Oracle® Cloud

Exadata Database Service on Dedicated Infrastructure

F92140-20
May 2024

Oracle Cloud Exadata Database Service on Dedicated Infrastructure,

F92140-20

Primary Authors: James Spiller, Neil Hebert, Nirmal Kumar

Contributors: Bryce Cracco, Sanjay Narvekar, Pravin Jha , Aneesh Khandelwal, Alan Williams, Bethany Lapaglia, Chandrashekhar Garud, Dileep Thiagarajan, Guruprasad Hegde, Jai Krishnani, Jeffery Wright, Vineet Kakani

Contents

3 Preparing for Exadata Cloud Infrastructure

4 Getting Started with Exadata Cloud Infrastructure Deployment

5 How-to Guides

6 Reference Guides for Exadata Cloud Infrastructure

1

Oracle Exadata Database Service on Dedicated Infrastructure Overview

This topic is an overview of the Exadata Cloud Infrastructure formerly Exadata Cloud Service.

- Oracle Exadata Database Service on Dedicated Infrastructure Description
  Learn how you can leverage the combined capabilities of Oracle Exadata and Oracle Cloud Infrastructure with Oracle Exadata Database Service on Dedicated Infrastructure
- Exadata Cloud Management Interfaces
  Exadata Cloud Infrastructure provides a variety of management interfaces to fit your use case and automation needs.

Oracle Exadata Database Service on Dedicated Infrastructure Description

Learn how you can leverage the combined capabilities of Oracle Exadata and Oracle Cloud Infrastructure with Oracle Exadata Database Service on Dedicated Infrastructure

- About Exadata Cloud Infrastructure
  Exadata Cloud Infrastructure allows you to leverage the power of Exadata in the cloud.
- Licensing Considerations for Oracle Exadata Database Service on Dedicated Infrastructure
  Subscription to Exadata Cloud Infrastructure can include all of the required Oracle Database software licenses, or you can choose to bring Oracle Database software licenses that you already own to Oracle Exadata Database Service on Dedicated Infrastructure.
- Supported Database Edition and Versions for Exadata Cloud Infrastructure
  Exadata Cloud Infrastructure databases require Enterprise Edition - Extreme Performance subscriptions or you can bring your own Oracle Enterprise Edition software licenses.
- Subscription Types
  Available subscription types.
- Metering Frequency and Per-Second Billing
  Per-Second billing, minimums and limitations on billing.
- Technical Architecture of Exadata Cloud Infrastructure Systems
  Exadata Cloud Infrastructure systems integrate Oracle's Exadata Database Machine hardware with the networking resources needed to securely connect to your organization’s on-premise network and to other services in the Oracle cloud.
- Scaling Options
  Introduction to the Scaling option on Exadata Cloud Infrastructure.
- System and Shape Configuration Options
  Review the list of Exadata System Shapes

About Exadata Cloud Infrastructure

Exadata Cloud Infrastructure allows you to leverage the power of Exadata in the cloud.

You can provision flexible X8M and X9M systems that allow you to add database compute servers and storage servers to your system as your needs grow. X8M and X9M systems offer RDMA over Converged Ethernet (RoCE) networking for high bandwidth and low latency, persistent memory (PMEM) modules, and intelligent Exadata software. X8M and X9M systems can be provisioned using an shape equivalent to a quarter rack X8 or X9M system, and then database and storage servers can be added at any time after provisioning. For more information on X8M and X9M systems, see Overview of X8M and X9M Scalable Exadata Infrastructure.

> **Note:**
>
> The RDMA software allows computers in a network to exchange data in main memory without involving the processor, cache, or OS of either computer. RDMA can improve throughput and performance because it frees up resources, and it can also facilitate a faster data transfer rate. RDMA over Converged Ethernet (RoCE) is the network protocol that allows RDMA over an Ethernet network.

X8 and X7 systems are also available in fixed-shapes (quarter, half, and full rack systems). These systems use InfiniBand networking, and do not have the ability to scale database and storage servers. You can also provision an Exadata base system, which has a smaller capacity than a quarter rack system.

For all Exadata Cloud Infrastructure instances, you can configure automatic backups, optimize for different workloads, and scale the OCPU and storage allocations as needed.

Note

Exadata Cloud Infrastructure instances launched on or after March 14, 2019 run Oracle Linux 7. Previously launched systems are running Oracle Linux 6. See To update the OS on all compute nodes of an Exadata Cloud Infrastructure instance for important information about updating existing Exadata DB system operating systems.

- Roles and Responsibilities for Exadata Cloud Infrastructure Software Maintenance
  Oracle is responsible for the base OS and hardware. The customer is responsible for the Guest VM OS, Grid Infrastructure, and the database software maintenance.

Roles and Responsibilities for Exadata Cloud Infrastructure Software Maintenance

Oracle is responsible for the base OS and hardware. The customer is responsible for the Guest VM OS, Grid Infrastructure, and the database software maintenance.

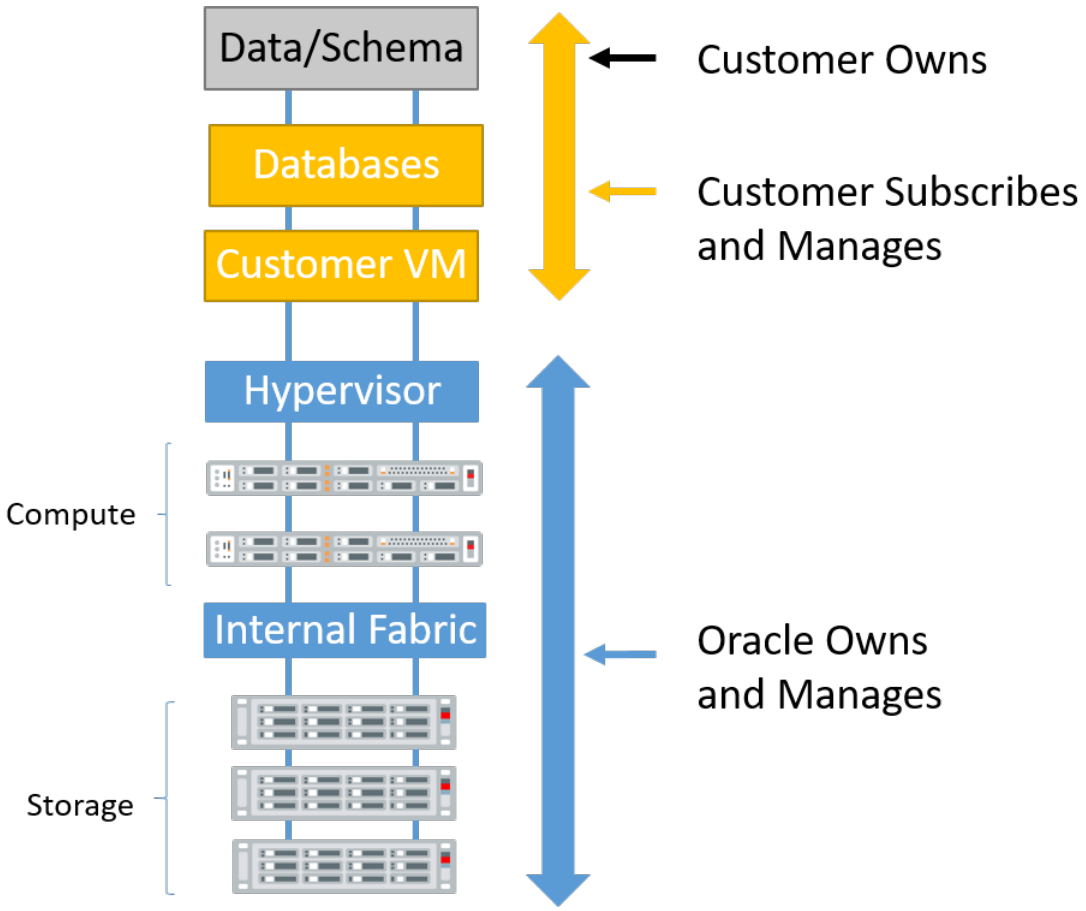

Customer Responsibilities

Customer is responsible for the Guest VM OS, Grid Infrastructure, and the database software maintenance.

The customer owns everything inside the database: data schema, and encryption keys.

Customer subscribes to database services

- Customer manages VMs and Databases using Cloud Automation (UI / APIs)
- Automation to create, delete, patch, backup, scale up/down, etc.
- Runs all supported Oracle Database versions from 11.2.0.4 to 19c
- Customer controls access to customer VM
- Customer can install and manage additional software in customer VM
- Oracle staff are not authorized to access customer VM

Oracle Responsibilities

Oracle owns and manages infrastructure. Oracle is responsible for the base OS and hardware.

- Hypervisor, physical database and storage servers, storage network

- Patching, security scans, security updates
- Monitoring and maintenance
- Customer is not authorized to access Oracle infrastructure.

Licensing Considerations for Oracle Exadata Database Service on Dedicated Infrastructure

Subscription to Exadata Cloud Infrastructure can include all of the required Oracle Database software licenses, or you can choose to bring Oracle Database software licenses that you already own to Oracle Exadata Database Service on Dedicated Infrastructure.

If you choose to include Oracle Database software licenses in your Oracle Exadata Database Service on Dedicated Infrastructure subscription, then the included licenses contain all of the features of Oracle Database Enterprise Edition, plus all of the database enterprise management packs, and all of the Enterprise Edition options, such as Oracle Database In-Memory and Oracle Real Application Clusters (Oracle RAC). Exadata Cloud Infrastructure also comes with cloud-specific software tools that assist with administration tasks, such as backup, recovery, and patching.

Supported Database Edition and Versions for Exadata Cloud Infrastructure

Exadata Cloud Infrastructure databases require Enterprise Edition - Extreme Performance subscriptions or you can bring your own Oracle Enterprise Edition software licenses.

The Enterprise Edition - Extreme Performance provides all the features of Oracle Database Enterprise Edition, plus all the database enterprise management packs and all the Enterprise Edition options, such as Oracle Database In-Memory and Oracle Real Application Clusters (Oracle RAC).

Exadata Cloud Infrastructure supports the following database versions:

- Oracle Database 23ai
- Oracle Database 19c (19.0)
- Oracle Database 12c Release 2 (12.2) (**Upgrade Support Required**)
- Oracle Database 12c Release 1 (12.1) (**Upgrade Support Required**)
- Oracle Database 11g Release 2 (11.2) (**Upgrade Support Required**)

> **Note:**
>
> - If you plan to run Oracle Database 19c on a cloud VM cluster or in the Exadata Cloud Infrastructure, you must specify version 19c when you create the resource. Earlier database versions are supported on a 19c cloud VM cluster and can be created at anytime. Cloud VM clusters created with earlier Oracle Database versions will not automatically support Oracle Database 19c.
> - For information on upgrading an existing 18c or earlier database to Oracle Database 19c, see Upgrading Exadata Databases.

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

Related Topics

- Release Schedule of Current Database Releases (Doc ID 742060.1)

Subscription Types

Available subscription types.

The available purchase models are:

- **Pay As You Go**
  Pay As You Go (PAYG) pricing lets customers quickly provision services with no commitment, and they're only charged for what they use. There's no upfront commitment and no minimum service period. Any cloud infrastructure (IaaS) and platform (PaaS) services consumed are metered and billed based on that consumption. If, during the services period of your order, Oracle makes new IaaS and PaaS services available within your cloud services account, Oracle will notify you of any fees that would apply to their activation and use. For more details, see our complete price list.
- **Annual Universal Credits**
  Oracle Annual Universal Credits enables customers to have the flexibility to use any Oracle Cloud Infrastructure and platform services at any time, in any region, to deliver faster time to market. Customers can commit to an amount of Oracle Annual Universal Credits that can be applied towards the future usage of eligible Oracle IaaS and PaaS cloud services. This payment option offers a significant savings across cloud services, combining cost reduction and a predictable monthly spend with a ramp up period as you onboard your workloads.

See the Universal Credit Pricing FAQ for more information.

Metering Frequency and Per-Second Billing

Per-Second billing, minimums and limitations on billing.

For each Exadata Cloud Infrastructure instance you provision, you are billed for the infrastructure for a minimum of 48 hours, and then by the second after that. Each OCPU you add to the system is billed by the second, with a minimum usage period of 1 minute. If you terminate the cloud VM cluster and do not terminate the cloud Exadata infrastructure resource, billing will continue for the infrastructure resource.

Technical Architecture of Exadata Cloud Infrastructure Systems

Exadata Cloud Infrastructure systems integrate Oracle's Exadata Database Machine hardware with the networking resources needed to securely connect to your organization's on-premise network and to other services in the Oracle cloud.

For a complete architectural overview of the components that make up an Exadata Cloud Infrastructure system, see Oracle Exadata Cloud Service (ExaCS) Technical Architecture. This interactive reference guides you through the key hardware and networking resources in the system, and provides technical specifications for the database (compute) and storage servers to help you plan for your deployment.

Scaling Options

Introduction to the Scaling option on Exadata Cloud Infrastructure.

Two kinds of scaling operations are supported for an Exadata Cloud Infrastructure:

- For X8M and X9M systems, the flexible shape allows you to add additional database and storage servers to the cloud Exadata infrastructure resource as you need them. See *Overview of X8M Scalable Exadata Infrastructure*.
- For X6, X7 and X8 Exadata DB systems, you can scale by moving the system to a different shape configuration, for example, from a quarter rack to a half rack.

For more information on each type of scaling, see *Scaling an Exadata Cloud Infrastructure Instance*.

- Scaling CPU cores within an Exadata Cloud Infrastructure instance
  If an Exadata Cloud Infrastructure instance requires more compute node processing power, you can scale up the number of enabled CPU cores symmetrically across all the nodes in the system as follows:
- Scaling X6, X7 and X8 Exadata Cloud Infrastructure Instances Configurations
  Scaling an Exadata X6, X7, or X8 Exadata Cloud Infrastructure instance by moving to a shape with more capacity enables you meet the needs of your growing workload.

Related Topics

- Overview of X8M and X9M Scalable Exadata Infrastructure
  Oracle Cloud Infrastructure scalable X8M and X9M Exadata cloud infrastructure model allows you to add additional database and storage servers after provisioning and create a system that matches your capacity needs.
- Scaling Resources within an Exadata Infrastructure Instance
  If an Exadata Cloud Infrastructure instance requires more resources, you can scale up the number of DB servers, or storage servers.

Scaling CPU cores within an Exadata Cloud Infrastructure instance

If an Exadata Cloud Infrastructure instance requires more compute node processing power, you can scale up the number of enabled CPU cores symmetrically across all the nodes in the system as follows:

The options for each of the shapes are:

NOT_SUPPORTED

You can scale the infrastructure by adding DB servers or Storage Servers, up to the infrastructure limits. For more information on adding compute and storage resources to an X8M or X9M system emabled for MVM, see *Scaling Exadata X8M or X9M Compute and Storage* .

NOT_SUPPORTED

You can scale CPU cores in multiples of the number of database servers currently provisioned for the cloud VM cluster. For example, if you have 6 database servers provisioned, you can add CPU cores in multiples of 6. At the time of provisioning, X8M or X9M systems have as few as 2 database servers or up to to 32 database servers. For more information on adding compute and storage resources to an X8M or X9M system, see *Scaling Exadata X8M or X9M Compute and Storage* .

NOT_SUPPORTED

All systems that are not X8M or X9M are fixed shape systems. For a base system or an X7 or X8 quarter rack, you can scale in multiples of 2 across the 2 database compute nodes. For an X7 or X8 half rack, you can scale in multiples of 4 across the 4 database compute nodes. For an X7 or X8 full rack, you can scale in multiples of 8 across the 8 database compute nodes.

For a non-metered service instances, you can temporarily modify the compute node processing power (bursting) or add compute node processing power on a more permanent basis. For a metered service instance, you can simply modify the number of enabled CPU cores.

You can provision an Exadata Cloud Infrastructure instance with zero CPU cores, or scale the service instance down to zero cores after you provision it. With zero cores, you are billed only for the infrastructure until you scale up the system. For detailed information about pricing, see Exadata Cloud Service Pricing.

> **Note:**
>
> OCPU scaling activities are done online with no downtime.

For information on CPU cores per configuration, see *Exadata Shape Configurations*. To learn how to scale a system, see *To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster or DB system* .

Related Topics

- Scaling Exadata X8M and X9M Compute and Storage
  The flexible X8M and X9M system model is designed to be easily scaled in place, with no need to migrate the database using a backup or Data Guard
- Exadata Shape Configuration
  This topic describes the available Exadata Cloud Infrastructure instance shapes in Oracle Cloud Infrastructure.
- To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster or DB system

Scaling X6, X7 and X8 Exadata Cloud Infrastructure Instances Configurations

Scaling an Exadata X6, X7, or X8 Exadata Cloud Infrastructure instance by moving to a shape with more capacity enables you meet the needs of your growing workload.

This is useful when a database deployment requires:

- Processing power that is beyond the capacity of the current system configuration.
- Storage capacity that is beyond the capacity of the current system configuration.
- A performance boost that can be delivered by increasing the number of available compute nodes.
- A performance boost that can be delivered by increasing the number of available Exadata Storage Servers.

You can move you workloads to a larger fixed shape (X7 and X8 hardware shapes), or move to the flexible X8M shape that allows for easy expansion of compute and storage resources as your workloads grow.

To assist with moving your database deployments between Exadata Cloud Infrastructure instances, you can restore a backup to a different service instance that has more capacity, or create a Data Guard association for your database in a service instance with more capacity, and then perform a switchover so that your new standby database assumes the primary role. To start the process, contact Oracle and request a service limit increase so that you can provision the larger service instance needed by your database.

System and Shape Configuration Options

Review the list of Exadata System Shapes

- Exadata Shape Configuration
  This topic describes the available Exadata Cloud Infrastructure instance shapes in Oracle Cloud Infrastructure.

Exadata Shape Configuration

This topic describes the available Exadata Cloud Infrastructure instance shapes in Oracle Cloud Infrastructure.

Exadata X9M and X8M shapes start with 2 database and 3 storage servers. Compute and/or storage servers can be added independently to these shapes up to a total of 32 DB servers and 64 storage servers. The initial minimum configuration for Exadata X6, X7, and X8 also starts with 2 database and 3 storage servers similar to the quarter rack shape. They are also offered in half and full rack shapes.

See the following sections for shape specifications:

- Exadata X9M
  The values in the table that follows represent the specifications for an X9M cloud instance with 2 database and 3 storage servers that has not been expanded.
- Exadata X8M
  The values in the table that follows represent the specifications for an X8M cloud instance with 2 database and 3 storage servers that has not been expanded.
- Exadata X8 Shapes
  The values in the table that follows represent the specifications for an X8 cloud instance with fixed Quarter, Half, and Full Rack shapes.
- Exadata X7 Shapes
  The values in the table that follows represent the specifications for an X7 cloud instance with fixed Quarter, Half, and Full Rack shapes.
- Exadata X6 Shapes
  The values in the table that follows represent the specifications for an X6 cloud instance with fixed Quarter, Half, and Full Rack shapes.
- Exadata Base System
  An Exadata Base System is a fixed shape similar in size to a Quarter Rack with some differences in capacity.

Exadata X9M

The values in the table that follows represent the specifications for an X9M cloud instance with 2 database and 3 storage servers that has not been expanded.

Independently add compute and/or storage servers up to a total of 32 DB servers and 64 storage servers.

- A single DB server contains 126 usable cores and 1390 GB memory.
- A single storage server contains 63.6 TB of usable disk storage capacity.

Property	Minimum Configuration
Number of DB servers per System	2
Number of Storage Servers per System	3
Minimum (Default) Number of Enabled CPU Cores	0
Total Usable Cores in DB Servers per System	252
Total Memory Available for VMs (GB)	2780
Max Usable Local Storage Per DB Server (GB)	2243
Max Usable File System Size Per VM (GB)	900
VM Image size minimum and default (GB)	244
Max Number of VM Clusters per System	8
Max Number of VMs per DB server	8
Total Flash Capacity (TB)	76.8
Total Usable Disk Storage Capacity (TB)	190

VM Image size minimum and default includes 60 GB for `/u02`.

A maximum of 8 VM Clusters can be created on a single system that contains greater than 2 DB servers. For more information, see Estimating How Much Local Storage You Can Provision to Your VMs and Scaling Local Storage.

Exadata X8M

The values in the table that follows represent the specifications for an X8M cloud instance with 2 database and 3 storage servers that has not been expanded.

Independently add compute and/or storage servers up to a total of 32 DB servers and 64 storage servers.

- A single DB server contains 50 usable cores and 1390 GB memory.
- A single storage server contains 49.9 TB of usable disk storage capacity.

Property	Minimum Configuration
Number of DB servers per System	2
Number of Storage Servers per System	3
Minimum (Default) Number of Enabled CPU Cores	0
Total Usable Cores in DB Servers per System	100
Total Memory Available for VMs (GB)	2780
Max Usable Local Storage Per DB Server (GB)	2243
Max Usable File System Size Per VM (GB)	900
VM Image size minimum and default (GB)	244
Max Number of VM Clusters per System	8
Max Number of VMs per DB server	8
Total Flash Capacity (TB)	76.8
Total Usable Disk Storage Capacity (TB)	149

VM Image size minimum and default includes 60 GB for `/u02`.

A maximum of 8 VM Clusters can be created on a single system that contains greater than 2 DB servers. For more information, see Estimating How Much Local Storage You Can Provision to Your VMs and Scaling Local Storage.

Exadata X8 Shapes

The values in the table that follows represent the specifications for an X8 cloud instance with fixed Quarter, Half, and Full Rack shapes.

Property	Quarter Rack	Half Rack	Full Rack
Shape Name	Exadata.Quarter3.100	Exadata.Half3.200	Exadata.Full3.400
Number of DB servers per System	2	4	8
Number of Storage Servers per System	3	6	12
Minimum Number (Default) of Enabled CPU Cores	0	0	0
Total Usable Cores in DB Servers per System	100	200	400
Total Memory Available (GB)	1440	2880	5760
Max Usable Local Storage (GB)	700	700	700
Total Flash Capacity (TB)	76.8	179.2	358.4
Total Usable Disk Storage Capacity (TB)	149	299	598

Exadata X7 Shapes

The values in the table that follows represent the specifications for an X7 cloud instance with fixed Quarter, Half, and Full Rack shapes.

Property	Quarter Rack	Half Rack	Full Rack
Shape Name	Exadata.Quarter2.92	Exadata.Half2.184	Exadata.Full2.368
Number of DB servers per System	2	4	8
Number of Storage Servers per System	3	6	12
Minimum Number (Default) of Enabled CPU Cores	0	0	0
Total Usable Cores in DB Servers per System	92	184	368
Total Memory Available (GB)	1440	2880	5760
Max Usable Local Storage (GB)	1000	1000	1000
Total Flash Capacity (TB)	76.8	153.6	307.2
Total Usable Disk Storage Capacity (TB)	106	212	424

Exadata X6 Shapes

The values in the table that follows represent the specifications for an X6 cloud instance with fixed Quarter, Half, and Full Rack shapes.

