#)

Database Utilities

Focus area to include database upgrade, all types of database migration, and utilities such as SQL*Loader and external tables

- [Oracle Data Pump loads partitioned table data in one operation](#)

Oracle Data Pump loads partitioned table data in one operation

Oracle Data Pump can optionally import the table data in all partitions of a table as one operation instead of a separate operation 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, 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.

Upgrade

- [AutoUpgrade for Oracle Database](#)

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 preupgrade 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 a single DBA 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 DBAs. 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](#)

RAC and Grid

- [General](#)

General

- [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](#)

Colocation Tag for Client Routing

The `COLOCATION_TAG` parameter is an alphanumeric string that you can use with the `CONNECT_DATA` parameter of the 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 Clusters still require 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 can help 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 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 Oracle Cluster Registry (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 (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 users can now specify that this service should fall back to a "preferred" instance when it becomes available and once it was failed over to an available instance with the last preferred instance going down.

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*