Property	Quarter Rack	Half Rack	Full Rack
Shape Name	Exadata.Quarter1.84	Exadata.Half1.168	Exadata.Full1.336
Number of DB servers per System	2	4	8
Number of Storage Servers per System	3	6	12
Minimum Number (Default) of Enabled CPU Cores	22	44	88
Total Usable Cores in DB Servers per System	84	168	336
Total Memory Available (GB)	1440	2880	5760
Max Usable Local Storage (GB)	200	200	200
Total Flash Capacity (TB)	38.4	76.8	153.6
Total Usable Disk Storage Capacity (TB)	73	168	336

> **Note:**
>
> Exadata X6 shapes must be provisioned using the License Included option. Bring-Your-Own-License (BYOL) is not supported with the X6 shapes.

Exadata Base System

An Exadata Base System is a fixed shape similar in size to a Quarter Rack with some differences in capacity.

Property	Configuration
Number of DB servers per System	2
Number of Storage Servers per System	3
Minimum Number of Enabled CPU Cores	0
Total Usable Cores in DB Servers per System	48
Total Memory Available (GB)	720
Max Usable Local Storage (GB)	900
Total Flash Capacity (TB)	38.4
Total Usable Disk Storage Capacity (TB)	73

For information on provisioning an Exadata Cloud Infrastructure instance, see Creating an Exadata Cloud Infrastructure Instance.

Exadata Cloud Management Interfaces

Exadata Cloud Infrastructure provides a variety of management interfaces to fit your use case and automation needs.

- Introduction to Exadata Cloud Management Interfaces
  The Exadata Cloud resources on Oracle Cloud Infrastructure (OCI) are created and managed through a variety of interfaces provided to fit your different management use cases.
- OCI Control Plane Interfaces for Exadata Cloud Infrastructure
  The OCI control plane accepts input from the OCI APIs, the OCI Console, and custom interfaces built with kits, tool and plugins provided to facilitate development and simplify the management of of OCI resources.
- Local VM Command-Line Interfaces
  In addition to the OCI REST-based APIs, CLI utilities located on the VM guests, provisioned as part of the VM clusters on the Exadata Cloud Infrastructure, are available to perform various lifecycle and administration operations.

Introduction to Exadata Cloud Management Interfaces

The Exadata Cloud resources on Oracle Cloud Infrastructure (OCI) are created and managed through a variety of interfaces provided to fit your different management use cases.

The various interfaces include:

- OCI Console interface and automation tools, see *Using the Console*
- Application Programming Interfaces (APIs)
- Command-Line Interfaces (CLIs)

The management interfaces are grouped into two primary categories:

- OCI Control Plane Interfaces
- Local Exadata Cloud VM CLIs

> **Note:**
>
> For more information and best practices on how these interfaces align for various Exadata Cloud database management use cases, refer to My Oracle Support note: *Exadata Cloud API/CLI Alignment Matrix (Doc ID 2768569.1)*.

Related Topics

- Oracle Database console overview
- Using the Console
- https://support.oracle.com/epmos/faces/DocContentDisplay?id=2768569.1

OCI Control Plane Interfaces for Exadata Cloud Infrastructure

The OCI control plane accepts input from the OCI APIs, the OCI Console, and custom interfaces built with kits, tool and plugins provided to facilitate development and simplify the management of of OCI resources.

The OCI APIs are typical REST APIs that use HTTPS requests and responses. The OCI Console, an intuitive, graphical interface for creating and managing your Exadata Cloud and other OCI resources, is one of the interfaces to the OCI APIs. When looking to develop automation utilizing the OCI APIs, a number of additional interfaces including: kits, tools and plug-ins, are provided to facilitate development and simplify the management of of OCI resources. A subset of these APIs applies to Exadata Cloud resources and the containing infrastructure. Each of these various interfaces provide the same functionality, all calling the OCI APIs, and are provided to enable flexibility and choice depending on preference and use case.

- **Command Line Interface (CLI):** The OCI CLI is a small footprint tool that you can use on its own or with the Console to perform Exadata Cloud resource tasks and other OCI tasks. The CLI provides the same core functionality as the Console, plus additional commands. Some of these, such as the ability to run scripts, extend the Console's functionality.
- **Software Development Kits (SDK):** OCI provides SDKs to enable you to develop custom solutions for your Exadata Cloud and other OCI based services and applications.
- **DevOps Tools and Plug-ins:** These tools can simplify provisioning and managing infrastructure, enable automated processes and facilitate development. Tools include the OCI Terraform Provider used with Resource Manager and OCI Ansible Collection.
- **Cloud Shell:** Cloud Shell is a free-to-use, browser-based terminal, accessible from the OCI Console, that provides access to a Linux shell with pre-authenticated OCI CLI and other useful developer tools. You can use the shell to interact with Exadata Cloud and other OCI resources, follow labs and tutorials, and quickly run OCI CLI commands.
- **Documentation: Appendix and Reference:** This general reference shows how to configure the SDKs and other developer tools to integrate with Oracle Cloud Infrastructure services.
- **Documentation: REST APIs:** This complete reference provides details on the Oracle Cloud Infrastructure REST APIs, including descriptions, syntax, endpoints, errors, and signatures. Exadata Cloud Infrastructure specific OCI REST APIs can be found throughout the documentation in the *Using the API* sections specific to each service:
  - *Using the API to Create Infrastructure Components*
  - *Using the API to Manage Exadata Cloud Service Instance*
  - *Using the API to manage database software images*
  - *Using the API to Create Oracle Database Home on Exadata Cloud Service*
  - *Using the API to Manage Oracle Database Home*
  - *Using the API to manage Databases*
  - *Using the API to Update the Grid Infrastructure on a VM Cluster Resources*
  - *Using the API to manage the I/O resources of an Exadata cloud VM cluster*
  - *Using the API to Patch an Exadata Cloud Service Instance*
  - *Using the API to upgrade Databases*
  - *Using the API to Manage Data Guard Associations*

- *Using the API to manage backups*

Related Topics

- Command Line Interface (CLI)
- Software Development Kits
- DevOps Tools and Plug-ins
- Terraform Provider
- Resource Manager
- Ansible Collection
- Cloud Shell
- Appendix and Reference
- REST APIs
- Using the API to Create Infrastructure Components
- Using the API to Manage Exadata Cloud Infrastructure Instance
- Using the API to manage database software images
  Use these API operations to manage database software images:
- Using the API to Create Oracle Database Home on Exadata Cloud Infrastructure
  To create an Oracle Database home, review the list of API calls.
- Using the API to Manage Oracle Database Home on Exadata Cloud Infrastructure
  Review the list of API calls to manage Oracle Database home.
- Using the API to manage Databases
- Using the API to Upgrade the Grid Infrastructure in a VM Cluster
- Using the API to manage the I/O resources of an Exadata cloud VM cluster
- Using the API to Patch an Exadata Cloud Infrastructure Instance
  Use these API operations to manage patching the following Exadata resources: cloud VM clusters, DB systems, databases, and Database Homes.
- Using the API to upgrade Databases
  Use the following APIs to manage database upgrades:
- Using the API to manage Data Guard associations
  Use these API operations to manage Data Guard associations on an Exadata Cloud Infrastructure instance:
- Using the API to manage backups

Local VM Command-Line Interfaces

In addition to the OCI REST-based APIs, CLI utilities located on the VM guests, provisioned as part of the VM clusters on the Exadata Cloud Infrastructure, are available to perform various lifecycle and administration operations.

The best practice is to use these utilities only when a corresponding Console command or OCI API is not available.

The utilities include:

- **dbaascli:** Use the `dbaascli` utility to perform various database lifecycle and administration operations on the Exadata Cloud Infrastructure such as

  - changing the password of a database user
  - starting a database
  - managing pluggable databases (PDBs)
- **bkup_api:** Use the `bkup_api` utility to perform various backup and recovery operations on the Exadata Cloud Infrastructure such as creating an on-demand backup of a complete database or an individual pluggable database (PDB), or to *customize backup settings* used by the automatic backup configuration

  > **Note:**
  >
  > `bkup_api` is deprecated. Use `dbaascli database backup`, `dbaascli pdb backup`, or `dbaascli pdb recover`

  instead.
- **ExaCLI:** Use the ExaCLI command-line utility to perform monitoring and management functions on Exadata storage servers in the Exadata Cloud.

These utilities are provided in addition to, and separate from, the OCI API-based interfaces listed above. To use the local VM command-line utilities, you must be connected to a virtual machine in an Exadata Cloud VM cluster and use the VM operating system user security, not the OCI user security, for execution. Most operations executed by these utilities sync their changes back to the OCI control plane using a process called `DB Sync`. However, there can be operations not synced with the control plane.

The cloud tooling software on the virtual machines, containing these CLI utilities, is automatically updated by Oracle on a regular basis. If needed, the tooling can be updated manually by following the instructions in *Updating Cloud Tooling Using dbaascli*.

Related Topics

- About Using the dbaascli Utility on Exadata Cloud Infrastructure
  You can use the dbaascli utility to perform various database lifecycle and administration operations on Exadata Cloud Infrastructure such as changing the password of a database user, starting a database, managing pluggable databases (PDBs), and more.
- Managing Exadata Database Backups by Using bkup_api
- To create a backup configuration file
- Monitoring and Managing Exadata Storage Servers with ExaCLI
  The ExaCLI command line utility allows you to perform monitoring and management functions on Exadata storage servers in an Exadata Cloud Infrastructure instance.
- Updating Cloud Tooling Using dbaascli
  To update the cloud tooling release for Oracle Exadata Database Service on Dedicated Infrastructure, complete this procedure.

2

What's New in Oracle Exadata Database Service on Dedicated Infrastructure

Oracle is constantly adding new capabilities to Exadata Cloud Infrastructure.

- Create and Manage Multiple Virtual Machines per Exadata System (MultiVM) and VM Cluster Node Subsetting
- VM Cluster and Database Health and Performance Metrics in the OCI Console
- Oracle Standard Tagging for Resources on Oracle Exadata Database Service on Dedicated Infrastructure
- Automatic Diagnostic Collection
- Exadata Database on Dedicated Infrastructure: Key Management Service for Cross Region Data Guard
- Concurrently Create or Terminate Oracle Databases in a VM Cluster
- VM Guest Exadata OS Image Major Version Update
- Database Service Events capability for Exadata Database
- Exadata Database on Dedicated Infrastructure: 'Create database from backup' now available for databases using customer-managed encryption
- Support for DB Home Minor Version Selection (N-3)
- Oracle Cloud Infrastructure Operations Insights Support for Oracle Cloud Databases
- Specify the Same SID for Primary and Standby Databases in Data Guard Association
- Exadata Cloud Infrastructure: Pluggable database lifecycle support
- Exadata Cloud Infrastructure: Set DB_UNIQUE_NAME and Oracle SID prefix during database creation
- Elastic Expansion
- Oracle Database: Encryption key options updated for bare metal, virtual machine, and Exadata Cloud Infrastructure databases
- Performance Hub Exadata Tab
- Exadata Cloud Infrastructure: custom SCAN listener port for VM cluster
- Performance Hub & metrics available for databases running in Exadata Cloud Infrastructure, bare metal DB systems, and virtual machine DB systems
- Maintenance advisory contacts for Exadata infrastructure
- Data Guard protection mode enhancement for Exadata, Bare Metal, and Virtual Machine Database Cloud services
- Exadata Cloud Infrastructure: Non-rolling infrastructure patching option now available
- Customer-managed encryption keys available with Oracle Data Guard-enabled databases in Exadata Cloud Infrastructure
- ExaDB-D OS/DomU Patching Project
- Oracle Cloud Infrastructure Vault service integration with Exadata Cloud Infrastructure
- Exadata Cloud Infrastructure: Oracle Database 19c upgrade feature available
- Create custom database software images for Exadata Cloud Infrastructure instances
- Exadata Cloud Infrastructure: grid infrastructure upgrade for cloud VM clusters
- Exadata Cloud Infrastructure: the flexible X8M shape now available
- Exadata Cloud Infrastructure: use an existing Database Home when setting up a Data Guard standby database

Use the Same Custom Software Image Across OCI Regions

- **Service**: Database
- **Release Date**: May 15, 2024

With this enhancement, you can use the software image created in one region in a different region while:

- updating a database software
- updating a Grid Infrastructure software
- provisioning a new Database Home
- provisioning a new database
- enabling a Data Guard association
- creating a database from a backup

Related Topics

- To update database software using custom database software image
  Use the following instructions to update database software using a custom database software image.
- To update Grid Infrastructure software using custom Grid Infrastructure software image
  Use the following instructions to update Grid Infrastructure software using a custom Grid Infrastructure software image.
- To create a new Database Home in an existing Exadata Cloud Infrastructure instance
  To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.
- To create a database in an existing Exadata Cloud Infrastructure instance
  This topic covers creating your first or subsequent databases.
- Using the Console to Enable Data Guard on an Exadata Cloud Infrastructure System
  Learn to enable Data Guard association between databases.
- To create a database from a backup

Ability to Increase the Size of Guest VM Local File Systems

- **Service**: Database
- **Release Date**: May 09, 2024

Currently, you can only increase or decrease the size of the `/u02` file system in the Guest VM. Now, using the OCI Console or API, you can increase the size of additional local file systems such as `/`, `/u01`, `/tmp`, `/var`, `/var/log`, `/var/log/audit`, and `/home`.

Related Topics

- Estimating How Much Local Storage You Can Provision to Your VMs
- The X8M and X9M Virtual Machine File System Structure Important File System and Sizes
- Scaling Local Storage
- To create a cloud VM cluster resource
  Create a VM cluster in an Exadata Cloud Infrastructure instance.

Create and Use Custom Software Images

- **Service**: Database
- **Release Date**: May 08, 2024

The ability to create a custom software image (Database and Grid Infrastructure) with all the required patches bundled together and certified in the customer environment will allow developers and database administrators to build an approved and reusable "gold image".

Related Topics

- Manage Software Images

Manage Serial Console Access to Oracle Exadata Database Service on Dedicated Infrastructure Systems

- **Service**: Database
- **Release Date**: May 07, 2024

> **Note:**
>
> The serial console feature requires (at a minimum) Exadata System Software 23.1.13. Once the necessary software is installed via Quarterly Maintenance and a reboot of your VMs takes place, you will be able to use the new serial console feature.

You can create and delete serial console connections to your Oracle Exadata Database Service on Dedicated Infrastructure systems to diagnose and resolve VM guest operating system issues using an SSH connection in case standard SSH access to the VMs is not possible.

Requirements: Exadata System Software 23.1.13 is the minimum required version. Also, make sure to review all prerequisites stated below, including setting a password for either the `opc` or the `root` user. Failure to make necessary changes to meet these requirements in advance will result in the inability to urgently connect to the serial console when the need arises when the VM is not otherwise accessible.

Related Topics

- Troubleshooting Virtual Machines Using Console Connections
  You can troubleshoot malfunctioning virtual machines using console connections. For example, a previously working Guest VM stops responding.
- Create the Virtual Machine Serial Console Connection
  Before you can make a local connection to the serial console, you need to create the virtual machine console connection.
- Resource-Types for Exadata Cloud Service Instances
- db-nodes
  Review the list of permissions and API operations for `db-nodes` resource-type.

- dbnode-console-connection
  Review the list of permissions and API operations for `dbnode-console-connection` resource-type.
- Serial Console Connection Event Types
  Review the list of event types that serial console connection emits.
- Viewing Audit Log Events
  Oracle Cloud Infrastructure Audit service provides records of API operations performed against supported services as a list of log events.
- Permissions Required for Each API Operation
  The following tables list the API operations for Exadata Cloud Infrastructure instances in a logical order, grouped by resource type.

Oracle Database 23ai on Exadata Database Service on Dedicated Infrastructure

- **Service**: Database
- **Release Date**: May 02, 2024

Oracle Database 23ai is a regular production release available on Oracle Exadata Database Service on Dedicated Infrastructure (ExaDB-D). With this release, you can perform all the lifecycle operations on the 23ai databases.

Related Topics

- To Upgrade the Oracle Grid Infrastructure of a Cloud VM Cluster
  Procedure for upgrading the Oracle Grid Infrastructure of a Cloud VM Cluster.
- Using the Console to Enable Data Guard on an Exadata Cloud Infrastructure System
  Learn to enable Data Guard association between databases.
- To create a cloud VM cluster resource
  Create a VM cluster in an Exadata Cloud Infrastructure instance.
- To create a new Database Home in an existing Exadata Cloud Infrastructure instance
  To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.

Enable Unified Auditing While Creating a Database Home

- **Service**: Database
- **Release Date**: April 30, 2024

With this enhancement, you can enable Unified Auditing, which has been available since Oracle Database version 12.2 while creating a database home.

- **Oracle Database version 12.1 or lower:** You cannot use the Unified Auditing framework. Instead, use the Traditional Audit - legacy Oracle Database audit framework.
- **Oracle Database version 12.2 or higher:** You can enable Unified Auditing from the OCI Console. For Oracle Database versions 12.2 or higher but lower than version 23, the **Unified Auditing** check box is not selected by default. However, it is selected by default for Oracle Database version 23.

> **Note:**
>
> You cannot disable Unified Auditing after provisioning the Database Home.

Related Topics

- To create a new Database Home in an existing Exadata Cloud Infrastructure instance
  To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.

Provision a VM Cluster with Either an OL7 or OL8-Based Image

- **Service**: Database
- **Release Date**: February 28, 2024

With this enhancement, you can provision a VM cluster with either an OL7-based image or an OL8-based one if the infrastructure is X9 or prior.

Related Topics

- To create a cloud VM cluster resource
  Create a VM cluster in an Exadata Cloud Infrastructure instance.

Enhancement to the OCI Console to Remove Database and Storage Servers

- **Service**: Database
- **Release Date**: January 29, 2024

With this enhancement, you can:

- Scale down your Exadata Infrastructure resources by changing your server count to a lower value than the current assignment.
- Scaling down to a lower count is supported for both DB and Storage Servers.
- Database servers will be removed if there are no VMs running on them.

> **Note:**
>
> You will not be able to choose the DB Server to remove. This functionality will automatically remove Database Servers in which there are no VMs.

- Storage server will be removed if the server has not been used to expand Exadata Infrastructure storage.
- Remove a VM from a provisioned VM cluster in a non multi-VM-enabled infrastructure. The procedure is similar to terminating a VM from a VM Cluster in a multi-VM-enabled infrastructure.

The 'Add Capacity' step runs as part of the storage server scale-up workflow, creates disk groups, and rebalances data across all storage servers. For more information, see *Scale VM Resources in Multi VM Enabled Infrastructure*.

Related Topics

- Scale VM Resources in Multi VM Enabled Infrastructure
  Increase or decrease the OCPUs, memory, storage or local disk size (`/u02`) storage available to a VM cluster

Enable Data Guard Across Different VCNs or Compartments in the Same OCI Region

- **Service**: Database
- **Release Date**: January 16, 2024

With this enhancement, you can enable Data Guard if the Exadata Cloud Infrastructure is in different VCNs or compartments in the same OCI region.

Enhancement to Pluggable Database (PDB) Management

- **Service**: Database
- **Release Date**: October 18, 2023

With this enhancement, you can restore, refresh, and relocate a Pluggable Database (PDB).

Related Topics

- Create and Manage Exadata Pluggable Databases
  You can create and manage pluggable databases (PDBs) in Exadata Cloud Infrastructure using the Console and APIs.
- Using the console to create pluggable database
- Using the console to relocate a pluggable database
- Cloning an Exadata Pluggable Database
  You can create local, remote, and refreshable clones.
- Restoring an Exadata Pluggable Database
  You can perfrom in-place and out of place restore of an Exadata pluggable database.
- Permissions Required for Each API Operation
  The following tables list the API operations for Exadata Cloud Infrastructure instances in a logical order, grouped by resource type.
- pluggable-databases (PDBs)
  Review the list of permissions and API operations for `pluggable-databases` resource-type.
- Pluggable Database Event Types
  These are the event types that Oracle pluggable databases in Oracle Cloud Infrastructure emit.

Manage Administrator (SYS User) and TDE Wallet Passwords

- **Service**: Database
- **Release Date**: September 28, 2023

With this enhancement, you can manage the administrator and TDE wallet passwords.

> **Note:**
>
> Changing a TDE wallet password for Oracle Key Vault (OKV) or OCI Vault Key management-enabled databases is currently not supported.

Related Topics

- To manage SYS user and TDE Wallet passwords
  Learn to manage administrator (SYS user) and TDE wallet passwords.

Backup and Restore from a Standby Database in a Data Guard Environment

- **Service**: Database
- **Release Date**: August 24, 2023

This enhancement:

- Enables customers to offload backups to the standby database in a Data Guard environment thereby freeing up resources in the production database environment.
- Allows customers to schedule automatic backups on the standby database in a Data Guard environment and configure retention period and backup schedules.
- Enables customers to create a database in another Availability Domain (AD) within the same region from a backup of the standby database.
- Allows customers to restore and recover a standby database using a backup of the standby database.
- Provides the flexibility to take backups only on the primary database, only on the standby database, or both.
- Allows customers to create a manual full backup of a standby database.
- Enables customers to enable or disable backup on the standby database only if the backup destination of the primary database is Object Storage.

Note:

- You cannot change the backup destination of the primary database to Autonomous Recovery Service if the backup destination of the primary and standby databases is Object Storage.
  To change the backup destination of the primary database to Autonomous Recovery Service, first, disable backup on the standby database.
- You cannot use standby-side backups to perform restore/recover operations on the primary database.
- Switchover scenarios:
  - If automatic backups was configured on the primary with backup destination of Object Storage, upon switchover, the backups will continue on the new standby database
  - If automatic backups was configured on the primary with backup destination of Autonomous Recovery Service, upon switchover, backup and restore will be disabled on the new standby database
  - If automatic backups was configured on the standby with backup destination of Object Storage, upon switchover, the backups will continue on the new primary database
- Failover scenarios:
  - If automatic backups was configured on the primary with backup destination of Object Storage or Autonomous Recovery Service, upon failover the backups be disabled on the new *Disabled Standby* database
  - If automatic backups was configured on the standby with backup destination of Object Storage, upon failover the backups will continue on the new primary database

Related Topics

- To enable automatic backups on a standby database
  Learn to enable automatic backups on a standby database.
- To restore a database
- To create a database from the latest backup

Cancel a Running Full or Incremental Backup

- **Service**: Database
- **Release Date**: August 21, 2023

You now have the ability to cancel an ongoing backup, allowing you to free up system resources. You will no longer have to call the operations team to have this backup job canceled.

As part of the Create Database workflow and independently (after the database has been created), you may enable Automatic Backup and select the desired backup destination. Depending on the backup destination selected, you may have one or more full backups and several incremental backups. Once any of these backups have started, you will not have the option to cancel that backup midway.

This feature allows you to cancel any running backup (automatic or standalone) from the OCI console or via OCI API.

You can also:

- Cancel a manual backup, which is triggered when you click the **Create backup** button **Note:** All manual backups are full backups.
- Delete a canceled manual backup

Related Topics

- Using the Console to Manage Backups
- To view backup status
- To cancel a backup

Autonomous Recovery Service as the Default Backup Destination

Service: Database

Release Date: August 17, 2023

This Console enhancement sets Autonomous Recovery Service as the default backup destination for automatic backups in all regions and includes default limits automatically without having to request them.

For more information about Service Limits, Quotas, and Usage, see Autonomous Recovery Service limits.

Related Topics

- To configure automatic backups for a database

Exadata Fleet Update

- **Services:** Database
- **Release Date:** August 02, 2023

Exadata Fleet Update simplifies, standardizes, and enhances the Oracle Database and Grid Infrastructure patching experience. Exadata Fleet Update achieves this by grouping components based on the customers' business needs into collections that can be patched as one entity within a given maintenance cycle.

Exadata Fleet Update brings this patching engine to OCI as a native cloud service, accessible from the OCI Console, OCI API, and via the OCI CLI.

Exadata Fleet Update is available free of charge on Oracle’s Exadata Database Service including Cloud@Customer (ExaDB-C@C) and Exadata Database Service on Dedicated Infrastructure (ExaDB-D).

For more information, see:

- Exadata Fleet Update Overview
- Exadata Fleet Update service AP

Update Guest VM (domU) Operating System to Oracle Linux 8

Service: Database

Release Date: August 01, 2023

Update the Guest VM operating system to Oracle Linus 8 using the Console or API. This enhancement is limited to Exadata X7, X8M, and X9M systems.

Related Topics

- Updating an Exadata Cloud VM Cluster Operating System
  Exadata VM cluster image updates allow you to update the OS image on your Exadata cloud VM cluster nodes in an automated manner from the OCI console and APIs.
- Supported Software Versions and Update Restrictions
  Minimum requirements for updating to Exadata image release 23.1.0.0.0 (Oracle Linux 8-based image):
- Updating the Operating System using the Console
- Add a VM to a VM Cluster
  Add a Virtual Machine to a VM Cluster

Use a Backup to Create a Database Across Availability Domains within the Same Region

Service: Database

Release Date: July 26, 2023

With this enhancement, when AD is up, you can:

- use an existing backup and restore it to create a database (out-of-place restore) either within the same availability domain or in a different availability domain within the same region, regardless of whether the backup was created with backup destination Object Storage or Autonomous Recovery Service
- restore a backup taken on either a database that was configured using host-based wallets (local wallet) or OCI Vault

Related Topics

- To create a database from a backup
- To create a database from the latest backup

Interim Software Updates

- **Services**: Database
- **Release Date**: June 07, 2023

This feature enables cloud-only customers to download one-off patches from the OCI console and API. There is no option to apply the downloaded patch via console and API. To apply these patches, customers must log in to their VM and run the patch apply utility.

> **Note:**
>
> To be able to download interim software update, you should at least have an ExaDB-D infrastructure provisioned.

Downloading one-off patches does not replace Database Software Image (DSI) creation. Customers must continue to use Database Software Images (DSI) to build and deploy their customized images.

Related Topics

- Interim Software Updates
  For authorized environments, learn how to download interim software updates.
- Permissions Required for Each API Operation
  The following tables list the API operations for Exadata Cloud Infrastructure instances in a logical order, grouped by resource type.
- oneoffPatch
  Review the list of permissions and API operations for `oneoffPatch` resource-type.
- Interim Software Updates Event Types
  These are the event types that Interim Software Updates in Oracle Cloud Infrastructure emit.

Enhanced Controls to Configure Automatic Full (L0) and Incremental (L1) Backups

Service: Database

Release Date: May 17, 2023

Enabling Automatic Backup during the Create Database workflow or as a separate step afterward starts the first full backup ("initial L0") immediately.

Similarly, for subsequent full backups (future L0) and daily incremental backups (L1), you can specify a time window but cannot change the day of the week when these backups must begin.

Future L0 and L1 backups will begin during the 2-hour scheduling window that the user selects for the database during which the automatic backup process will begin. There are 12 scheduling windows to choose from, each starting at an even-numbered hour. For example, one window runs from 4:00-6:00 AM, and the next from 6:00-8:00 AM. Backup jobs do not necessarily complete within the scheduled window. If you do not specify a window, the default 6-hour backup window of 00:00 to 06:00 is chosen. In this case, the time zone corresponds to the region of the Exadata Cloud infrastructure instance.

Here are the current defaults for the backup destinations, Object Storage Service, and Autonomous Recovery Service:

- **Initial full L0 backup:** Immediate
- **Subsequent full L0 backups:** Every Sunday
- **Daily incremental L1 backups:** Every Monday - Saturday

With these enhanced controls, you can:

1. Aside from configuring the initial L0 backup to start immediately, you can also specify whether you want the initial L0 backup to start immediately or according to the L0 schedule.
2. Choose a time window for the future full backups to start.
3. Choose a time window for the incremental backups to start, which can be different from the time window for the L0 backups.
   The time windows will remain the same, 2-hour scheduling windows and the default 6-hour window.

Related Topics

- To create a database in an existing Exadata Cloud Infrastructure instance
  This topic covers creating your first or subsequent databases.
- To configure automatic backups for a database

Configure Oracle Database Autonomous Recovery Service as a Backup Destination

Service: Database

Release Date: April 12, 2023

Oracle Database Autonomous Recovery Service provides an optimized policy-driven automatic backup and recovery system for the Exadata Database on Dedicated Infrastructure. It also offers a real-time data protection feature that enables protected databases with zero data loss recovery in the event of a database failure. Since Real-time data protection is an extra cost option, you can choose to enable or disable it.

Related Topics

- Overview of Oracle Database Autonomous Recovery Service
- Network Requirements for Oracle Database Autonomous Recovery Service
  Oracle Database Autonomous Recovery Service requires a registered Recovery Service subnet dedicated to backup and recovery operations in your database virtual cloud network (VCN).
- Create a Service Gateway to Object Storage
  In the OCI Console, create a service gateway to Object Storage. The service gateway is required for automation updates and configuration metadata.
- Creating Protection Policies
  Recovery Service uses protection policies to control database backup retention in Oracle Cloud.
- Managing Exadata Database Backups
  Automatic Exadata database backups are managed by Oracle Cloud Infrastructure. You configure this by using the Console or the API.
- Managed Backup Types and Usage Information
  There are two types of automatic Exadata database backups: Autonomous Recovery Service, and Oracle Object Storage.
- Prerequisites for Backups on Exadata Cloud Infrastructure
- To configure automatic backups for a database
- To create a database in an existing Exadata Cloud Infrastructure instance
  This topic covers creating your first or subsequent databases.

- To view details of a Protected Database
  To view the details of a Protected Database, use this procedure.
- To designate Autonomous Recovery Service as a Backup Destination for an Existing Database
  To designate Autonomous Recovery Service as a Backup Destination for an existing database, use this procedure.
- To terminate a database
- Using the Console to Terminate a VM Cluster
  Before you can terminate a VM cluster, you must first terminate the databases that it contains.
- Recovering an Exadata Database from Backup Destination
  This topic explains how to recover an Exadata database from a backup stored in either Object Storage or Autonomous Recovery Service by using the Console or the API.
- Using the Console to restore a database
  You can use the Console to restore the database from a backup in a backup destination that was created by using the Console.
- To restore a database
- To create a cloud VM cluster resource
  Create a VM cluster in an Exadata Cloud Infrastructure instance.

Application VIP Support

Service: Database

Release Date: April 12, 2023

The VM Cluster now supports attaching and detaching Application Virtual IP Addresses.

Related Topics

- About Application VIP
  Oracle Exadata Database Service on Dedicated Infrastructure fully supports creating additional Virtual IP Addresses on an Exadata VM Cluster.
- To Attach a Virtual IP Address
  Attach a Virtual IP address from a VM cluster using this procedure.
- To Detach a Virtual IP Address
  Attach a Virtual IP address from a VM cluster using this procedure.
- Resource-Types for Exadata Cloud Service Instances
- application-vips
  Review the list of permissions and API operations for `application-vips` resource-type.
- Permissions Required for Each API Operation
  The following tables list the API operations for Exadata Cloud Infrastructure instances in a logical order, grouped by resource type.
- Application VIP Event Types
  These are the event types that Application VIPs in Oracle Cloud Infrastructure emit.

Monthly ExaDB-D Infrastructure Security Maintenance

Service: Database

Release Date: March 01, 2023

Security maintenance, performed alongside the quarterly maintenance, is executed once a month and includes fixes for vulnerabilities with CVSS scores greater than or equal to 7.

> **Note:**
>
> The Infrastructure Security Maintenance implementation will be rolled out only to the `SCL`, `YUL`, `ZRH`, `AUH`, `DXB`, `YNY`, `HYD`, `ORD`, `MAD`, `MTZ`, and `ARN` regions. It will be rolled out to other regions in a phased manner.

Related Topics

- About Oracle-managed Exadata Cloud Infrastructure Maintenance
  Oracle performs patches and updates to all of the Oracle-managed system components on Exadata Cloud Infrastructure.
- Overview of Monthly Security Maintenance
  Security maintenance, performed alongside the quarterly maintenance, is executed in months when important security updates are needed and includes fixes for vulnerabilities across all CVSS scores.
- To set the automatic quarterly maintenance schedule for Exadata Cloud Infrastructure
- To view the maintenance history of an Exadata Cloud Infrastructure resource
  This task describes how to view the maintenance history for a cloud Exadata infrastructure or DB system. resource.

Identity and Access Management (IAM) Integration with Oracle Exadata Database Service on Dedicated Infrastructure

Service: Database

Release Date: January 24, 2023

With the latest Release Update, you can now configure the database in a virtual machine cluster to use Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) authentication and authorization to allow IAM users to access the database with IAM credentials.

As of this release, IAM authentication and authorization:

- will be available on newly provisioned and existing the databases patched to 19.17. This feature is not available on Oracle Database release 21c.
- can not be used with databases configured with Data Guard.

Related Topics

- Connect Identity and Access Management (IAM) Users to Oracle Exadata Database Service on Dedicated Infrastructure
  You can configure Oracle Exadata Database Service on Dedicated Infrastructure to use Oracle Cloud Infrastructure Identity and Access Management (IAM) authentication and authorization to allow IAM users to access an Oracle Database with IAM credentials.

Exadata Cloud Infrastructure: Private DNS

- **Services:**Database
- **Release Date:** January 18, 2023

Allow users to choose the private view and private zone while provisioning a new VM cluster for ExaCS. All the underlying resources of the VCN, including those of ExaDB-D, should be created in the same private zone. The private zones can be associated with subnets inside the VCN. This configuration cannot be changed later.

Private DNS resolver is going to resolve queries in customer VCN and queries coming from the on-premise networks. Ability to provide the DNS address and seed that into DB resources using conditional forwarding, is provided by the private DNS feature. With a private resolver, customers can resolve the A-record across different VCNs (with VCN peering- local/remote).

Related Topics

- Configure Private DNS
  Prerequistes needed to use Private DNS
- To create a cloud VM cluster resource
  Create a VM cluster in an Exadata Cloud Infrastructure instance.
- To view details about private DNS configuration

Enhanced Infrastructure Maintenance Controls

- **Services**: Database
- **Release Date**: January 18, 2023

The Exadata Cloud Infrastructure Oracle-managed infrastructure maintenance now allows greater control and visibility including:

- Choice of rolling and non-rolling maintenance methods.
- Ability to perform custom actions before maintenance on each database server by having the automated maintenance wait before shutting down VMs until the maintenance is resumed or the configured timeout is reached.
- Visibility into the database server update order.
- Granular tracking of the maintenance progress at a component level.

Related Topics

- To create a Cloud Exadata infrastructure resource
- Cloud Infrastructure Maintenance Updates
  Oracle performs the updates to all of the Oracle-managed infrastructure components on Exadata Cloud Infrastructure.
- To view or edit the properties of the next scheduled quarterly maintenance for Exadata Cloud Infrastructure
  Review and change the properties of the Exadata Cloud Infrastructure scheduled quarterly maintenance.
- Using the API to Manage Exadata Cloud Infrastructure Maintenance Controls
  Use these API operations to manage Exadata Cloud Infrastructure maintenance controls and resources.

- Oracle Exadata Database Service on Dedicated Infrastructure Event Types
  The events in this section are emitted by the cloud Exadata infrastructure resource

Database Management Support for Pluggable Databases in Oracle Exadata Database Service on Dedicated Infrastructure

Service: Database

Release Date: January 11, 2023

You can now enable Database Management for Pluggable Databases (PDBs) on Oracle Exadata Database Service on Dedicated Infrastructure, and use Database Management features for monitoring, performance management, and tuning.

Related Topics

- Monitor Metrics to Diagnose and Troubleshoot Problems with Pluggable Databases
  Enable Database Management service to view metrics to diagnose and troubleshoot problems with pluggable databases.
- databases (CDBs)
  Review the list of permissions and API operations for `databases` resource-type.
- pluggable-databases (PDBs)
  Review the list of permissions and API operations for `pluggable-databases` resource-type.
- Permissions Required for Each API Operation
  The following tables list the API operations for Exadata Cloud Infrastructure instances in a logical order, grouped by resource type.

Microsoft Azure Active Directory Integration with Oracle Cloud Infrastructure Databases

Service: Database

Release date: January 10, 2023

Oracle Exadata Database on Dedicated Infrastructure now can accept Azure AD tokens to access the database. Azure AD users can access the database directly using their Azure AD token, and applications can use their service tokens to access the database.

Azure AD integration will be available for databases patched to 19.17 and above. This feature is not available on Oracle Database release 21c.

Related Topics

- Authenticating and Authorizing Microsoft Azure Active Directory Users for Oracle Databases
  An Oracle Database instance can be configured for Microsoft Azure AD users to connect using Azure `OAuth2` access tokens.

Create and Manage Multiple Virtual Machines per Exadata System (MultiVM) and VM Cluster Node Subsetting

- **Services:** Database

- **Release Date:** Starting November 9, 2022 (the release date varies by region)

NOT_SUPPORTED

Slice Exadata resources into multiple virtual machines. Define up to 8 multiple virtual machine (VM) clusters on an Oracle Exadata Database Service on Dedicated Infrastructure, and specify how the overall system resources are allocated to them.

VM Cluster Node Subsetting enables you to allocate a subset of database servers to new and existing VM clusters to enable maximum flexibility in the allocation of compute (CPU, memory, local storage) resources.

> **Note:**
>
> For existing Exadata Infrastructure, MultiVM will be enabled as part of your next scheduled maintenance run after MultiVM migration on December 20, 2022. All newly provisioned Exadata Infrastructure after the release of MVM on November 15, 2022, will have MultiVM enabled.

Related Topics

- To create a cloud VM cluster resource
- Scale VM Resources in Multi VM Enabled Infrastructure
- Exadata Cloud Infrastructure VM Cluster Event Types
- Overview of VM Cluster Node Subsetting
- Add a VM to a VM Cluster
- Using the Console to View a List of DB Servers on an Exadata Infrastructure
- Using the Console to Terminate a VM Cluster
- Terminate a VM from a VM Cluster
- Permissions Required for Each API Operation
- dbServers
- VM Node Subsetting Event Types
- Adding a VM to a VM Cluster Fails
- CPU Offline Scaling Fails

VM Cluster and Database Health and Performance Metrics in the OCI Console

- **Services:** Database
- **Release Date:** October 7, 2022

With this release Oracle will provide health metrics for databases and VM clusters in the Oracle Cloud Infrastructure (OCI) console.

> **Note:**
>
> When there is a network problem and Oracle Trace File Analyzer (TFA) is unable to post metrics, TFA will wait for one hour before attempting to retry posting the metrics. This is required to avoid creating a backlog of metrics processing on TFA.
>
> Potentially one hour of metrics will be lost between network restore and the first metric posted.

Related Topics

- Monitor Metrics for VM Cluster Resources
- Prerequisites for Using Metrics
- View Metrics for VM Cluster
- View Metrics for a Database
- View Metrics for VM Clusters in a Compartment
- View Metrics for Databases in a Compartment
- Metrics for Oracle Exadata Database Service on Dedicated Infrastructure in the Monitoring Service
  This article describes the metrics emitted by the Exadata Cloud Infrastructure Database service in the `oci_database_cluster` and `oci_database` namespaces for Oracle Databases.
- Metrics for Exadata Cloud Infrastructure in the Database Management Service
  Database Management provides comprehensive database performance diagnostics and management capabilities for Oracle Databases.

Oracle Standard Tagging for Resources on Oracle Exadata Database Service on Dedicated Infrastructure

- **Services**: Database
- **Release Date**: September 15, 2022

Exadata Cloud Infrastructure resources can now be tagged using Oracle Standard tags according to your organizational scheme. By tagging resources, you can group them, manage costs, and gain insight into how they are being used.

Related Topics

- Tagging Oracle Exadata Database Service on Dedicated Infrastructure Resources
  Tagging is a powerful foundational service for Oracle Cloud Infrastructure (OCI) that enables users to search, control access, and do bulk actions on a set of resources based on the tag.

Automatic Diagnostic Collection

- **Services**: Database
- **Release Date**: August 31, 2022

This feature extends the Database Service Events feature implementation that enables you to get notified about health issues with your Oracle Databases or other components on the Guest VM. With this enhancement, you can allow:

- Oracle to proactively collect detailed health metrics for diagnosis and issue resolution
- Oracle to reactively collect Incident logs and trace files on demand for a deeper diagnosis and issue resolution

Collecting Guest VM events, health metrics, incident logs, and trace files, will help Oracle to enhance service operations as well as provide proactive support by early detection and correlation.

Related Topics

- Overview of Automatic Diagnostic Collection
  By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
- Incident Logs and Trace Files
  This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.
- Health Metrics
  Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.
- Database Service Events
  The Database Service emits events, which are structured messages that indicate changes in resources.
- To create a cloud VM cluster resource
  Create a VM cluster in an Exadata Cloud Infrastructure instance.
- Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.

Exadata Database on Dedicated Infrastructure: Key Management Service for Cross Region Data Guard

- **Services:** Database
- **Release Date:** Aug 8, 2022

You can now have the encryption keys used for the primary and standby databases to be available in the primary and standby regions respectively so that it provides protection against a single point of failure for the OCI Vault key. This is possible if the keys are in OCI Virtual Private Vault. So, cross-region Data Guard can be set up between two databases if their keys are residing in a virtual private vault (VPV) and are managed by the OCI Vault service.

Related Topics

- To administer Vault encryption keys
  Use this procedure to rotate the Vault encryption key or or change the encryption management configuration.

- Customer-Managed Keys in Exadata Cloud Infrastructure
  Customer-managed keys for Exadata Cloud Infrastructure is a feature of Oracle Cloud Infrastructure (OCI) Vault service that enables you to encrypt your data using encryption keys that you control.
- Prerequisites for Using Oracle Data Guard with Exadata Cloud Infrastructure
  An Exadata Cloud Infrastructure Oracle Data Guard implementation requires two existing Exadata VM Clusters: one containing an existing database that is to be duplicated by Data Guard, and one that will house the new standby database by Data Guard.

Concurrently Create or Terminate Oracle Databases in a VM Cluster

- **Services**: Database
- **Release Date**: August 3, 2022

With this enhancement, you can now concurrently create or terminate Oracle databases even if the VM cluster is in the Updating state.

- The number of databases that can be created on a cluster depends on the available memory on the VMs. For each database, by default, 12.6 GB (7.6 GB for SGA and 5 GB for PGA) is allocated if the VM has greater than 60 GB of memory. If the VM has less than or equal to 60 GB, then 6.3 GB (3.8 GB for SGA and 2.5 GB for PGA) is allocated. Also, Grid Infrastructure and ASM consume some memory, approximately 2 to 4 GB.
- A database that is being created cannot be terminated. You can, however, terminate other databases in the VM Cluster.

VM Guest Exadata OS Image Major Version Update

- **Services:** Database
- **Release Date:** July 11, 2022

In addition to performing minor version updates to the Exadata VM Cluster images, you can update to a new major version if the currently installed version is 19.2 or higher. For example, if the VM cluster is on version 20, then you can update it to version 21.

Related Topics

- Supported Software Versions and Update Restrictions
  Minimum requirements for updating to Exadata image release 23.1.0.0.0 (Oracle Linux 8-based image):

Database Service Events capability for Exadata Database

- **Services:** Database
- **Release Date:** July 8, 2022

This feature allows customers to use OCI Console or API/CLI/SDK/Terraform to receive event notifications about health issues with your Oracle Databases or other components on the Guest VM.

Customers currently have basic lifecycle management events like backup begin, backup end, patching begin, etc. We are extending that capability to include a comprehensive set of Database Service events to help customers troubleshoot issues.

Database Service Events monitor guest VM operations and conditions and generate diagnostic notifications for customers by leveraging the existing OCI Events service and Notification mechanisms in their tenancy. Customers can then create topics and subscribe to these topics through email, functions, streams, etc. For more information about using the Events Notification Service, see *Notifications Overview*

Key Customer Benefits

- Ability to receive notifications for Guest VM operations via an opt-in mechanism.
- Allows customers to proactively address issues before they may become serious.

OCI Console Experience

Customers can navigate to the VM Cluster details page from the OCI console menu by selecting Oracle Database → **Oracle Exadata Database Service on Dedicated Infrastructure** → a specific VM Cluster to enable Diagnostics Notification for a VM Cluster.

Related Topics

- Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.
- To create a cloud VM cluster resource
  Create a VM cluster in an Exadata Cloud Infrastructure instance.
- Overview of Database Service Events
- Receive Notifications about Database Service Events
  Subscribe to the Database Service Events and get notified.
- Database Service Event Types
  Review the list of event types that the Database Service emits.
- Temporarily Restrict Automatic Diagnostic Collections for Specific Events
  Use the `tfactl blackout` command to temporarily suppress automatic diagnostic collections.
- To create a cloud VM cluster resource
  Create a VM cluster in an Exadata Cloud Infrastructure instance.
- Notifications Overview

Exadata Database on Dedicated Infrastructure: 'Create database from backup' now available for databases using customer-managed encryption

- **Services:** Database
- **Release Date:** June 29, 2022

Oracle Exadata Database Service on Dedicated Infrastructure (ExaDB-D): Now allows you to create a database from backup, when the backup is of the database using customer-managed encryption. This is in addition to the existing ability to create a database from backup, when the backup is of the database using oracle-managed encryption

Related Topics

- To create a database from a backup

Support for DB Home Minor Version Selection (N-3)

- **Services**: Database
- **Release Date**: May 23, 2022

Provision a DB Home using a major version and RU version of your choice.

While provisioning, if you opt to use **Oracle Provided Database Software Images** as the image type, then you can use the **Display all available versions** switch to choose from all available PSUs and RUs. The most recent release for each major version is indicated with a **latest** label.

For the Oracle Database major version releases available in Oracle Cloud Infrastructure, images are provided for the current version plus the three most recent older versions (N through N - 3). For example, if an instance is using Oracle Database 19c, and the latest version of 19c offered is 19.8.0.0.0, images available for provisioning are for versions 19.8.0.0.0, 19.7.0.0, 19.6.0.0 and 19.5.0.0.

Related Topics

- To create a database in an existing Exadata Cloud Infrastructure instance
  This topic covers creating your first or subsequent databases.
- To create a new Database Home in an existing Exadata Cloud Infrastructure instance
  To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.

Oracle Cloud Infrastructure Operations Insights Support for Oracle Cloud Databases

- **Services**:Support for Oracle Database Cloud Service Databases
- **Release Date**: March 22, 2022

Operations Insights now allows you to use the Capacity Planning and SQL Warehouse functionality to gain insight into Oracle Databases deployed in Oracle Cloud (Bare Metal, Virtual Machine VM, and Exadata Cloud Infrastructure).

Related Topics

- Enabling Database Cloud Service Databases
- Oracle Cloud Infrastructure Operations Insights
  Oracle Cloud Infrastructure Operations Insights allows you to use the Capacity Planning and SQL Warehouse functionality to gain insight into Oracle Databases deployed in Oracle Cloud (Bare Metal, Virtual Machine VM, and Exadata Cloud Infrastructure).

Specify the Same SID for Primary and Standby Databases in Data Guard Association

- **Services**: Database

- **Release Date**: March 14, 2022

The same SID prefix used for the primary database can now also be used for the standby database when creating a Data Guard Association.

Related Topics

- Using the Console to Enable Data Guard on an Exadata Cloud Infrastructure System
  Learn to enable Data Guard association between databases.

Exadata Cloud Infrastructure: Pluggable database lifecycle support

- **Services**: Database
- **Release Date**: Jan. 12, 2022

You can now create and manage pluggable databases (PDBs) in Exadata Cloud Infrastructure using the OCI console and APIs. See Create and Manage Exadata Pluggable Databases for details.

Exadata Cloud Infrastructure: Set DB_UNIQUE_NAME and Oracle SID prefix during database creation

- **Services**: Database
- **Release Date**: January 12, 2022

You can now specify the `DB_UNIQUE_NAME` value and the Oracle SID prefix when creating a new Oracle Database in Exadata Cloud Infrastructure. You can also set these values when creating a standby database in an Oracle Data Guard association. See the following topics for instructions:

To create a database in an existing Exadata Cloud Infrastructure instance

Using the Console to Enable Data Guard on an Exadata Cloud Infrastructure System

Elastic Expansion

Services: Database

Release Date: December 23, 2021

With elastic provisioning and expansion, you can dynamically increase your CPU and storage capacity to meet your growing workload requirements.

Expand the infrastructure capacity on-demand by scaling up the infrastructure with additional database or storage servers without being constrained by the standard supported shapes. You can allocate CPU and storage capacity available on X8M and X9M servers up to the system limits when you provision new VM Clusters on the infrastructure, or to already deployed VM Clusters without disrupting the current running workloads.

For more information, see:

- *To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster or DB system*
- *To add compute and storage resources to a flexible cloud Exadata infrastructure resource*

- *Using the API to Manage Exadata Cloud Infrastructure Instance*
- *cloud-exadata-infrastructures*
- *Storage Expansion Event Types*

Related Topics

- Scaling CPU cores within an Exadata Cloud Infrastructure instance
  If an Exadata Cloud Infrastructure instance requires more compute node processing power, you can scale up the number of enabled CPU cores symmetrically across all the nodes in the system as follows:
- Exadata X9M
  The values in the table that follows represent the specifications for an X9M cloud instance with 2 database and 3 storage servers that has not been expanded.
- Exadata X8M
  The values in the table that follows represent the specifications for an X8M cloud instance with 2 database and 3 storage servers that has not been expanded.
- The Cloud Exadata Infrastructure Resource
  The infrastructure resource is the top-level (parent) resource.
- Resources to Be Created
- To create a Cloud Exadata infrastructure resource
- To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster or DB system
- Scaling Exadata X8M and X9M Compute and Storage
  The flexible X8M and X9M system model is designed to be easily scaled in place, with no need to migrate the database using a backup or Data Guard
- Permissions Required for Each API Operation
  The following tables list the API operations for Exadata Cloud Infrastructure instances in a logical order, grouped by resource type.
- cloud-exadata-infrastructures
  Review the list of permissions and API operations for `cloud-exadata-infrastructures` resource-type.

Oracle Database: Encryption key options updated for bare metal, virtual machine, and Exadata Cloud Infrastructure databases

- **Services**: Database
- **Release Date**: Nov. 17, 2021

When provisioning a new Oracle Database or a new virtual machine or bare metal DB system, the TDE wallet password parameter is now optional, and it is not used if you configure an Exadata Cloud Infrastructure database to use customer-managed keys with the OCI Vault service. For more information on creating Oracle Databases and virtual machine DB systems, see the following topics in the documentation:

- Creating Bare Metal and Virtual Machine DB Systems
- Creating Databases (bare metal DB systems)
- Creating and Managing Exadata Databases

For information on using the Vault service to store and manage encryption keys and other secrets used by OCI resources, see the Vault service documentation.

Performance Hub Exadata Tab

- **Services**: Database
- **Release Date**: Nov. 16, 2021

The Exadata tab provides a unified view of Oracle Exadata hard disk and flash performance statistics with deep insight into the health and performance of all components such as the Oracle databases, Oracle Exadata storage cells, and Automatic Storage Management (ASM). It is available for Exadata Cloud deployments.and external databases that use Exadata infrastructure. For more information, see Using Performance Hub to Analyze Database Performance.

Exadata Cloud Infrastructure: custom SCAN listener port for VM cluster

- **Services**: Database
- **Release Date**: Nov. 3, 2021

You can now specify a custom SCAN listener port for your Exadata cloud VM cluster. See SCAN Listener Port Setting for more information.

Performance Hub & metrics available for databases running in Exadata Cloud Infrastructure, bare metal DB systems, and virtual machine DB systems

- **Services**: Database
- **Release Date**: Sept. 9, 2021

You can now use the Performance Hub tool and view metrics on cloud databases that run on the following systems: Exadata Cloud Infrastructure instances, bare metal DB systems, and virtual machines DB systems. This feature provides additional monitoring and management functions to these databases. For more information, see Analyzing Exadata Cloud Service Database Performance and Analyzing Virtual Machine / Bare Metal Database Performance .

Maintenance advisory contacts for Exadata infrastructure

- **Services**: Database
- **Release Date**: July 28, 2021

You can choose to specify up to 10 valid email addresses to which Oracle sends maintenance notifications when updates are made to an Exadata infrastructure. The email addresses you specify are used only for service-related operational issues. See Oracle-Managed Infrastructure Maintenance Updates for more information.

Data Guard protection mode enhancement for Exadata, Bare Metal, and Virtual Machine Database Cloud services

- **Services**: Database
- **Release Date**: July 21, 2021

You can now specify the Data Guard protection mode for Exadata Cloud Infrastructure instances and bare metal and virtual machine DB systems. See the following topics for more information:

- Use Oracle Data Guard with Exadata Cloud Infrastructure
- Using Oracle Data Guard (for bare metal and virtual machine DB systems)

Exadata Cloud Infrastructure: Non-rolling infrastructure patching option now available

- **Services**: Database
- **Release Date**: May 26, 2021

You can now configure Exadata infrastructure patching to take place in a non-rolling fashion across the database nodes. This option allows you to reduce the total time your system is undergoing quarterly maintenance, but does involve system downtime. See Oracle-Managed Infrastructure Maintenance Updates for more information.

Customer-managed encryption keys available with Oracle Data Guard-enabled databases in Exadata Cloud Infrastructure

- **Services**: Database
- **Release Date**: April 16, 2021

Customer-managed keys for Exadata Cloud Infrastructure is an encryption key management service that enables you to encrypt your data using encryption keys that you control. You can use customer-managed keys on databases you provision in Exadata Cloud Infrastructure that are enabled with Oracle Data Guard.

ExaDB-D OS/DomU Patching Project

- **Services**: Database
- **Release Date**: Feb. 5, 2021

DomU OS Patching is a feature which enables ExaDB-D customers to upgrade the Exadata OS image on their domU nodes in an automated manner from their OCI console and APIs. The following information explains about a recent change in the feature that could not be added to the docs in time for release, but will be added soon.

Rollback required if the patch fails.

On multi-node systems, if one of the nodes fails the patch, you must roll back all nodes to get them all on the same version. Then run precheck, fix any problems, and run the patch again.

Example: If you run precheck on Monday and all nodes pass, but do not apply the patch until Wednesday, it is possible that one or more of the nodes may fail the patch because of changes on the nodes or a maintenance conflict.

To prevent this from happening, Oracle recommends that you run precheck right before applying the patch.

For more information, see Updating an Exadata Cloud Service VM Cluster Operating System.

Oracle Cloud Infrastructure Vault service integration with Exadata Cloud Infrastructure

- **Services**: Database, Key Management, Vault
- **Release Date**: Dec. 18, 2020

Oracle Cloud Infrastructure Vault service integration with Exadata Cloud Infrastructure enables database encryption with customer-managed keys. For more information, see Customer-Managed Keys in Exadata Cloud Service.

Exadata Cloud Infrastructure: Oracle Database 19c upgrade feature available

- **Services**: Database
- **Release Date**: Dec. 3, 2020
- **API Versions Affected**: 20160918

You can now upgrade Exadata Cloud Infrastructure databases to Oracle Database version 19c using the Oracle Cloud Infrastructure Console or API. For information and instructions, see Upgrading Exadata Databases.

Create custom database software images for Exadata Cloud Infrastructure instances

- **Services**: Database
- **Release Date**: Nov. 10, 2020

You can now create custom Oracle Database software images to use for provisioning Database Homes and and patching databases in Exadata Cloud Infrastructure instances. For more information, see Oracle Database Software Images.

Exadata Cloud Infrastructure: grid infrastructure upgrade for cloud VM clusters

- **Services**: Database

- **Release Date**: Oct. 15, 2020

You can now upgrade the grid infrastructure (GI) of an Exadata Cloud Infrastructure VM cluster using the Console. For more information, see Upgrading Exadata Grid Infrastructure.

Exadata Cloud Infrastructure: the flexible X8M shape now available

- **Services**: Database
- **Release Date**: Oct. 15, 2020

You can now provision an Exadata Cloud Infrastructure instance using the flexible X8M shape. This shape allows you to expand your system after provisioning, as your databases grow and you need more storage servers, compute servers, or both. For more information, see Overview of X8M Scalable Exadata Infrastructure.

Exadata Cloud Infrastructure: use an existing Database Home when setting up a Data Guard standby database

- **Services**: Database
- **Release Date**: Oct. 15, 2020

You can now choose to use an existing Database Home when setting up a Data Guard standby database in your Exadata Cloud Infrastructure instance. See Using Oracle Data Guard with Exadata Cloud Service Instances for information on setting up Data Guard for your Exadata databases.

3

Preparing for Exadata Cloud Infrastructure

Review OCI as well as the site, network and storage requirements to prepare and deploy Exadata Cloud Infrastructure in your data center.

- Oracle Cloud Infrastructure (OCI) Requirements for Oracle Exadata Database Service on Dedicated Infrastructure
  Learn the basic concepts to get started using Oracle Cloud Infrastructure.
- Network Setup for Exadata Cloud Infrastructure Instances
  This topic describes the recommended configuration for the VCN and several related requirements for the Exadata Cloud Infrastructure instance.
- Creating Protection Policies
  Recovery Service uses protection policies to control database backup retention in Oracle Cloud.
- Storage Configuration
  The storage space inside the Exadata storage servers is configured for use by Oracle Automatic Storage Management (ASM) When you launch an Exadata Cloud Infrastructure instance.

Oracle Cloud Infrastructure (OCI) Requirements for Oracle Exadata Database Service on Dedicated Infrastructure

Learn the basic concepts to get started using Oracle Cloud Infrastructure.

Exadata Cloud Infrastructure is managed by the Oracle Cloud Infrastructure (OCI) control plane. The Exadata Cloud Infrastructure resources are deployed in your OCI Tenancy.

Before you can provision Exadata Cloud Infrastructure infrastructure, your Oracle Cloud Infrastructure tenancy must be enabled to use Oracle Exadata Database Service on Dedicated Infrastructure. Review the information in this publication for further details.

The following tasks are common for all OCI deployments, refer to the links in the Related Topics to find the associated Oracle Cloud Infrastructure documentation.

- Getting Started with OCI.
  If you are new to OCI, learn the basic concepts to get started by following the *OCI Getting Started Guide* .
- Setting Up Your Tenancy.
  After Oracle creates your tenancy in OCI, an administrator at your company will need to perform some set up tasks and establish an organization plan for your cloud resources and users. The information in this topic will help you get started.
- Managing Regions
  This topic describes the basics of managing your region subscriptions.
- Managing Compartments
  This topic describes the basics of working with compartments.
- Managing Users
  This topic describes the basics of working with users.

- Managing Groups
  This topic describes the basics of working with groups.
- Required IAM Policy for Exadata Cloud Infrastructure
  Review the identity access management (IAM) policy for provisioning Oracle Exadata Database Service on Dedicated Infrastructure systems.

Related Topics

- OCI Getting Started Guide
- Setting Up Your Tenancy
- Managing Regions
- Managing Compartments
- Managing Users
- Managing Groups

Required IAM Policy for Exadata Cloud Infrastructure

Review the identity access management (IAM) policy for provisioning Oracle Exadata Database Service on Dedicated Infrastructure systems.

A **policy** is an IAM document that specifies who has what type of access to your resources. It is used in different ways:

- An individual statement written in the policy language
- A collection of statements in a single, named "policy" document, which has an Oracle Cloud ID (OCID) assigned to it
- The overall body of policies your organization uses to control access to resources

A **compartment** is a collection of related resources that can be accessed only by certain groups that have been given permission by an administrator in your organization.

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console, or the REST API with a software development kit (SDK), a command-line interface (CLI), or some other tool. If you try to perform an action, and receive a message that you don’t have permission, or are unauthorized, then confirm with your tenancy administrator the type of access you've been granted, and which compartment you should work in.

For administrators: The policy in "Let database admins manage DB systems" lets the specified group do everything with databases, and related database resources.

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

For more details on writing policies specific to Exadata Cloud@Customer resources see "Policy Details for Exadata Cloud Infrastructure".

Related Topics

- Let database admins manage DB systems
- Getting Started with Policies
- Common Policies
- Policy Details for the Database Services

- Policy Details for Exadata Cloud Infrastructure
  This topic covers details for writing policies to control access to Exadata Cloud Infrastructure resources.

Creating Protection Policies

Recovery Service uses protection policies to control database backup retention in Oracle Cloud.

Protection Policies provide automated retention management for protected databases, satisfying requirements for regulated environments. Each protected database must be associated with one protection policy.

A protection policy determines the maximum period (in days) allowed to retain backups created by Recovery Service. Based on your business requirements, you can assign separate policies for each protected database or use a single policy across all protected databases in a VCN. For more information, see *About Configuring Protection Policies*.

To use the Oracle Cloud Infrastructure (OCI) Console to configure and manage protection policies, follow the steps outlined in *Creating a Protection Policy*.

Related Topics

- About Configuring Protection Policies
- Create Protection Policy

Network Setup for Exadata Cloud Infrastructure Instances

This topic describes the recommended configuration for the VCN and several related requirements for the Exadata Cloud Infrastructure instance.

Before you set up an Exadata Cloud Infrastructure instance, you must set up a virtual cloud network (VCN) and other Networking service components.

- VCN and Subnets
  To launch an Exadata Cloud Infrastructure instance, you must have a Virtual Cloud Network and at least two subnets:
- Node Access to Object Storage: Static Route
- Service Gateway for the VCN
  Your VCN needs access to both Object Storage for backups and Oracle YUM repos for OS updates.
- Security Rules for the Oracle Exadata Database Service on Dedicated Infrastructure
  This section lists the security rules to use with Exadata Cloud Infrastructure.
- Ways to Implement the Security Rules
  Learn how to implement security rules within your VCN using the networking service.
- Network Requirements for Oracle Database Autonomous Recovery Service
  Oracle Database Autonomous Recovery Service requires a registered Recovery Service subnet dedicated to backup and recovery operations in your database virtual cloud network (VCN).

VCN and Subnets

To launch an Exadata Cloud Infrastructure instance, you must have a Virtual Cloud Network and at least two subnets:

To launch an Exadata Cloud Infrastructure instance, you must have a Virtual Cloud Network, at least two subnets and select the type of DNS resolver you will use:

- A VCN in the region where you want the Exadata Cloud Infrastructure instance
- At least two subnets in the VCN. The two subnets are:
  - Client subnet
  - Backup subnet
- Choose which method of DNS name resolution you will use. See *Choices for DNS in Your VCN*

Note

For Exadata Cloud Infrastructure instances using The New Exadata Cloud Infrastructure Resource Model , networking is configured on the cloud VM cluster resource.

In general, Oracle recommends using **regional subnets** , which span all **availability domains** in the region. For more information, see Overview of VCNs and Subnets.

You will create custom route tables route tables for each subnet. You will also create security rulessecurity rules to control traffic to and from the client network and backup network of the Exadata compute nodes (for The Cloud VM Cluster Resource, nodes are called virtual machines). More information follows about those items.

- Option 1: Public Client Subnet with Internet Gateway
  This option can be useful when doing a proof-of-concept or development work.
- Option 2: Private Subnets
  Oracle recommends private subnets for a production system.
- Requirements for IP Address Space
  IP addresses must not overlap, especially when Exadata Cloud Infrastructure VM Clusters (and thus VCNs) are in more than one region.
- Configuring a Static Route for Accessing the Object Store
- Setting Up DNS for an Exadata Cloud Infrastructure Instance
  DNS lets you use host names instead of IP addresses to communicate with an Exadata Cloud Infrastructure instance.
- DNS: Short Names for the VCN, Subnets, and Exadata Cloud Infrastructure instance
- DNS: Between On-Premises Network and VCN
  Oracle recommends using a private DNS resolver to enable the use of hostnames when on-premises hosts and VCN resources communicate with each other.
- Configure Private DNS
  Prerequistes needed to use Private DNS

Related Topics

- Choices for DNS in Your VCN
- Overview of VCNs and Subnets
- About Regions and Availability Domains
- Availability Domains and Your VCN

Option 1: Public Client Subnet with Internet Gateway

This option can be useful when doing a proof-of-concept or development work.

You can use this setup in production if you want to use an **internet gateway** with the VCN, or if you have services that run only on a public network and need access to the database. See the following diagram and description.

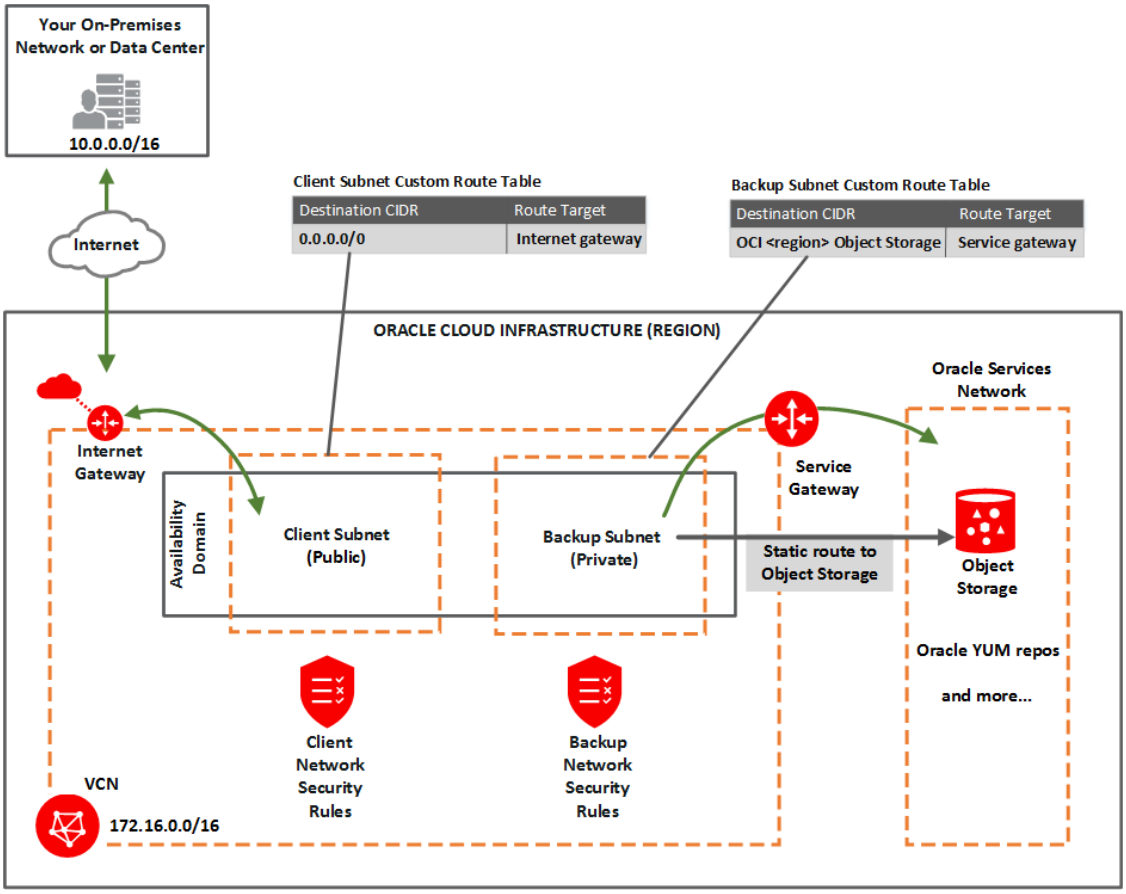

You set up:

- Subnets:
  - *Public* client subnet (*public* means that the resources in the subnet can have public IP addresses at your discretion).
  - *Private* backup subnet (*private* means that the resources in the subnet cannot have public IP addresses and therefore cannot receive incoming connections from the internet).
- Gateways for the VCN:
  - Internet gateway (for use by the client subnet).
  - Service gateway (for use by the backup subnet). Also see Option 1: Service Gateway Access to OCI Services .
- Route tables:
  - Custom route table for the public client subnet, with a route for 0.0.0.0/0, and target = the internet gateway.
  - Separate custom route table for the private backup subnet, with a route rule for the service CIDR labels (see about CIDR labels under Overview of Service Gateways and Available Sevice CIDR labels, and target = the service gateway. Also see Option 1: Service Gateway Access to OCI Services.

- Security rules to enable the desired traffic to and from the Exadata virtual machines compute nodes. See Security Rules for the Oracle Exadata Database Service on Dedicated Infrastructure.
- Node Access to Object Storage: Static Route on the Exadata Cloud Service instance's compute nodes (to enable access to OCI services by way of the backup subnet).

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

Option 2: Private Subnets

Oracle recommends private subnets for a production system.

Both subnets are private and cannot be reached from the internet. See the following diagram and description.

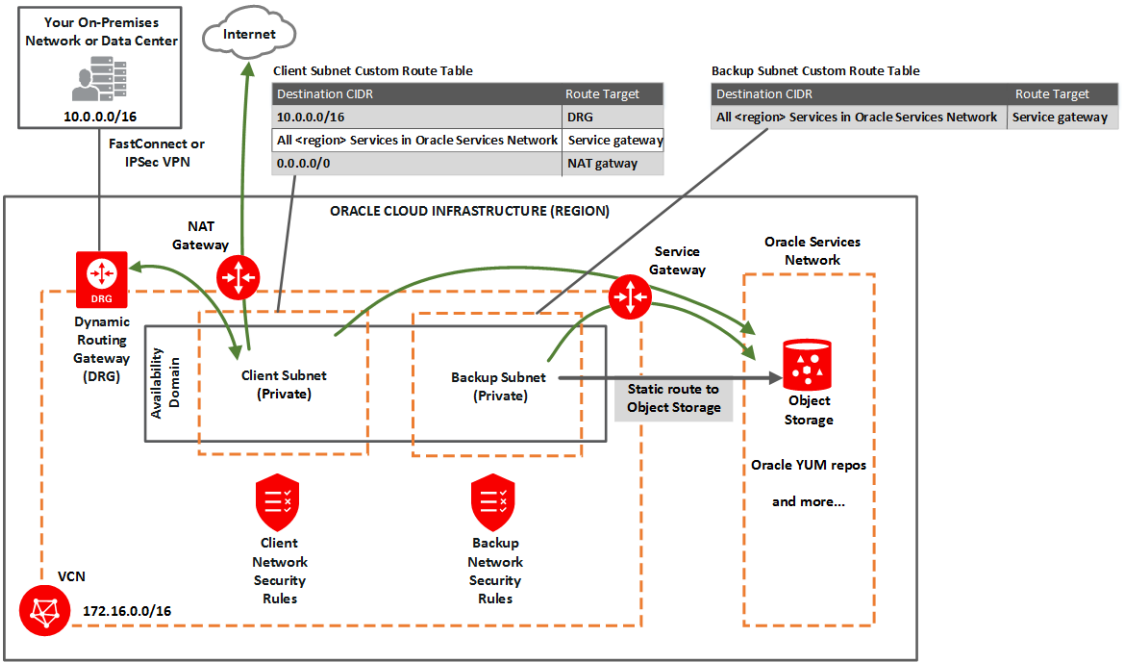

You set up:

- Subnets:
  - *Private* client subnet.
  - *Private* backup subnet.
- Gateways for the VCN:
  - Dynamic routing gateway (DRG), with a FastConnect or IPSec VPN to your on-premises network (for use by the client subnet).
  - Service gateway For use by the backup and client subnets to reach OCI Services, such as Object Storage for backups, YUM for OS updates, IAM (Identitiy Access Management) and OCI Vault (KMS Integration) Also see Option 2: Service Gateway Access to Both Object Storage and YUM Repos.

  - NAT gateway(*optional*) For use by the client subnet to reach public endpoints not supported by the service gateway.

- Route tables:
  - Custom route table for the private client subnet, with the following rules:
    - A rule for the on-premises network's CIDR, and target = DRG.
    - A rule for the service CIDR label called **All <region> Services in Oracle Services Network**, and target = the service gateway. The *Oracle Services Network* is a conceptual network in Oracle Cloud Infrastructure that is reserved for Oracle services. The rule enables the client subnet to reach the regional Oracle YUM repository for OS updates. Also see Option 2: Service Gateway Access to Both Object Storage and YUM Repos.
    - Optionally, a rule for 0.0.0.0/0, and target = NAT gateway.
  - Separate custom route table for the private backup subnet, with one rule:
    - The same rule as for the client subnet: for the service CIDR label called **All <region> Services in Oracle Services Network**, and target = the service gateway. This rule enables the backup subnet to reach the regional Object Storage for backups.
- Security rules to enable the desired traffic to and from the Exadata nodes. See Security Rules for the Exadata Cloud Service instance.
- Optionally add a Static route on the compute nodes to other OCI services (for VM clusters, the virtual machines) to enable access, if the services are only reachable on the backup subnet and not via. the client subnet, e.g. when using a NAT Gateway.

Requirements for IP Address Space

IP addresses must not overlap, especially when Exadata Cloud Infrastructure VM Clusters (and thus VCNs) are in more than one region.

If you're setting up Exadata Cloud Infrastructure VM Clusters (and thus VCNs) in more than one region, make sure the IP address space of the VCNs does not overlap. This is important if you want to set up disaster recovery with Oracle Data Guard.

For Exadata X8 and lower, the two subnets you create for the Exadata Cloud Infrastructure VM Clusters must not overlap with 192.168.128.0/20.

For Exadata X8M and X9M, IP addresses (100.106.0.0/16 and 100.107.0.0/16) are used for the interconnect.

The following table lists the minimum required subnet sizes, depending on the Exadata VM Infrastructure for each VM Cluster size. For the client subnet, each node requires four IP addresses, and in addition, three addresses are reserved for Single Client Access Names (SCANs). For the backup subnet, each node requires three addresses.

Rack Size	Client Subnet: # Required IP Addresses	Client Subnet: Minimum Size Note	Backup Subnet: # Required IP Addresses	Backup Subnet: Minimum Size Note
Base System or Quarter Rack per VM Cluster	(4 addresses * 2 nodes) + 3 for SCANs = 11	/28 (16 IP addresses)	3 address * 2 nodes = 6	/28 (16 IP addresses)
Half Rack per VM Cluster	(4 * 4 nodes) + 3 = 19	/27 (32 IP addresses)	3 * 4 nodes = 12	/28 (16 IP addresses)

Rack Size	Client Subnet: # Required IP Addresses	Client Subnet: Minimum Size Note	Backup Subnet: # Required IP Addresses	Backup Subnet: Minimum Size Note
Full Rack per VM Cluster	(4* 8 nodes) + 3 = 35	/26 (64 IP addresses)	3 * 8 nodes = 24	/27 (32 IP addresses)
Flexible infrastructure systems (X8M and higher) per VM Cluster	4 addresses per database node (minimum of 2 nodes) + 3 for SCANs	Minimum size determined by total number of IP addresses needed for database nodes	3 address per database node (minimum of 2 nodes)	Minimum size determined by total number of IP addresses needed for database nodes

Note:

The Networking service reserves three IP addresses in each subnet. Allocating a larger space for the subnet than the minimum required (for example, at least /25 instead of /28) can reduce the relative impact of those reserved addresses on the subnet's available space.

Configuring a Static Route for Accessing the Object Store

All the traffic in an Exadata Cloud Infrastructure instance is, by default, routed through the data network. To route backup traffic to the backup interface (BONDETH1), you need to configure a static route on *each* of the compute nodes in the cluster.
For instructions, see Node Access to Object Storage: Static Route.

Setting Up DNS for an Exadata Cloud Infrastructure Instance

DNS lets you use host names instead of IP addresses to communicate with an Exadata Cloud Infrastructure instance.

You can use the **Internet and VCN Resolver** (the DNS capability built into the VCN) as described in DNS in Your Virtual Cloud Network. Oracle recommends using a VCN Resolver for DNS name resolution for the client subnet. It automatically resolves the Swift endpoints required for backing up databases, patching, and updating the cloud tooling on an Exadata instance.

DNS: Short Names for the VCN, Subnets, and Exadata Cloud Infrastructure instance

For the nodes to communicate, the VCN must use the Internet and VCN Resolver. The Internet and VCN resolver enables hostname assignment to the nodes, and DNS resolution of those hostnames by resources in the VCN.
The Internet and VCN resolver enables round robin resolution of the database's SCANs. It also enables resolution of important service endpoints required for backing up databases, patching, and updating the cloud tooling on an Exadata Cloud Infrastructure instance. The Internet and VCN Resolver is the VCN's default choice for DNS in the VCN. For more information, see DNS in Your Virtual Cloud Network and also DHCP Options.

When you create the VCN, subnets, and Exadata, you must carefully set the following identifiers, which are related to DNS in the VCN:

- VCN domain label
- Subnet domain label

- Hostname prefix for the Exadata Cloud Infrastructure instance's cloud VM cluster or DB system resource

These values make up the node's fully qualified domain name (FQDN):

`<hostname_prefix>-######.<subnet_domain_label>.<vcn_domain_label>.oraclevcn.com`

For example:

```
exacs-abcde1.clientpvtad1.acmevcniad.oraclevcn.com
```

In this example, you assign `exacs` as the hostname prefix when you create the cloud VM cluster or DB system. The Database service automatically appends a hyphen and a five-letter string with the node number at the end. For example:

- Node 1: `exacs-abcde1.clientpvtad1.acmevcniad.oraclevcn.com`
- Node 2: `exacs-abcde2.clientpvtad1.acmevcniad.oraclevcn.com`
- Node 3: `exacs-abcde3.clientpvtad1.acmevcniad.oraclevcn.com`
- And so on

Requirements for the hostname prefix:

- Recommended maximum: 12 characters. For more information, see the example under the following section, "Requirements for the VCN and subnet domain labels".
- Cannot be the string *localhost*

Requirements for the VCN and subnet domain labels:

- Recommended maximum: 14 characters each. The actual underlying requirement is a total of 28 characters *across both domain labels* (excluding the period between the labels). For example, both of these are acceptable: `subnetad1.verylongvcnphx` or `verylongsubnetad1.vcnphx`. For simplicity, the recommendation is 14 characters each.
- No hyphens or underscores.
- Recommended: include the region name in the VCN's domain label, and include the availability domain name in the subnet's domain label.
- In general, the FQDN has a maximum total limit of 63 characters. Here is a safe general rule:

  `<12_chars_max>-######.<14_chars_max>.<14_chars_max>.oraclevcn.com`

The preceding maximums are not enforced when you create the VCN and subnets. However, if the labels exceed the maximum, the Exadata deployment fails.

DNS: Between On-Premises Network and VCN

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

See Private DNS resolvers for information on creating and using private resolvers. For a reference architecture see Use private DNS in your VCN in the Oracle Architecture Center.

Configure Private DNS

Prerequistes needed to use Private DNS

Private view and private zone must be created before launching DB system provisioning. For details see *Private DNS resolvers*

Forwarding to another DNS server should be set up beforehand in the DNS console. This can be done by going to the VCN's resolver, and creating the endpoint and then the rules. See *DNS in Your Virtual Cloud Network*

Related Topics

- Private DNS resolvers
- DNS in Your Virtual Cloud Network

Node Access to Object Storage: Static Route

To be able to back up databases, and patch and update cloud tools on an Exadata Cloud Infrastructure instance, you must configure access to Oracle Cloud Infrastructure Object Storage. Regardless of how you configure the VCN with that access (for example, with a service gateway), you may also need to configure a static route to Object Storage on each of the compute nodes in the cluster. This is only required if you are not using automatic backups. If you are using customized backups using the backup APIs, then you must route traffic destined for Object Storage through the backup interface (BONDETH1). This is not necessary if you are using the automatic backups created with the Console, APIs, or CLIs.

> **Caution:**
>
> You must configure a static route for Object Storage access on each compute node in an Exadata Cloud Infrastructure instance if you *are not* creating automatic backups with the Console, APIs, or CLIs. Otherwise, attempts to back up databases, and patch or update tools on the system, can fail.

> **Note:**
>
> When you enable the first automatic backup for a database the static route configuration will be automatically done on the service.
>
> If you want to patch the service before creating a database, the manual static route is required to be able to patch the GI or DB Home.
>
> The static route may also be required to access other services (IAM, KMS) if these are not reachable via client subnet and only the backup subnet uses the setting to access all servcies within a region.

- Object Storage IP allocations
- To configure a static route for Object Storage access

Object Storage IP allocations

Oracle Cloud Infrastructure Object Storage uses the CIDR block IP range 134.70.0.0/16 for all regions.

As of June 1, 2018, Object Storage no longer supports the following discontinued IP ranges. Oracle recommends that you remove these older IP addresses from your access-control lists, firewall rules, and other rules after you have adopted the new IP ranges.

The **discontinued** IP ranges are:

- Germany Central (Frankfurt): 130.61.0.0/16
- UK South (London): 132.145.0.0/16
- US East (Ashburn): 129.213.0.0/16
- US West (Phoenix): 129.146.0.0/16

To configure a static route for Object Storage access

1. SSH to a compute node in the Exadata Cloud Infrastructure instance.

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

2. Log in as opc and then sudo to the root user. Use `sudo su -` with a hyphen to invoke the root user's profile.

   ```
   login as: opc

   [opc@dbsys ~]$ sudo su -
   ```

3. Identify the gateway configured for the BONDETH1 interface.

   ```
   [root@dbsys ~]# grep GATEWAY /etc/sysconfig/network-scripts/ifcfg-bondeth1
   |awk -F"=" '{print $2}'


   10.0.4.1
   ```

4. Add the following static rule for BONDETH1 to the `/etc/sysconfig/network-scripts/route-bondeth1` file:

   ```
   10.0.X.0/XX dev bondeth1 table 211
   default via <gateway> dev bondeth1 table 211
   134.70.0.0/17 via <gateway_from_previous_step> dev bondeth1
   ```

5. Restart the interface.

   ```
   [root@dbsys ~]# ifdown bondeth1; ifup bondeth1;
   ```

   The file changes from the previous step take effect immediately after the ifdown and ifup commands run.

6. Repeat the preceding steps on *each* compute node in the Exadata Cloud Infrastructure instance.

Service Gateway for the VCN

Your VCN needs access to both Object Storage for backups and Oracle YUM repos for OS updates.

Depending on whether you use Option1: Public Client Subnet with Internet Gateway or Option 2: Private Subnets described previously, you use the service gateway in different ways. See the next two sections.

- Option 1: Service Gateway Access to OCI Services
  You configure the *backup subnet* to use the service gateway for access only to Object Storage.
- Option 2: Service Gateway Access to Both Object Storage and YUM Repos
  You configure *both the client subnet and backup subnet* to use the service gateway for access to the Oracle Services Network, which includes both Object Storage and the Oracle YUM repos.

Related Topics

- Option 1: Public Client Subnet with Internet Gateway
  This option can be useful when doing a proof-of-concept or development work.
- Option 2: Private Subnets
  Oracle recommends private subnets for a production system.

Option 1: Service Gateway Access to OCI Services

You configure the *backup subnet* to use the service gateway for access only to Object Storage.

As a reminder, here's the diagram for option 1:

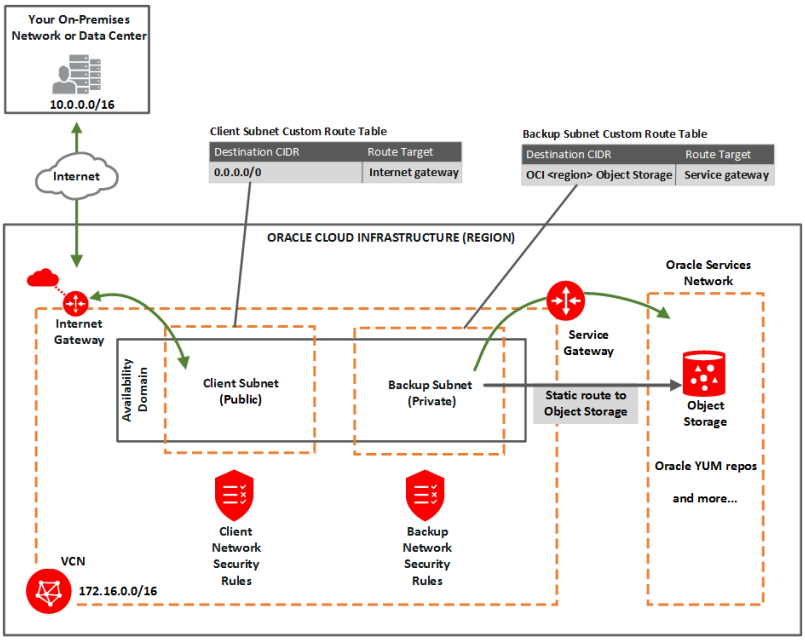

In general, you must:

- Perform the *tasks for setting up a service gateway on a VCN*, and specifically enable the service CIDR label called **OCI *<region>* Object Storage**.
- In the task for updating routing, add a route rule to the *backup* subnet's custom route table. For the destination service, use **OCI *<region>* Object Storage** and target = the service gateway.
- In the task for updating security rules in the subnet, perform the task on the *backup* network's network security group (NSG) or custom security list. Set up a security rule with the destination service set to **OCI *<region>* Object Storage**. See *Rule Required Specifically for the Backup Network* Rule Required Specifically for the Backup Network .

Related Topics

- Tasks for Setting Up a Service Gateway on a VCN in the Console

- Rule Required Specifically for the Backup Network
  The following security rule is important for the backup network because it enables the DB system to communicate with Object Storage through the service gateway (and optionally with the Oracle YUM repos if the client network doesn't have access to them).

Option 2: Service Gateway Access to Both Object Storage and YUM Repos

You configure *both the client subnet and backup subnet* to use the service gateway for access to the Oracle Services Network, which includes both Object Storage and the Oracle YUM repos.

> **Note:**
>
> See this known issues for information about accessing Oracle YUM services through the service gateway.

As a reminder, here's the diagram for option 2:

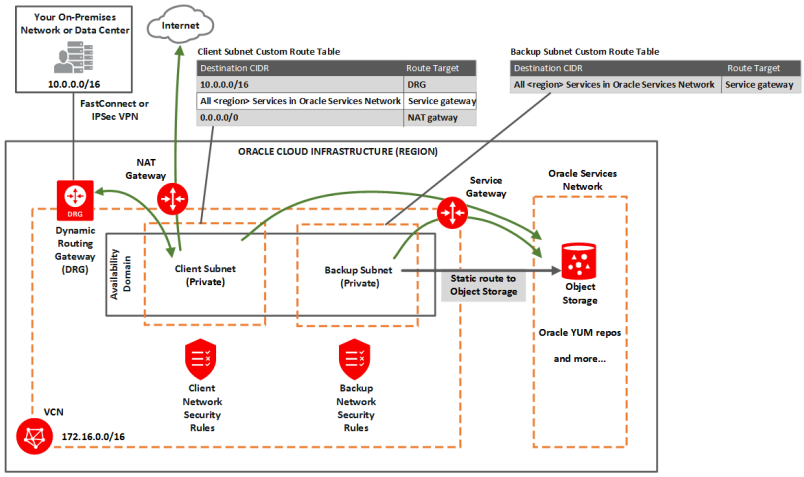

In general, you must:

- Perform the *tasks for setting up a service gateway on a VCN*, and specifically enable the service CIDR label called **All *<region>* Services in Oracle Services Network**.
- In the task for updating routing in each subnet, add a rule to each subnet's custom route table. For the destination service, use **All *<region>* Services in Oracle Services Network** and target = the service gateway.
- In the task for updating security rules for the subnet, perform the task on the *backup* network's network security group (NSG) or custom security list. Set up a security rule with the destination service set to **OCI *<region>* Object Storage**. See *Security Rules for the Oracle Exadata Cloud Service* Security Rules for the Oracle Exadata Database Service on Dedicated Infrastructure. Note that the client subnet already has a broad egress rule that covers access to the YUM repos.

Here are a few additional details about using the service gateway for option 2:

- Both the client subnet and backup subnet use the service gateway, but to access different services. You cannot enable both the **OCI *<region>* Object Storage** service CIDR label and the **All *<region>* Services in Oracle Services Network** for the service gateway. To cover the needs of both subnets, you must enable **All *<region>* Services in Oracle**

Services Network for the service gateway. The VCN can have only a single service gateway.

- Any route rule that targets a given service gateway must use an enabled service CIDR label and not a CIDR block as the destination for the rule. That means for option 2, the route tables for both subnets must use **All *<region>* Services in Oracle Services Network** for their service gateway rules.
- Unlike route rules, security rules can use either *any* service CIDR label (whether the VCN has a service gateway or not) or a CIDR block as the source or destination CIDR for the rule. Therefore, although the backup subnet has a route rule that uses **All *<region>* Services in Oracle Services Network**, the subnet can have a security rule that uses **OCI *<region>* Object Storage**. See *Security Rules for the Exadata Cloud Service instance*.

Related Topics

- Oracle Service Gateway
- Tasks for Setting up a Service Gateway on a VCN
- Security Rules for the Oracle Exadata Database Service on Dedicated Infrastructure
  This section lists the security rules to use with Exadata Cloud Infrastructure.

Security Rules for the Oracle Exadata Database Service on Dedicated Infrastructure

This section lists the security rules to use with Exadata Cloud Infrastructure.

Security rules control the types of traffic allowed for the client network and backup network of the Exadata's compute nodes. The rules are divided into three sections.

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

> **Note:**
>
> For X8M and X9M systems, Oracle recommends that all ports on the client subnet need to be open for ingress and egress traffic. This is a requirement for adding additional database servers to the system.

Rules Required for Both the Client Network and Backup Network

This section has several general rules that enable essential connectivity for hosts in the VCN.

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

General ingress rule 1: Allows SSH traffic from anywhere

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

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

General ingress rule 2: Allows Path MTU Discovery fragmentation messages

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

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

General ingress rule 3: Allows connectivity error messages within the VCN

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

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** Your VCN's CIDR
- **IP Protocol:** ICMP
- **Type:** 3
- **Code:** All

General egress rule 1: Allows all egress traffic

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

Rules Required Specifically for the Client Network

The following security rules are important for the client network.

Important:

- Client ingress rules 1 and 2 only cover connections initiated from within the client subnet. If you have a client that resides *outside the VCN*, Oracle recommends setting up two *additional* similar rules that instead have the **Source CIDR** set to the public IP address of the client.
- Client ingress rules 3 and 4 and client egress rules 1 and 2 allow TCP and ICMP traffic inside the client network and enable the nodes to communicate with each other. If TCP connectivity fails across the nodes, the Exadata cloud VM cluster or DB system resource fails to provision.

Client ingress rule 1: Allows ONS and FAN traffic from within the client subnet

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

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

Client ingress rule 2: Allows SQL*NET traffic from within the client subnet

This rule is for SQL*NET traffic and is required in these cases:

- If you need to enable client connections to the database
- If you plan to use Oracle Data Guard
- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** Client subnet's CIDR
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 1521
- **Description:** An optional description of the rule.

Client egress rule 1: Allows all TCP traffic inside the client subnet

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** 0.0.0.0/0
- **IP Protocol:** TCP
- **Source Port Range:** All

- **Destination Port Range:** 22
- **Description:** An optional description of the rule.

Client egress rule 2: Allows all egress traffic (allows connections to the Oracle YUM repos)

Client egress rule 3 is important because it allows connections to the Oracle YUM repos. It is redundant with the general egress rule in this topic (and in the default security list). It is optional but recommended in case the general egress rule (or default security list) is inadvertently changed.

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** 0.0.0.0/0
- **IP Protocol:** All
- **Description:** An optional description of the rule.

Rule Required Specifically for the Backup Network

The following security rule is important for the backup network because it enables the DB system to communicate with Object Storage through the service gateway (and optionally with the Oracle YUM repos if the client network doesn't have access to them). It is redundant with the general egress rule in this topic (and in the default security list). It is optional but recommended in case the general egress rule (or default security list) is inadvertently changed.

Backup egress rule: Allows access to Object Storage

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** Service
- **Destination Service:**
  - The service CIDR label called **OCI *<region>* Object Storage**
  - If the client network does not have access to the Oracle YUM repos, use the service CIDR label called **All *<region>* Services in Oracle Services Network**
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 443 (HTTPS)
- **Description:** An optional description of the rule.

- Rules Required for Both the Client Network and Backup Network
  This topic has several general rules that enable essential connectivity for hosts in the VCN.
- Rules Required Specifically for the Client Network
  The following security rules are important for the client network.
- Rule Required Specifically for the Backup Network
  The following security rule is important for the backup network because it enables the DB system to communicate with Object Storage through the service gateway (and optionally with the Oracle YUM repos if the client network doesn't have access to them).
- Rules Required for Events Service
  The compute instance must have either a public IP address or a service gateway to be able to send compute instance metrics to the Events service.

- Rules Required for Monitoring Service
  The compute instance must have either a public IP address or a service gateway to be able to send compute instance metrics to the Monitoring service.

Rules Required for Both the Client Network and Backup Network

This topic has several general rules that enable essential connectivity for hosts in the VCN.

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

- General ingress rule 1: Allows SSH traffic from anywhere
- General ingress rule 2: Allows Path MTU Discovery fragmentation messages
- General ingress rule 3: Allows connectivity error messages within the VCN
  This rule enables the hosts in the VCN to receive connectivity error messages from each other.
- General egress rule 1: Allows all egress traffic

Related Topics

- default security list

General ingress rule 1: Allows SSH traffic from anywhere

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** 0.0.0.0/0
- **IP Protocol:** SSH
- **Source Port Range:** All
- **Destination Port Range:** 22

General ingress rule 2: Allows Path MTU Discovery fragmentation messages

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

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

General ingress rule 3: Allows connectivity error messages within the VCN

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

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

General egress rule 1: Allows all egress traffic

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

If the customer enables notification of Data Plane Guest VM Events, the default egress rule is sufficient.

Rules Required Specifically for the Client Network

The following security rules are important for the client network.

> **Note:**
>
> - For X8M systems, Oracle recommends that all ports on the client subnet need to be open for ingress and egress traffic. This is a requirement for adding additional database servers to the system.
> - Client ingress rules 1 and 2 only cover connections initiated from within the client subnet. If you have a client that resides *outside the VCN*, Oracle recommends setting up two *additional* similar rules that instead have the **Source CIDR** set to the public IP address of the client.
> - Client ingress rules 3 and 4 and client egress rules 1 and 2 allow TCP and ICMP traffic inside the client network and enable the nodes to communicate with each other. If TCP connectivity fails across the nodes, the Exadata cloud VM cluster or DB system resource fails to provision.

- Client ingress rule 1: Allows ONS and FAN traffic from within the client subnet
  The first rule is recommended and enables the Oracle Notification Services (ONS) to communicate about Fast Application Notification (FAN) events.
- Client ingress rule 2: Allows SQL*NET traffic from within the client subnet
  This rule is for SQL*NET traffic and is required in these cases:
- Client egress rule 1: Allows all TCP traffic inside the client subnet
  This rule is for SQL*NET traffic as noted.
- Client egress rule 2: Allows all egress traffic (allows connections to the Oracle YUM repos)
  Client egress rule 3 is important because it allows connections to the Oracle YUM repos.

Client ingress rule 1: Allows ONS and FAN traffic from within the client subnet

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

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

Client ingress rule 2: Allows SQL*NET traffic from within the client subnet

This rule is for SQL*NET traffic and is required in these cases:

- If you need to enable client connections to the database
- If you plan to use Oracle Data Guard
- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** Client subnet's CIDR
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 1521
- **Description:** An optional description of the rule.

Client egress rule 1: Allows all TCP traffic inside the client subnet

This rule is for SQL*NET traffic as noted.

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

Client egress rule 2: Allows all egress traffic (allows connections to the Oracle YUM repos)

Client egress rule 3 is important because it allows connections to the Oracle YUM repos.

It is redundant with the general egress rule 1: Allow all egress traffic (and in the *default security list*). It is optional but recommended in case the general egress rule (or default security list) is inadvertently changed.

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** 0.0.0.0/0
- **IP Protocol:** All
- **Description:** An optional description of the rule.

Related Topics

- default security list

Rule Required Specifically for the Backup Network

The following security rule is important for the backup network because it enables the DB system to communicate with Object Storage through the service gateway (and optionally with the Oracle YUM repos if the client network doesn't have access to them).

It is redundant with the *general egress rule 1: Allows all egress traffic* in (and in the). It is optional but recommended in case the general egress rule (or default security list) is inadvertently changed.

- Backup egress rule: Allows access to Object Storage

Related Topics

- General egress rule 1: Allows all egress traffic
- default security list

Backup egress rule: Allows access to Object Storage

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** Service
- **Destination Service:**
  - The service CIDR label called **OCI *<region>* Object Storage**
  - If the client network does not have access to the Oracle YUM repos, use the service CIDR label called **All *<region>* Services in Oracle Services Network**
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 443 (HTTPS)
- **Description:** An optional description of the rule.

Rules Required for Events Service

The compute instance must have either a public IP address or a service gateway to be able to send compute instance metrics to the Events service.

The default egress rules are sufficient to to allow the compute instance to send compute instance metrics to the Events service.

If the instance does not have a public IP address, set up a service gateway on the virtual cloud network (VCN). The service gateway lets the instance send compute instance metrics to the Events service without the traffic going over the internet. Here are special notes for setting up the service gateway to access the Events service:

- When creating the service gateway, enable the service label called **All <region> Services in Oracle Services Network**. It includes the Events service.
- When setting up routing for the subnet that contains the instance, set up a route rule with **Target Type** set to **Service Gateway**, and the **Destination Service** set to **All <region> Services in Oracle Services Network**.

  For detailed instructions, see Access to Oracle Services: Service Gateway.

Rules Required for Monitoring Service

The compute instance must have either a public IP address or a service gateway to be able to send compute instance metrics to the Monitoring service.

The default egress rules are sufficient to to allow the compute instance to send compute instance metrics to the Monitoring service.

If the instance does not have a public IP address, set up a service gateway on the virtual cloud network (VCN). The service gateway lets the instance send compute instance metrics to the Monitoring service without the traffic going over the internet. Here are special notes for setting up the service gateway to access the Monitoring service:

- When creating the service gateway, enable the service label called **All <region> Services in Oracle Services Network**. It includes the Monitoring service.
- When setting up routing for the subnet that contains the instance, set up a route rule with **Target Type** set to **Service Gateway**, and the **Destination Service** set to **All <region> Services in Oracle Services Network**.

  For detailed instructions, see Access to Oracle Services: Service Gateway.

Ways to Implement the Security Rules

Learn how to implement security rules within your VCN using the networking service.

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

- Network security groups
- Security lists

For a comparison of the two methods, see Comaprison of Security Lists and Network Security Groups.

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

If you use network security groups

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

1. Create an NSG for the client network. Add the following security rules to that NSG:
   - The rules listed in Rules Required for Both the Client Network and Backup Network
   - The rules listed in Rules Required Specifically for the Client Network
2. Create a separate NSG for the backup network. Add the following security rules to that NSG:
   - The rules listed in Rules Required for Both the Client Network and Backup Network

- The rules listed in Rules Required Specifically for the Client Network

3. When the database administrator is Creating an Exadata Cloud Infrastructure Instance, they must choose several networking components (for example, which VCN and subnets to use):
   - When they choose the client subnet, they can also choose which NSG or NSGs to use. Make sure they choose the client network's NSG.
   - When they choose the backup subnet, they can also choose which NSG or NSGs to use. Make sure they choose the backup network's NSG.

You could instead create a separate NSG for the general rules. Then when the database administrator chooses which NSGs to use for the client network, make sure they choose both the general NSG and the client network NSG. Similarly for the backup network, they choose both the general NSG and the backup network NSG.

If you use security lists

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

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

1. Configure the client subnet to use the required security rules:
   a. Create a custom security list for the client subnet and add the rules listed in Rules Required Specifically for the Client Network.
   b. Associate the following two security lists with the client subnet:
      - VCN's *default security list* with all its default rules. This automatically comes with the VCN. By default it contains the rules in Rules Required for Both the Client Network and Backup Network.
      - The new custom security list you created for the client subnet.
2. Configure the backup subnet to use the required security rules:
   a. Create a custom security list for the backup subnet and add the rules listed in Rule Required Specifically for the Backup Network.
   b. Associate the following two security lists with the backup subnet:
      - VCN's *default security list* with all its default rules. This automatically comes with the VCN. By default it contains the rules in Rules Required for Both the Client Network and Backup Network.
      - The new custom security list you created for the backup subnet.

Later when the database administrator creates the Exadata Cloud Service instance, they must choose several networking components. When they select the client subnet and backup subnet that you've already created and configured, the security rules are automatically enforced for the nodes created in those subnets.

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

Network Requirements for Oracle Database Autonomous Recovery Service

Oracle Database Autonomous Recovery Service requires a registered Recovery Service subnet dedicated to backup and recovery operations in your database virtual cloud network (VCN).

To use Recovery Service for backups, follow the steps outlined in *Configuring your Tenancy for Recovery Service*.

- Create a Service Gateway to Object Storage
  In the OCI Console, create a service gateway to Object Storage. The service gateway is required for automation updates and configuration metadata.

Related Topics

- Configuring your Tenancy for Recovery Service

Create a Service Gateway to Object Storage

In the OCI Console, create a service gateway to Object Storage. The service gateway is required for automation updates and configuration metadata.

1. Open the navigation menu. Click **Networking**, and then click **Virtual Cloud Networks**.
2. Select the VCN where your database services to be backed up are located.
3. On the resulting Virtual Cloud Network Details page, under **Resources**,click **Service Gateways**.
4. Click **Create Service Gateway** and provide the following details.
   a. **Name**: A descriptive name for the service gateway. It doesn't have to be unique. Avoid entering confidential information.
   b. **Compartment**: The compartment where you want to create the service gateway, if different from the compartment you're currently working in.
   c. **Services**: Select the service CIDR Label, `All <region> Services in Oracle Services Network` from the drop-down list.
   d. **Tags:** (advanced option) If you have permissions to create a resource, then you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.
5. Click **Create Service Gateway**.

Wait for the gateway to be created before proceeding to the next step.

6. Under **Resources**, click **Route Tables**.

   Route Table Association: You can associate a specific VCN route table with this gateway. If you associate a route table, afterward the gateway must always have a route table associated with it. You can modify the rules in the current route table or replace them with another route table.

7. Click the **Route Table** name that is being used by the subnet for Recovery Service.

8. In the resulting Route Table Details page, click **Add Route Rules** in the **Route Rules** section.

   When you configure a service gateway for a particular service CIDR label, you must also create a route rule that specifies that label as the destination and the target as the service gateway. You do this for each subnet that needs to access the gateway.

9. In the resulting Add Route Rules dialog, enter the following details:

   a. **Target Type**: Service Gateway.

   b. **Destination Service**: The service CIDR label that is enabled for the gateway. `All <region> Services in Oracle Services Network`

   c. **Target Service Gateway**: Select the name that you provided in step 4.

   d. **Description**: An optional description of the rule.

10. Click **Add Route Rules**.

Related Topics

- Resource Tags

Storage Configuration

The storage space inside the Exadata storage servers is configured for use by Oracle Automatic Storage Management (ASM) When you launch an Exadata Cloud Infrastructure instance.

By default, the following ASM disk groups are created:

- The DATA disk group is intended for the storage of Oracle Database data files.
- The RECO disk group is primarily used for storing the Fast Recovery Area (FRA), which is an area of storage where Oracle Database can create and manage various files related to backup and recovery, such as RMAN backups and archived redo log files.
- The `/acfs` file systems contain system files that support various operations. You should not store custom files, Oracle Database data files, or backups inside the ACFS disk groups. Custom ACFS mounts can be created using the DATA ASM disk group for files that are not service-related.

The disk group names contain a short identifier string that is associated with your Exadata Database machine environment. For example, the identifier could be C2, in which case the DATA disk group would be named DATAC2, the RECO disk group would be named RECOC2, and so on.

In addition, you can create a SPARSE disk group. A SPARSE disk group is required to support Exadata snapshots. Exadata snapshots enable space-efficient clones of Oracle databases that can be created and destroyed very quickly and easily. Snapshot clones are often used for development, testing, or other purposes that require a transient database.

Note that you cannot change the disk group layout after service creation.

- Impact of Configuration Settings on Storage
  Your choices concerning database Backups or sparse disk groups profoundly affect how storage space in the Exadata storage servers is allocated to the ASM and sparse disk groups

Impact of Configuration Settings on Storage

Your choices concerning database Backups or sparse disk groups profoundly affect how storage space in the Exadata storage servers is allocated to the ASM and sparse disk groups

If you choose to perform database backups to the Exadata storage, or to create a sparse disk group, or to do both, the storage space allocation in the Exadata storage servers will be affected.

The table that follows shows the approximate percentages of storage allocated for DATA, RECO, and SPARSE disk groups for each possible configuration.

Configuration Settings	DATA Disk Group	RECO Disk Group	SPARSE Disk Group
Database backups on Exadata storage: No **Sparse disk group: No**	80 %	20 %	0 %
Database backups on Exadata storage: Yes **Sparse disk group: No**	40 %	60 %	0 %
Database backups on Exadata storage: No **Sparse disk group: Yes**	60 %	20 %	20 %
Database backups on Exadata storage: Yes **Sparse disk group: Yes**	35 %	50 %	15 %

4

Getting Started with Exadata Cloud Infrastructure Deployment

After completing the preparation tasks in Preparing for Exadata Cloud Infrastructure, get started with deploying your Exadata Cloud Infrastructure system following these procedures.

- Tagging Oracle Exadata Database Service on Dedicated Infrastructure Resources
  Tagging is a powerful foundational service for Oracle Cloud Infrastructure (OCI) that enables users to search, control access, and do bulk actions on a set of resources based on the tag.
- Overview of X8M and X9M Scalable Exadata Infrastructure
  Oracle Cloud Infrastructure scalable X8M and X9M Exadata cloud infrastructure model allows you to add additional database and storage servers after provisioning and create a system that matches your capacity needs.
- Creating an Exadata Cloud Infrastructure Instance
  This topic explains how to create an Oracle Exadata Cloud Infrastructure instance. It also describes how to configure required access to the Oracle Cloud Infrastructure Object Storage service and set up DNS.
- Cloud Infrastructure Maintenance Updates
  Oracle performs the updates to all of the Oracle-managed infrastructure components on Exadata Cloud Infrastructure.
- Connecting to an Exadata Cloud Infrastructure Instance
  This topic explains how to connect to an Exadata Cloud Infrastructure instance using SSH or SQL Developer.
- Best Practices for Exadata Cloud Infrastructure Instances
  Oracle recommends that you follow these best practice guidelines to ensure the manageability of your Exadata Cloud Infrastructure instance:
- Moving to Oracle Cloud Using Zero Downtime Migration

Tagging Oracle Exadata Database Service on Dedicated Infrastructure Resources

Tagging is a powerful foundational service for Oracle Cloud Infrastructure (OCI) that enables users to search, control access, and do bulk actions on a set of resources based on the tag.

Importance of Tagging

Using the Oracle Cloud Infrastructure (OCI) tagging system, you can tag resources per your organizational scheme allowing you to group resources, manage costs, and give insights into usage. Tags also help you build a governance model around security and Maximum Availability Architecture (MAA). As your organization expands its cloud footprint, it can become challenging to keep track of the deployment architectures, security best practices, MAA, application tier, and so on. Using metadata tags to identify workload attributes can help keep up with the security and availability of your tenancy without cost overruns.

To enable customers to manage OCI resources securely and cost-effectively, Oracle provides a set of pre-defined tags in line with best practices for tagging resources. These tags are grouped into two namespaces, the `oracleStandard` namespace, and the `OracleApplicationName` namespace. You can think of a tag namespace as a container for your tag keys.

Consider a scenario where your organization has multiple cloud resources such as Exadata Infrastructure, VM Cluster, DB Home, Oracle Database and VM Cluster Networks across multiple compartments in your tenancy. Suppose you wish to track these cloud resources for specific purposes, report on them, or take bulk actions. In that case, you will need a system that lets you group these resources based on different criteria such as environment, criticality, target users, application, etc. You can achieve this by applying appropriate tags to these resources.

For example, you may tag all resources in your development stack with `Oracle-Standard.Environment=Dev` or for a business-critical application stack set `Oracle-Standard.Criticality=High` or `Extreme`. In the event of service disruptions due to various reasons, you would then be able to quickly identify all OCI resources associated with an application or business function or be able to separate critical and non-critical workloads.

Tagging can also help you deploy optimized configurations based on workload attributes identified via tags. For example, database deployments for the PeopleSoft application require a specific configuration. Setting the `ApplicationName` and `AppMajorVersion` tags while deploying an Oracle Database can ensure that the database is configured and ready for the particular application, for example, PeopleSoft out of the box.

Moreover, integration with the Cloud Advisor OCI service can provide you with direct, deep insight into how well your cloud services adhere to the corporate guidelines and help your management govern with a vision. See *Cloud Advisor Overview* for more details.

Adding Tags

You can tag resources using the Oracle Cloud Infrastructure (OCI) console, command-line interface, or SDK.

There are many cloud resources that can be tagged in an Oracle Exadata Database Service on Dedicated Infrastructure deployment. Exadata Infrastructure, VM Cluster, DB Home, Oracle Database, Autonomous Exadata VM Cluster, Autonomous Container Database, Autonomous Database, and VM Cluster Networks are some of them. Tags can either be applied while creating the resources or modified later. For example, you can apply tags to an Autonomous Container Database (ACD) while provisioning the ACD or add them later from its **Details** page.

See *How Tagging Works* for more details on using tags. Tagging integrates with Oracle Cloud Infrastructure authorization system. You can use IAM policy controls to enable delegation or restriction of tag manipulation. See *Authentication and Authorization* to learn about the permissions required to work with defined and free-form tags. (Required) Enter introductory text here, including the definition and purpose of the concept.

> **Tip:**
>
> For a "try it out" tutorial that demonstrates implementing tags in Oracle Autonomous Database, refer to *Lab 14: Oracle Standard Tags* in *Oracle Autonomous Database Dedicated for Fleet Administrators Workshop* on Oracle LiveLabs.

Your tenancies come with a library of standard tags that would apply to most resources. These tags are currently available as a set of Tag Namespaces that your governance administrators

can deploy. OCI best practices recommend applying these tags to all resources a standard tag can be applied to. Besides reporting and governance, OCI service automation can deliver workload-specific optimizations based on standard tag values.

For example, database deployments for the PeopleSoft application require a specific configuration. By setting the appropriate application tag key in the `Oracle-ApplicationName` tag namespace while deploying an Autonomous Database, can ensure that the database is configured ready for the particular application, for example, PeopleSoft out of the box.

Figure 4-1 Tagging Example

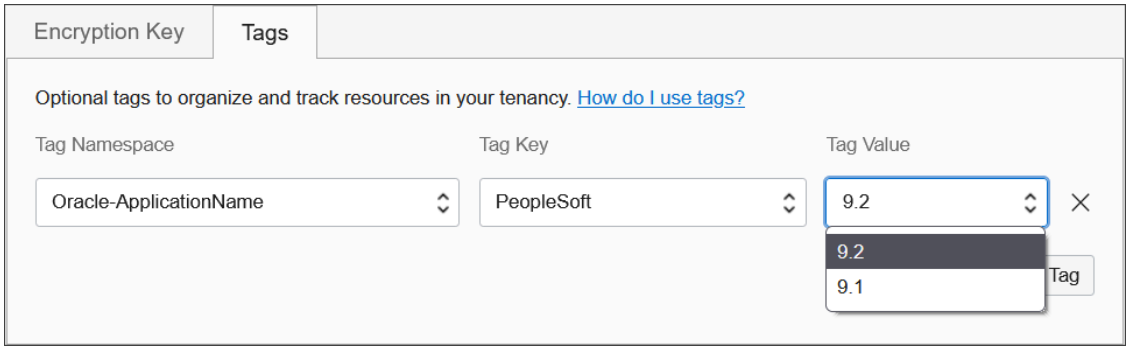

Oracle Standard Tags

Your tenancy governance administrators can deploy the standard tags at the tenancy level and may also mark certain tags as required, thereby enforcing tags on resources in those compartments. The following are the standard tags defined in the namespace called `OracleStandard`. For more information about importing standard tags, see *To import standard tags* under the *Managing Tag Namespaces* section.

Table 4-1 Oracle Standard Tags

Tag Key	Tag Value Options	Description
`OracleStandard.Criticality`	• Extreme • High • Medium • Low	Enables tiering of resources in line with corporate application classification standards. Customer governance can use this tag for reporting and ensuring resources are configured as per the guideline for the tier they belong to. For example, a database resource with `OracleStandard.Criticality` set to Extreme or High may require the highest availability SLA and may need to be configured with Autonomous Data Guard.

Table 4-1 (Cont.) Oracle Standard Tags

Tag Key	Tag Value Options	Description
`OracleStandard.Environment`	• Dev • Test • Prod • Pre-Prod • Staging • Trial • Sandbox • User Testing	Denotes a resource lifecycle. In the case of databases, it helps determine consolidation density, database distribution across containers, set maintenance plans, and manage clones.
`OracleStandard.Sensitivity`	• Public • Internal • Sensitive • Highly Sensitive • Extremely Sensitive	An application or database classification tag. `OracleStandard.Sensitivity` set to Highly Sensitive may indicate that an access control list or certain Network Security Group (NSG) enforcement is mandatory to restrict access.
`OracleStandard.Regulation`	Refer to *List of Compliance Regulations* for values.	Denotes one or more compliance regulations that a resource must adhere to. Tag administrators may add values to the list from the OCI Governance and Administration console. Refer to *Using Predefined Values* for more details.
`OracleStandard.TargetUsers`	• Public • Customers • Partners • Company • Division • Department • Workgroup	Denotes the end users of a resource. Another form of resource classification that helps determine target users and allows governance teams to set corporate standards based on user or application type.
`OracleStandard.EndUserCount`	• 1 • 10 • 100 • 1000 • 10000 • 100000 • 1000000 • 1000000 • 10000000	An approximate count of end-users. This tag helps determine the number of users impacted or the blast radius during an availability or security event. This also helps prioritize recovery efforts in the event of major outages affecting a large number of cloud resources.
`OracleStandard.OwnerEmail`	Free form tag. For example *john.smith@acme.com* or *app_support_grp@acme.com*	Denotes the email address of the resource owner.

Table 4-1 (Cont.) Oracle Standard Tags

Tag Key	Tag Value Options	Description
`OracleStandard.Org`	• HR • Finance • Marketing • Sales • Legal • R&D • Customer Suppport • Internal Support • Manufacturing	Identifies the customer's line of business or department that owns or uses the resource. This may help with cost aggregation reports and determining usage across business units.Tag administrators may add relevant values to the list from the OCI Governance and Administration console. Refer to *Using Predefined Values* for more details.
`OracleStandard.CostCenter`	• 12345 • WebMarketing	Freeform field for cost center.
`OracleStandard.RecoveryTimeObjectiveMinutes`	0-10080	Time in minutes. Denotes the maximum time within which the resource is required to recover from a failure.
`OracleStandard.RecoveryPointObjectiveMinutes`	0-1440	Time in minutes. Maximum data loss tolerance for a data store resource such as a database or a storage device.

List of Compliance Regulations

Table 4-2 List of Compliance Regulations

Regulation	Description
PCI DSS	Payment Card Industry Data Security Standard
HIPAA	Health Insurance Portability and Accountability Act
ISO	International Standards Organization
SOC1	System and Organization Controls 1
SOC 2	System and Organization Controls 2
FedRamp	Federal Risk and Authorization Management Program
GLBA	Gramm–Leach–Bliley Act
CCPA	California Consumer Privacy Act
SOX	Sarbanes Oxley
NIST	National Institute of Standards and Technology - Cyber Security
FISMA	Federal Information Security Management
HITECH	Health Information Technology for Economic and Clinical Health Act
FERPA	Family Educational Rights and Privacy Act (Student privacy)
FACTA	Fair and Accurate Credit Transaction Act

Table 4-2 (Cont.) List of Compliance Regulations

Regulation	Description
Texas HB300	Texas Medical Records Privacy Act
CIS	Center for Internet Security
CJIS	Criminal Justice Information Services Security Policy
C-TPAT	Customs-Trade Partnership Against Terrorism
COPPA	Children's Online Privacy Protection Act
PIPED Act, or PIPEDA	Personal Information Protection and Electronic Documents Act
GDPR	General Data Protection Regulation
PIPL	Personal Information Protection Law

Oracle Application Name Tags

Table 4-3 Oracle Application Name Tags

Tag Key	Tag Value Options	Description
Hyperion	• 11.2 • 11.1	Denotes the version of the Hyperion application.
JD Edwards	• 9.2 • 9.1 • 9.0	Denotes the version of the JD Edwards application.
Oracle_E-Business_Suite	• 12.2 • 12.1 • 12.1 • 11i	Denotes the version of the Oracle E-Business Suite application.
PeopleSoft	• 9.2 • 9.1	Denotes the version of the PeopleSoft application.
Siebel	• 8.2 • 8.1	Denotes the version of the Siebel application.
Other_Oracle_Application	Free form tag in string format.	Can be used to denote any application other than those listed above. You can enter the application name as a string value.

Related Topics

- To Import standard tags
- Cloud Advisor Overview
- Oracle Autonomous Database Dedicated for Fleet Administrators Workshop
- How Tagging Works
- Authentication and Authorization
- Managing Tag Namespaces

- Using Predefined Values

Cloud Infrastructure Maintenance Updates

Oracle performs the updates to all of the Oracle-managed infrastructure components on Exadata Cloud Infrastructure.

You may manage contacts who are notified regarding infrastructure maintenance, set a maintenance window to determine the time your quarterly infrastructure maintenance will begin, and also view scheduled maintenance runs and the maintenance history of your Exadata Cloud Infrastructure in the Oracle Cloud Infrastructure Console. For details regarding the infrastructure maintenance process and configuring the maintenance controls refer to the following:

- About Oracle-managed Exadata Cloud Infrastructure Maintenance
  Oracle performs patches and updates to all of the Oracle-managed system components on Exadata Cloud Infrastructure.
- Overview of the Quarterly Infrastructure Maintenance Process
  By default, infrastructure maintenance updates the Exadata database server hosts in a rolling fashion, followed by updating the storage servers.
- Overview of Monthly Security Maintenance
  Security maintenance, performed alongside the quarterly maintenance, is executed in months when important security updates are needed and includes fixes for vulnerabilities across all CVSS scores.
- Understanding Monthly and Quarterly Maintenance in the Same Month
- Using the Console to Configure Oracle-Managed Infrastructure Updates
  Software updates are scheduled quarterly and monthly. You can use the the console to schedule and plan for them.
- Monitor Infrastructure Maintenance Using Lifecycle State Information
  The lifecycle state of your Exadata Infrastructure resource enables you to monitor when the maintenance of your infrastructure resource begins and ends.
- Receive Notifications about Your Infrastructure Maintenance Updates
  There are two ways to receive notifications. One is through email to infrastructure maintenance contacts and the other one is to subscribe to the maintenance events and get notified.
- Managing Infrastructure Maintenance Contacts
  Learn to manage your Exadata infrastructure maintenance contacts.

About Oracle-managed Exadata Cloud Infrastructure Maintenance

Oracle performs patches and updates to all of the Oracle-managed system components on Exadata Cloud Infrastructure.

Oracle patches and updates include the physical database server hosts, Exadata Storage Servers, Network Fabric Switches, management switch, power distribution units (PDUs), integrated lights-out management (ILOM) interfaces, and Control Plane Servers. This is referred to as infrastructure maintenance.

The frequency of the updates depends on the region type, as follows:

- Commercial regions: Oracle performs full quarterly infrastructure updates and monthly security infrastructure updates.
- Government regions: Oracle performs monthly full infrastructure maintenance updates.

In all but rare exceptional circumstances, you receive advance communication about these updates to help you plan for them. If there are corresponding recommended updates for your VM cluster virtual machines (VMs), then Oracle provides notifications about them.

Wherever possible, scheduled updates are performed in a manner that preserves service availability throughout the update process. However, there can be some noticeable impact on performance and throughput while individual system components are unavailable during the update process.

For example, database server patching typically requires a reboot. In such cases, wherever possible, the database servers are restarted in a rolling manner, one at a time, to ensure that the service remains available throughout the process. However, each database server is unavailable for a short time while it restarts, and the overall service capacity diminishes accordingly. If your applications cannot tolerate the restarts, then take mitigating action as needed. For example, shut down an application while database server patching occurs.

> **Note:**
>
> Customers using Exadata Database on Dedicated Infrastructure in Oracle Cloud Infrastructure (OCI) US Government (OC2) and US Department of Defense (OC3) regions can use the OCI console to reschedule monthly and quarterly patching events.
>
> At this time specifying a maintenance schedule, all so known as “Setting Patch Management Schedule for Exadata Cloud Infrastructure”, is still not available in the OCI US Government (OC2) and US DOD (OC3) realms for Exadata patch management. For more information on Exadata Database on Dedicated Infrastructure on Patch Management Rescheduling can be found here.

Overview of the Quarterly Infrastructure Maintenance Process

By default, infrastructure maintenance updates the Exadata database server hosts in a rolling fashion, followed by updating the storage servers.

You can also choose non-rolling maintenance to update database and storage servers. The non-rolling maintenance method first updates your storage servers at the same time, then your database servers at the same time. Although non-rolling maintenance minimizes maintenance time, it incurs full system downtime while the storage servers and database servers are being updated.

Rolling infrastructure maintenance begins with the Exadata database server hosts. For the rolling maintenance method, database servers are updated one at a time. Each of the database server host's VMs is shut down, the host is updated, restarted, and then the VMs are started, while other database servers remain operational. This rolling maintenance impacts older applications not written to handle a rolling instance outage. This process continues until all servers are updated.

After database server maintenance is complete, storage server maintenance begins. For the rolling maintenance method, storage servers are updated one at a time and do not impact VM cluster VM's availability. However, the rolling storage server maintenance can result in reduced IO performance as storage servers are taken offline (reducing available IO capacity) and resynced when brought back online (small overhead on database servers). Properly sizing the database and storage infrastructure to accommodate increased work distributed to database and storage servers not under maintenance will minimize (or eliminate) any performance impact.

> **Note:**
>
> While databases are expected to be available during the rolling maintenance process, the automated maintenance verifies Oracle Clusterware is running but does not verify that all database services and pluggable databases (PDBs) are available after a server is brought back online. The availability of database services and PDBs after maintenance can depend on the application service definition. For example, a database service, configured with certain preferred and available nodes, may be relocated during the maintenance and wouldn't automatically be relocated back to its original node after the maintenance completes. Oracle recommends reviewing the documentation on *Achieving Continuous Availability for Your Applications on Exadata Cloud Systems* to reduce the potential for impact to your applications. By following the documentation's guidelines, the impact of infrastructure maintenance will be only minor service degradation as database servers are sequentially updated.

Oracle recommends that you follow the *Maximum Availability Architecture (MAA)* best practices and use Data Guard to ensure the highest availability for your critical applications. For databases with Data Guard enabled, Oracle recommends that you separate the maintenance windows for the infrastructure instances running the primary and standby databases. You may also perform a switchover prior to the maintenance operations for the infrastructure instance hosting the primary database. This allows you to avoid any impact on your primary database during infrastructure maintenance.

Prechecks are performed on the Exadata Cloud Infrastructure components prior to the start of the maintenance window. The goal of the prechecks is to identify issues that may prevent the infrastructure maintenance from succeeding. The Exadata infrastructure and all components remain online during the prechecks. An initial precheck is run approximately 5 days prior to the maintenance start and another precheck is run approximately 24 hours prior to maintenance start. If the prechecks identify an issue that requires rescheduling the maintenance notification is sent to the maintenance contacts.

> **Note:**
>
> Do not perform major maintenance operations on your databases or applications during the patching window, as these operations could be impacted by the infrastructure maintenance operations

- Time Estimates for Quarterly Maintenance Windows
  The time taken to update infrastructure components varies depending on the number of database servers and storage servers in the Exadata infrastructure, the maintenance method and whether custom action has been enabled.

Related Topics

- Achieving Continuous Availability For Your Applications
- Maximum Availability Architecture (MAA) Best Practices

Time Estimates for Quarterly Maintenance Windows

The time taken to update infrastructure components varies depending on the number of database servers and storage servers in the Exadata infrastructure, the maintenance method and whether custom action has been enabled.

The approximate times provided are estimates. Time for custom action, if configured, is not included in the estimates. Database server maintenance time may vary depending on the time required to shutdown each VM before the update and then start each VM and associated resources after the update of each node before proceeding to the next node. The storage server maintenance time will vary depending on the time required for the ASM rebalance, which is not included in the estimates below. If issues are encountered during maintenance this may also delay completion beyond the approximate time listed. In such a situation, if Oracle cloud operations determine resolution would extend beyond the expected window, they will send a notification and may reschedule the maintenance.

> **Note:**
>
> The timeframes mentioned below can change if Oracle cloud operations determine that additional maintenance work is needed. If additional time is necessary, Oracle will send a customer notification in advance to inform customers that additional time will be required for the next quarterly maintenance window.

Table 4-4 Approximate Times for Exadata Infrastructure Maintenance

Exadata Shape Configuration	Rolling Patching Method (Approximate Time)	Non-Rolling Patching Method (Approximate Time)
Quarter rack	5-6 hours	4-7 hours
Half rack	10 hours	4-7 hours
Full rack	20 hours	4-7 hours
Flexible shapes (X8M and higher)	1.5 hours per compute node + 1 hour per storage node	4-7 hours

Overview of Monthly Security Maintenance

Security maintenance, performed alongside the quarterly maintenance, is executed in months when important security updates are needed and includes fixes for vulnerabilities across all CVSS scores.

> **Note:**
>
> For more information about the CVE release matrix, see Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1).
> To view the CVE release matrix specific to an Exadata Infrastructure version, click the Exadata version, for example, Exadata 23. Version-specific CVE release matrices are listed in the Notes column of the table.

Security maintenance, when needed, is scheduled to be applied during a 21-day window that begins between the 18th-21st of each month and will run till the 9th-12th of the next month. Customers will receive notification of the proposed schedule at least 7 days before the start of the monthly maintenance window and can reschedule monthly maintenance to another date in the window if desired. The monthly security maintenance process updates database servers to fix critical security vulnerabilities and critical product issues. Monthly maintenance also updates storage servers to an Exadata Storage Software image that resolves known security vulnerabilities and product issues.

Updates to database servers are applied online via Ksplice technology, and have no impact to workloads running on the compute (database) servers, as database server security updates are applied online to the host server while your VM and all processes within the VM, including databases, remain up and running. Servers and VMs are not restarted. Updates to storage servers are applied in a rolling fashion. As with quarterly maintenance, the impact of rebooting storage servers should be minimal to applications.

While updating your services infrastructure, some operations including memory, and storage scaling, operating system and Grid Infrastructure patching (including prechecks), and elastic expansion of compute and storage servers may be blocked.

> **Note:**
>
> Only VM startup and shutdown operations are supported during monthly infrastructure maintenance.

Please plan to defer these operations until after the updates are complete. Application of security updates takes about 15 minutes per DB server host, plus 60 minutes per storage server depending on the I/O activity. If you attempt an affected operation, the console will notify you of the ongoing security updates. No software is updated in the guest VMs.

Related Topics

- https://support.oracle.com/rs?type=doc&id=888828.1

Understanding Monthly and Quarterly Maintenance in the Same Month

Special considerations are made when both quarterly and monthly security maintenance are scheduled to run in the same month. Quarterly maintenance will reapply any security fixes already applied by security maintenance, and neither quarterly nor monthly maintenance will apply a storage server update if the existing storage server version is the same or newer than the version contained in the update.

- The contents of the updates applied during quarterly maintenance are determined at the start of the maintenance quarter and use the latest Exadata release from the month prior to the start of the maintenance quarter. If any additional security fixes are available at that time, those updates are included in the quarterly maintenance. That image is then used throughout the quarter. For example, the January release is used for quarterly maintenance in Feb, March, and April.
- When quarterly maintenance is applied it is possible there are security updates previously installed on the database servers are not included in the quarterly maintenance release to be applied. In that case, the automation will apply the same security fixes to new release installed by the quarterly maintenance so there will not be any regression in security fixes. If the current image on the storage server is the same or newer than that to be applied by the quarterly or monthly security maintenance, that maintenance will be skipped for the storage servers.

If quarterly maintenance is scheduled within 24 hours of the time the monthly is scheduled, the scheduled monthly maintenance will be skipped and the monthly update will instead be applied immediately following the quarterly maintenance.

- When scheduled at the same time, the monthly update is executed immediately following the completion of the quarterly maintenance.

- If monthly maintenance is scheduled to begin 0-24 hours ahead of the quarterly maintenance, then the monthly maintenance will not execute as scheduled, but instead, wait and be executed immediately following the quarterly maintenance. If the quarterly maintenance is subsequently rescheduled, then the monthly security maintenance will begin immediately. Oracle, therefore, recommends scheduling quarterly and monthly maintenance at the same time. As a result, if you reschedule the quarterly at the last moment, the monthly maintenance will run at the scheduled time instead of immediately upon editing the schedule. You can also reschedule the monthly security maintenance when rescheduling the quarterly maintenance as long as you keep the monthly within the current maintenance window. Monthly maintenance can be rescheduled to another time in the maintenance window, but cannot be skipped.

Monthly Security Maintenance before Quarterly Maintenance

- To apply security maintenance before quarterly maintenance, reschedule the monthly security maintenance to occur more than 24 hours prior to the quarterly maintenance. The security maintenance will online apply security patches to the database servers with no impact to applications, and apply an update to the storage servers with minimal to no impact (may be slight performance degradation) on applications. The quarterly maintenance will follow as scheduled, and will perform rolling maintenance on the database servers, which will impact applications not written to handle a rolling reboot. As part of the quarterly maintenance, it will apply the same security updates to the database server that are already installed on the system (no security regression).
- If you are concerned about getting the latest security updates applied, schedule the monthly security maintenance to run after the new monthly maintenance window opens (usually on the 21st of the month).
- The impact of the monthly security maintenance rebooting the storage servers should be minimal, so impact to the applications during this month will only be due to the restart of the database servers during the quarterly maintenance. However, if you must coordinate a maintenance window with your end users for the security maintenance, this will require two maintenance windows.

Quarterly Maintenance before Monthly Security Maintenance

- To run the quarterly maintenance before the monthly security maintenance, reschedule the security maintenance to run no earlier than 24 hours before the quarterly maintenance is scheduled to start. The security maintenance will be deferred until the quarterly maintenance is completed. The quarterly maintenance will perform rolling maintenance on the database servers, which will impact applications not written to handle a rolling reboot. The quarterly maintenance may or may not skip the storage server patching. That depends on if it is newer or older than the release currently installed. In most cases, the version installed should be newer than the version associated with the quarterly maintenance. Exceptions to this rule may occur if it is the first month of a maintenance quarter, or you skipped the security maintenance in one or more prior months. The security maintenance will run either immediately after the quarterly maintenance is completed, or when scheduled, whichever is later. It will apply online updates to the database servers (no application impact) and will likely update the storage servers in a rolling manner. In some corner cases. the quarterly maintenance may contain the same storage server release as the security maintenance and the security maintenance storage server updates will be skipped.
- The impact to end users of running the quarterly maintenance before the security maintenance should be roughly the same as running the security maintenance first. The quarterly maintenance will be a disruptive event, but the security maintenance rebooting the storage servers should cause minimal disruption, and the security maintenance is applied to the database servers online. However, if you must coordinate a maintenance

window with your end users for the security maintenance, this will require two maintenance windows. You can schedule those two maintenance windows to run back-to-back, to appear as single maintenance window to end users. To do this, reschedule the security maintenance to start at the same time (or up to 24 hours prior) as the quarterly maintenance. The security maintenance will be deferred until the quarterly maintenance is completed. Assuming you have been regularly applying monthly security maintenance, the storage servers will be skipped by the quarterly maintenance and will be updated by the security maintenance immediately upon the completion of the quarterly maintenance.

Minimizing Maintenance Windows

- To minimize the number of maintenance windows (you have to negotiate those with end users), schedule the quarterly maintenance and monthly maintenance at the same time. The security maintenance will be blocked. The quarterly maintenance will update the database servers in a rolling manner and will most likely skip the storage server. The security maintenance will follow up immediately and update the database servers online and the storage servers in a rolling manner. The result is a single database and storage server restart in a single maintenance window.
- There are two exceptions to this. 1. If the quarterly and monthly maintenance contain the same storage server release, the quarterly maintenance will apply the storage server update, and the security maintenance will be skipped. From your perspective, this is still a single rolling reboot in a single maintenance window. 2. The currently installed release on the storage servers is older than that contained in the quarterly maintenance, which in turn is older than that in the security maintenance. That would cause the quarterly maintenance to update the storage, and then the security maintenance to do it as well. This can only happen if you skipped a prior month's security maintenance, because it requires the current image to be at least 2 months out of date. In such a scenario, you may want to schedule the security maintenance first and then the quarterly maintenance. This would result in one storage server reboot, but two distinct maintenance windows — the first for the security maintenance, and then later the quarterly maintenance.
- To minimize the impact to your end users, always apply the monthly security updates, and in months where both are scheduled, schedule them at the same time.

> **Note:**
>
> If the Exadata Infrastructure is provisioned before Oracle schedules the security maintenance, then it will be eligible for security maintenance.
> Any time before the scheduled monthly Exadata Infrastructure maintenance, you can reschedule it.

Using the Console to Configure Oracle-Managed Infrastructure Updates

Software updates are scheduled quarterly and monthly. You can use the the console to schedule and plan for them.

Full Exadata Cloud Infrastructure software updates are scheduled on a quarterly basis for commercial regions, and monthly for government regions. In addition, important security updates are scheduled monthly. While you cannot opt-out of these infrastructure updates, Oracle alerts you in advance through the Cloud Notification Portal and allows scheduling flexibility to help you plan for them.

For quarterly infrastructure maintenance, you can set a maintenance window to determine when the maintenance will begin. You can also edit the maintenance method, enable custom

action, view the scheduled maintenance runs and the maintenance history, and manage maintenance contacts in the in the Exadata Infrastructure Details page of the Oracle Cloud Infrastructure Console.

- View or Edit Quarterly Infrastructure Maintenance Preferences for Exadata Cloud Infrastructure
  To edit your Oracle Exadata Database Service on Dedicated Infrastructure infrastructure maintenance preferences, be prepared to provide values for the infrastructure configuration. The changes you make will only apply to future maintenance runs, not those already scheduled.
- To set the automatic monthly maintenance schedule for Exadata Cloud Infrastructure
- To view or edit the properties of the next scheduled quarterly maintenance for Exadata Cloud Infrastructure
  Review and change the properties of the Exadata Cloud Infrastructure scheduled quarterly maintenance.
- To view the maintenance history of an Exadata Cloud Infrastructure resource
  This task describes how to view the maintenance history for a cloud Exadata infrastructure or DB system. resource.
- View and Edit Quarterly Maintenance While Maintenance is In Progress or Waiting for Custom Action
  While maintenance is in progress, you can enable or disable custom action and change the custom action timeout. While maintenance is waiting for custom action, you can resume the maintenance prior to the timeout or extend the timeout.

View or Edit Quarterly Infrastructure Maintenance Preferences for Exadata Cloud Infrastructure

To edit your Oracle Exadata Database Service on Dedicated Infrastructure infrastructure maintenance preferences, be prepared to provide values for the infrastructure configuration. The changes you make will only apply to future maintenance runs, not those already scheduled.

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.

   > **Note:**
   >
   > Specifying maintenance preferences is not available in Government regions.

2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to edit.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Edit Maintenance Preferences**.
   Edit Maintenance Preferences page is displayed.

> **Note:**
>
> Changes made to maintenance preferences apply only to future maintenance, not the maintenance that has already been scheduled. To modify scheduled maintenance, see *View or Edit a Scheduled Maintenance for Exadata Cloud Infrastructure*.

6. On the **Edit Maintenance Preferences** page, configure the following:
   - **Choose a maintenance method:**
     - **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
     - **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.
   - **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.
     - **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.

       Default: 30 minutes

       Maximum: 120 minutes
   - **Maintenance schedule:**
     - **No preference:** The system assigns a date and start time for infrastructure maintenance.
     - **Specify a schedule:** Choose your preferred month, week, weekday, start time, and lead time for infrastructure maintenance.

       > **Note:**
       >
       > Specifying a maintenance schedule is not available in Government regions.

       * Under **Maintenance months**, specify at least one month for each quarter during which Exadata infrastructure maintenance will take place. You can select more than one month per quarter. If you specify a long lead time for advanced notification (for example, 4 weeks), you may wish to specify 2 or 3 months per quarter during which maintenance runs can occur. This will ensure that your maintenance updates are applied in a timely manner after accounting for your required lead time. Lead time is discussed in the following steps.
       * Optional. **Under Week of the month**, specify which week of the month, maintenance will take place. Weeks start on the 1st, 8th, 15th, and 22nd days

of the month, and have a duration of 7 days. Weeks start and end based on calendar dates, not days of the week. Maintenance cannot be scheduled for the fifth week of months that contain more than 28 days. If you do not specify a week of the month, Oracle will run the maintenance update in a week to minimize disruption.

* Optional. **Under Day of the week**, specify the day of the week on which the maintenance will occur. If you do not specify a day of the week, Oracle will run the maintenance update on a weekend day to minimize disruption.
* Optional. **Under Start hour**, specify the hour during which the maintenance run will begin. If you do not specify a start hour, Oracle will pick the least disruptive time to run the maintenance update.
* Under **Lead Time**, specify the minimum number of weeks ahead of the maintenance event you would like to receive a notification message. Your lead time ensures that a newly released maintenance update is scheduled to account for your required minimum period of advanced notification.

7. Click **Save Changes**.
   If you switch from rolling to non-rolling maintenance method, then **Confirm Non-rolling Maintenance Method** dialog is displayed.
   a. Enter the name of the infrastructure in the field provided to confirm the changes.
   b. Click **Save Changes**.

To set the automatic monthly maintenance schedule for Exadata Cloud Infrastructure

This task describes how to set the maintenance schedule for a cloud Exadata infrastructure resource.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Navigate to the Cloud Exadata infrastructure you want to access:

   In the **Oracle Exadata Database Service on Dedicated Infrastructure** section, click **Exadata Infrastructure**. In the list of infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.

   On the resource details page, under **Maintenance**, click the **View** link in the **Next Security Maintenance** field. The Next Security Maintenance will only show a scheduled time if a Monthly maintenance is scheduled. Customers will get a minimum of 7 days advance notice prior to monthly maintenance, and will then be able to reschedule during the 3 week window.
3. In the **Cloud Exadata Infrastructure** pane, you will see the **Scheduled Start Time**
4. In the **Scheduled Start Time** field, click **Edit**.
5. In the **Edit Maintenance Start Time** dialog, enter a new date and time in the **Scheduled start time** field. The Maintenance run can be rescheduled to start within: 3 Weeks, 15th to 6th of next month, any blocked dates/weekends, any other validations will block the UI/API.
6. **Save Changes**.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

To view or edit the properties of the next scheduled quarterly maintenance for Exadata Cloud Infrastructure

Review and change the properties of the Exadata Cloud Infrastructure scheduled quarterly maintenance.

NOT_SUPPORTED

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**."
2. Navigate to the Cloud Exadata infrastructure or DB system you want to access:
   Cloud Exadata infrastructure (new resource model): Under **Exadata at Oracle Cloud**, click **Exadata Infrastructure**. In the list of infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.

   DB systems: Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
3. On the resource details page, under **Maintenance**, click the **View** link in the **Next Quarterly Maintenance** field.
   The Exadata Infrastructure Maintenance page is displayed.
4. On the **Exadata Infrastructure Maintenance** page, scheduled maintenance details are listed.
   Target DB Server Version and **Target Storage Server Version**: These fields display the Exadata software version to be applied by the scheduled maintenance. The version applied will be the most recent certified update for Exadata infrastructures in the cloud. If the next quarterly update is not yet certified when the maintenance is scheduled, then the versions may show "LATEST" until the new quarterly update becomes available. Once the update becomes available the new version will be displayed.

   To find information on the Database Server Exadata software version or the Storage Server Exadata software version, see My Oracle Support note *Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1)*.

   For each scheduled Exadata Infrastructure resource maintenance event, the Maintenance page lists the following details:

   - The status of the event
   - The OCID of the event
   - The scheduled start time and date of the event
   - Click **Patch Now** to start the maintenance event immediately. When prompted, click **Run Maintenance** to confirm that you want to start the event now.

   If a maintenance event is already in progress on one or more of the VM Clusters hosted by an Exadata Infrastructure resource when a maintenance event on that resource is to start, the Exadata Infrastructure resource maintenance event is queued and begins immediately after all VM Cluster maintenance events complete.
5. To change the next scheduled maintenance settings, click **Edit Maintenance Run**.
6. On the **Edit Maintenance** page, do the following:
   - Select a maintenance method, **Rolling** or **Non-rolling**.

> **Note:**
>
> If you select the **Non-rolling** option, components will be updated simultaneously, resulting in full system downtime.

- **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.
  - **Custom action timeout (in minutes):** Maximum timeout available to perform custom action before starting maintenance on the DB Servers.
    Default: 30 minutes

    Minimum: 15 minutes

    Maximum: 120 minutes
- To reschedule the next maintenance run, enter a date and time in the **Scheduled Start time** field.
  The following restrictions apply:
  - You can reschedule the infrastructure maintenance to a date no more than 180 days from the prior infrastructure maintenance. If a new maintenance release is announced prior to your rescheduled maintenance run, the newer release will be applied on your specified date. You can reschedule your maintenance to take place earlier than it is currently scheduled. You cannot reschedule the maintenance if the current time is within 2 hours of the scheduled maintenance start time.
  - Oracle reserves certain dates each quarter for internal maintenance operations, and you cannot schedule your maintenance on these dates.
- Click **Save Changes**.

7. To view estimated maintenance time details for various components, click the **View** link is displayed in the Total Estimated Maintenance Time field.
   The **View** link is displayed in the **Total Estimated Maintenance Time** field only if the Maintenance Method is **Rolling**.

   The **Estimated Maintenance Time Details** page is displayed with details that include:

   - Total Estimated Maintenance Time
   - Database Servers Estimated Maintenance Time
   - Storage Servers Estimated Maintenance Time
   - Order in which components are updated. In rolling maintenance, components are updated in the sequence displayed

   a. To view the number of VMs that will be restarted as part of Database Server maintenance, click the **Show details** link.
      The **VM Location** dialog is displayed.

   b. In the **VM Cluster Name** field, you can find out what VM cluster a particular VM belongs to.

c. Click **Close**.

8. Click **Close** to close the **Estimated Maintenance Time Details** page.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.
- Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1)

To view the maintenance history of an Exadata Cloud Infrastructure resource

This task describes how to view the maintenance history for a cloud Exadata infrastructure or DB system. resource.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Navigate to the Cloud Exadata infrastructure or DB system you want to access:

   Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata Infrastructure**. In the list of infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.
3. On the resource details page, under **Bare Metal,VM, and Exadata**, click the **Maintenance History**.
4. The Maintenance jobs, State and type of patching are displayed.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.
- My Oracle Support note Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1)

View and Edit Quarterly Maintenance While Maintenance is In Progress or Waiting for Custom Action

While maintenance is in progress, you can enable or disable custom action and change the custom action timeout. While maintenance is waiting for custom action, you can resume the maintenance prior to the timeout or extend the timeout.

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to edit.

   The **Infrastructure Details** page displays information about the selected Oracle Exadata infrastructure.

> **Note:**
>
> **Maintenance In Progress** status is displayed in the **Next Maintenance** field.

View and Edit Maintenance While Maintenance is In Progress

1. Click the **View** link in the **Next Maintenance** field.

   The **Exadata Infrastructure Maintenance** page is displayed.

2. Click **Edit Maintenance Run**.

   Edit Maintenance page is displayed.

> **Note:**
>
> You can only make edits to the custom action configuration, not the maintenance method or scheduled start time. Enabling or disabling the custom action or modifying the custom action timeout while maintenance is in progress will apply to all database servers that have yet to be updated.

3. On the **Edit Maintenance** page, do the following:

   - **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.
     - **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.
       Default: 30 minutes

       Minimum: 15 minutes

       Maximum: 120 minutes

4. Click **Save Changes**.

   If you have configured the rolling maintenance method, then the **View** link is displayed in the **Total Estimated Maintenance Time** field.

   a. Click **View**.
   Estimated Maintenance Time Details page is displayed with details that include:

      - Total Estimated Maintenance Time
      - Database Servers Estimated Maintenance Time
      - Storage Servers Estimated Maintenance Time
      - Network Switches Estimated Maintenance Time
      - Order in which components are updated. In rolling maintenance, components are updated in the sequence displayed.

b. Click **Close**.

View and Edit Maintenance While Maintenance is Waiting for Custom Action

1. Click the **View** link in the **Next Maintenance** field.

   Exadata Infrastructure Maintenance page is displayed.

   > **Note:**
   >
   > - Editing a maintenance run is not available while waiting for custom action.
   > - While maintenance is waiting for custom action, an information block is displayed. The information block is removed after the maintenance resumes.

2. On the information block, do one of the following:

   a. Click **Resume Maintenance Now** to resume the maintenance, proceeding to the next database server.
   Resume Maintenance dialog is displayed. Click **Resume Maintenance Now**.

   b. Click **Extend Custom Action Timeout**.
   You can extend timeout multiple times within the maximum allowable time of 2 hours. If you try extending beyond the maximum limit, then the system displays the Cannot Extend Custom Action Timeout dialog indicating that the custom action timeout has already been extended to the maximum allowable 2 hours and you cannot extend it further.

Monitor Infrastructure Maintenance Using Lifecycle State Information

The lifecycle state of your Exadata Infrastructure resource enables you to monitor when the maintenance of your infrastructure resource begins and ends.

In the Oracle Cloud Infrastructure Console, you can see lifecycle state details messages on the **Exadata Infrastructure Details** page when a tooltip is displayed beside the **Status** field. You can also access these messages using the `ListCloudExadataInfrastructures` API, and using tools based on the API, including *SDKs* and the *OCI CLI*.

During infrastructure maintenance operations, you can expect the following:

- If you specify a maintenance window, then patching begins at your specified start time. The infrastructure resource's lifecycle state changes from **Available** to **Maintenance in Progress**.

  > **Note:**
  >
  > The prechecks are now done prior to the start of the maintenance.

- When Exadata database server maintenance starts, the infrastructure resource's lifecycle state is **Maintenance in Progress**, and the associated lifecycle state message is, **The underlying infrastructure of this system (dbnodes) is being updated.**
- When storage server maintenance starts, the infrastructure resource's lifecycle state is **Maintenance in Progress**, and the associated lifecycle state message is, **The underlying infrastructure of this system (cell storage) is being updated and this will not impact Database availability.**

- After storage server maintenance is complete, the networking switches are updated one at a time, in a rolling fashion.
- When maintenance is complete, the infrastructure resource's lifecycle state is **Available**, and the Console and API-based tools do not provide a lifecycle state message.

Related Topics

- ListCloudExadataInfrastructures
- Software Development Kits and Command Line Interface
- Command Line Interface (CLI)

Receive Notifications about Your Infrastructure Maintenance Updates

There are two ways to receive notifications. One is through email to infrastructure maintenance contacts and the other one is to subscribe to the maintenance events and get notified.

Oracle schedules maintenance run of your infrastructure based on your scheduling preferences and sends email notifications to all your infrastructure maintenance contacts. You can login to the console and view details of the schedule maintenance run. Appropriate maintenance related events will be generated as Oracle prepares for your scheduled maintenance run, for example, schedule reminder, patching started, patching end, and so on. For more information about all maintenance related events, see *Oracle Cloud Exadata Infrastructure Events*. In case, if there are any failures, then Oracle reschedules your maintenance run, generates related notification, and notifies your infrastructure maintenance contacts.

For more information about Oracle Cloud Infrastructure Events, see *Overview of Events*. To receive additional notifications other than the ones sent to infrastructure maintenance contacts, you can subscribe to infrastructure maintenance events and get notified using the Oracle Notification service, see *Notifications Overview*.

Related Topics

- Oracle Exadata Database Service on Dedicated Infrastructure Events
  Exadata Cloud Infrastructure resources emit events, which are structured messages that indicate changes in resources.
- Overview of Events
- Notifications Overview
- Managing Infrastructure Maintenance Contacts
  Learn to manage your Exadata infrastructure maintenance contacts.

Managing Infrastructure Maintenance Contacts

Learn to manage your Exadata infrastructure maintenance contacts.

- To manage maintenance contacts in an Exadata Cloud Infrastructure
  Manage contacts for Exadata infrastructure maintenance notifications using the Console.

To manage maintenance contacts in an Exadata Cloud Infrastructure

Manage contacts for Exadata infrastructure maintenance notifications using the Console.

To prevent an Exadata infrastructure administrator from being overwhelmed by system update notifications, you can specify up to 10 email addresses of people to whom maintenance notifications are sent.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. In the **Oracle Exadata Database Service on Dedicated Infrastructure** section, click **Exadata Infrastructure** to display a list of Exadata infrastructures in the default compartment. You can select a different compartment from the **Compartment** drop-down located in the **List Scope** section.
3. In the list of Exadata infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.
4. In the **Maintenance** section, click **Manage** in the **Customer Contacts** field to display the Manage Contacts dialog.
5. Click the **Add Contacts** button to display a field in which to enter a valid email address. You can have up to 10 maintenance contacts for each Exadata infrastructure.
6. To edit an email address, in the Manage Contacts dialog, select the box preceding the email address you want to edit and click the **Edit** button.
7. To remove an email address from the list, in the Manage Contacts dialog, select the box preceding the email address you want to remove and click the **Remove** button.

Overview of X8M and X9M Scalable Exadata Infrastructure

Oracle Cloud Infrastructure scalable X8M and X9M Exadata cloud infrastructure model allows you to add additional database and storage servers after provisioning and create a system that matches your capacity needs.

Adding capacity to the X8M or X9M Exadata cloud infrastructure has two workflows:

- One for infrastructures created after the release date for Create and Manage Multiple Virtual Machines per Exadata System (MultiVM) and VM Cluster Node Subsetting,
  See *Using the Console to Scale the Resources on a VM Cluster*
- Another workflow for infrastructures created before the release date for Create and Manage Multiple Virtual Machines per Exadata System (MultiVM) and VM Cluster Node Subsetting
  See *Using the Console to Scale Infrastructure Compute and Storage*
- The Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances are provisioned with an infrastructure model that uses two resources, the **cloud Exadata infrastructure** resource, and the **cloud VM cluster** resource.
- The Cloud Exadata Infrastructure Resource
  The infrastructure resource is the top-level (parent) resource.
- The Cloud VM Cluster Resource
  The VM cluster is a child resource of the infrastructure resource.
- Additional Exadata Cloud Infrastructure Instance Resources
  The new Exadata resource model retains the rest of the resource types found in DB systems: Oracle Databases, database backups, Data Guard Associations, Work Requests, Oracle Database Homes, and database server nodes (also called "virtual machines").
- The X8M and X9M Virtual Machine File System Structure Important File System and Sizes
- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

Related Topics

-
-

The Exadata Cloud Infrastructure Resource Model

Exadata Cloud Infrastructure instances are provisioned with an infrastructure model that uses two resources, the **cloud Exadata infrastructure** resource, and the **cloud VM cluster** resource.

Existing Exadata DB systems do not use this infrastructure model can be easily switched to the new resource model with no downtime. For instructions on switching, see Switch an Exadata DB system to the new Exadata resource model.

The Cloud Exadata Infrastructure Resource

The infrastructure resource is the top-level (parent) resource.

At the infrastructure level, you control the number of database and storage servers. You also control Exadata system maintenance scheduling at the Exadata infrastructure level. This resource is created using the CreateCloudExadataInfrastructure API.

See *Scaling Exadata X8M and X9M Compute and Storage* for information on scaling the X8M or X9M cloud Exadata infrastructure resource.

> **Note:**
>
> After adding storage or database servers to the infrastructure resource, you must then add them to the system VM clusters to utilize the new capacity.

Related Topics

- Scaling Exadata X8M and X9M Compute and Storage
  The flexible X8M and X9M system model is designed to be easily scaled in place, with no need to migrate the database using a backup or Data Guard

The Cloud VM Cluster Resource

The VM cluster is a child resource of the infrastructure resource.

The VM cluster resource provides a link between your Exadata cloud infrastructure resource and Oracle Database. Networking, OCPU count, IORM, and Oracle Grid Infrastructure are configured and managed at the VM cluster level. This resource is created using theCreateCloudVmCluster API.

See To add database server or storage server capacity to a cloud VM cluster for information on adding available storage or database servers to the VM cluster. Note that you must add servers to the infrastructure resource before you can add capacity to the VM cluster.

Multi-VM enabled Infrastructure support multiple VM clusters in a single infrastructure

Exadata Cloud Infrastructure instances that are NOT Multi-VM enabled Infrastructure only support creating a single cloud VM cluster.

Related Topics

- About IORM
  The I/O Resource Management (IORM) feature allows you to manage how multiple databases share the I/O resources of an Oracle Exadata cloud VM cluster for systems using the new resource model or DB system

Additional Exadata Cloud Infrastructure Instance Resources

The new Exadata resource model retains the rest of the resource types found in DB systems: Oracle Databases, database backups, Data Guard Associations, Work Requests, Oracle Database Homes, and database server nodes (also called "virtual machines").

> **Note:**
>
> The database server file system for database server nodes (also known as "virtual machines") has changed with the X8M generation of hardware. See The X8M and X9M Virtual Machine File System Structure Important File System and Sizes for details on the X8M database server node file system.

The X8M and X9M Virtual Machine File System Structure Important File System and Sizes

Table 4-5 The X8M and X9M Virtual Machine File System Structure Important File System and Default Sizes

Filesystem Mount	Size	Configuration
/	15 GB	Max supported size 900 GB. File system size can only be increased.
/u01	20 GB	Max supported size 900 GB. File system size can only be increased.
/u01/../grid	50 GB	File system is not resizable.
/u02	60 GB	Max supported size 900 GB. File system size can be increased or decreased.
/acfs01	100 GB	Contact Oracle Support to resize.
/boot	509 MB	File system is not resizable.
/crashfiles	20 GB	File system is not resizable.
/home	4 GB	Max supported size 900 GB. File system size can only be increased.
/var	5 GB	Max supported size 900 GB. File system size can only be increased.

Table 4-5 (Cont.) The X8M and X9M Virtual Machine File System Structure Important File System and Default Sizes

Filesystem Mount	Size	Configuration
`/var/log`	18 GB	Max supported size 900 GB. File system size can only be increased.
`/var/log/audit`	3 GB	Max supported size 900 GB. File system size can only be increased.
`/tmp`	3 GB	Max supported size 900 GB. File system size can only be increased.

The New Exadata Cloud Infrastructure Resource Model

Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

In the new model, there are two resources, the **cloud Exadata infrastructure** resource, and the **cloud VM cluster** resource.

The X8M and X9M system models are only compatible with the new resource model. For provisioning new X7 and X8 systems, Oracle recommends using the new resource model so that your instance will not have to be switched to the new resource model later.

> **Note:**
>
> No new systems can be provisioned with the old DB system resource model/APIs after May 15th, 2021. Support for the old DB system resource model/APIs on existing systems will end on January 15th, 2021. After this date, old APIs will stop working and the only action available will be to list DB System details and perform the switch to the new API. Oracle recommends that you migrate your Exadata Cloud Infrastructure instances to the new resource model APIs as soon as possible. Converting to the new resource model does not involve any system downtime.

Existing Exadata DB systems can be easily switched to the new resource model with no downtime. For instructions on switching, see Switch an Exadata DB System to the New Resource Model and APIs.

Creating an Exadata Cloud Infrastructure Instance

This topic explains how to create an Oracle Exadata Cloud Infrastructure instance. It also describes how to configure required access to the Oracle Cloud Infrastructure Object Storage service and set up DNS.

When you create an Exadata Cloud Infrastructure instance using the Console or the API, the system is provisioned to support Oracle databases. Along with the Infrastructure, a VM cluster, an initial Database Home and database are created. You can create additional Database Homes and databases at any time by using the Console or the Oracle Cloud Infrastructure

API. The service creates an initial database based on the options you provide and some default options described later in this topic.

- Resources to Be Created
- Prerequisites for Creating an Cloud Exadata Infrastructure Instance
  You need a SSH key pair key and a Virtual Cloud Network (VCN) to create an infrastructure instance.
- Default Options for the Initial Database
  Default option simplify launching an Exadata Cloud Infrastructure instance in the Console and when using the API.
- Using the Console to Create Infrastructure Resources
  Console tasks required to create cloud resources

Resources to Be Created

You will provision Exadata Cloud Infrastructure infrastructure and VM cluster resources separately.

- **Cloud Exadata infrastructure** resource: The infrastructure resource is the top-level (parent) resource. At the infrastructure level, you control the number of database and storage servers. You also control Exadata system maintenance scheduling at the Exadata infrastructure level.
- **Cloud VM cluster** resource: The VM cluster is a child resource of the infrastructure resource, providing a link between your Exadata cloud infrastructure resource and Oracle Database. Networking, OCPU count, IORM (see About IORM, and Oracle Grid Infrastructure are configured and managed at the VM cluster level. To create a cloud VM cluster, you must have an existing Cloud Exadata infrastructure resource to house the VM cluster.

> **Note:**
>
> - Exadata Cloud Infrastructure only supports using the new resource model (consisting of separate Exadata infrastructure and VM cluster resources) to provision Exadata Cloud Infrastructure instances, regardless of the hardware shape family you are choosing (X7, X8, X8M, or X9M). The DB system resource model and APIs are deprecated for Exadata Cloud Infrastructure.
> - Multi-VM enabled Infrastructure supports the creation of up to 8 VM clusters in an Infrastructure. >> Exadata Infrastructures with X8M and above generation of DB Servers can support a maximum of 8 VM clusters across all DB Servers. Maximum number of clusters across the infrastructure depends on the resources available per DB server and is subject to the per DB Server maximum VM limit. For more information, see Overview of VM Cluster Node Subsetting.
> - An Exadata Cloud Service Infrastructure instance that is NOT Multi-VM enabled supports only one cloud VM cluster

Prerequisites for Creating an Cloud Exadata Infrastructure Instance

You need a SSH key pair key and a Virtual Cloud Network (VCN) to create an infrastructure instance.

- The proper IAM policy is required to proceed. See Required IAM Policy for Exadata Cloud Infrastructure
- The public key, in OpenSSH format, from the key pair that you plan to use for connecting to the system via SSH. A sample public key, abbreviated for readability, is shown below.

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

  For more information, see *Managing Key Pairs on Linux Instances* .
- A correctly configured virtual cloud network (VCN) to launch the system in. Its related networking resources (gateways, route tables, security lists, DNS, and so on) must also be configured as necessary for the system. For more information, see *Network Setup for Exadata Cloud Infrastructure Instances* .

Related Topics

- Managing Key Pairs on Linux Instances
- Network Setup for Exadata Cloud Infrastructure Instances
  This topic describes the recommended configuration for the VCN and several related requirements for the Exadata Cloud Infrastructure instance.

Default Options for the Initial Database

Default option simplify launching an Exadata Cloud Infrastructure instance in the Console and when using the API.

The following default options are used for the initial database:

- **Console Enabled:** False
- **Create Container Database:** False for version 11.2.0.4 databases. Otherwise, true.
- **Create Instance Only (for standby and migration):** False
- **Database Home ID:** Creates a database home
- **Database Language:** AMERICAN
- **Database Sizing Template:** odb2
- **Database Storage:** Automatic Storage Management (ASM)
- **Database Territory:** AMERICA
- **Database Unique Name:** The user-specified database name and a system-generated suffix, for example, dbtst_phx1cs.
- **PDB Admin Name:** pdbuser (Not applicable for version 11.2.0.4 databases.)

Using the Console to Create Infrastructure Resources

Console tasks required to create cloud resources

- To create a Cloud Exadata infrastructure resource
- To create a cloud VM cluster resource
  Create a VM cluster in an Exadata Cloud Infrastructure instance.

- Configuring Network Resources for Recovery Service
  Review the network requirements and configurations required to backup your Oracle Cloud databases to Recovery Service.

To create a Cloud Exadata infrastructure resource

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata Infrastructure**.
3. Click **Create Exadata Cloud Infrastructure**.
4. **Compartment:** Select a compartment for the Exadata infrastructure.
5. **Display name:** Enter a display name for the Exadata infrastructure. The name doesn't need to be unique. An Oracle Cloud Identifier (OCID) will uniquely identify the Cloud Exadata infrastructure resource. Avoid entering confidential information.
6. **Select an availability domain:** The availability domain in which the Exadata infrastructure resides.
7. **Select the Exadata Cloud Infrastructure model:** Select either a fixed-shape system (quarter, half, or full rack X7-2 or X8-2 shapes), or a scalable system (X8M-2 or X9M-2).

   X9M-2: If you select the flexible X9M-2 cloud infrastructure model, your initial Exadata Cloud Infrastructure instance can have a minimum of 2 database servers and 3 storage servers up to 32 database servers and 64 storage servers. After provisioning, you can scale the service instance as needed by adding additional storage servers, compute servers, or both.

   X8M-2: If you select the flexible X8M-2 cloud infrastrure model, your initial Exadata Cloud Infrastructure instance can have a minimum of 2 database servers and 3 storage servers (the equivalents of an X8 quarter rack shape) up to 32 database servers and 64 storage servers. After provisioning, you can scale the service instance as needed by adding additional storage servers, compute servers, or both.

   X7 and X8: If you select an X7 or X8 system, you are given the choice of provisioning a quarter, half, or full rack. See *Exadata Fixed Hardware Shapes: X6, X7, X8 and Exadata Base* for hardware and capacity details.

   Exadata Base: The Exadata base shape comes in a single configuration, and provides an economical alternative to provisioning a quarter rack system. See *Exadata Fixed Hardware Shapes: X6, X7, X8 and Exadata Base*
8. If you selected a flexible shape (X8M-2 or X9M-2) , specify the **Compute and storage configuration**. You can specify **Database servers** from minimum of 2 up to 32. You can specify **Storage servers** from minimum of 3 up to 64.
9. Select the time zone: Choose one of :
   - **UTC**
   - **Select another timezone**
   - (Browser detected) timezone
10. In the **Provide maintenance details** pane.
    The default values for Maintenance Method and Maintenance Schedule are shown. Click **Edit Maintenance Preferences** to change the values

    In the **Edit Maintenance Preferences** dialog, under the
    Configure maintenance method: Click one of **Rolling** or **Non-rolling**.

- **Rolling**: By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
- **Non-rolling**: Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.

You can also select **Enable custom action before performing maintenance on DB servers**. Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. When the custom action is enabled the custom action timeout field appears.

- **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.

  Default: 30 minutes

  Maximum: 120 minutes

In the **Maintenance Schedule** section keep the default setting of **No preference** to let the system assign a date and start time for infrastructure maintenance, or select **Specify a Schedule**

a. Click the **Specify a schedule** radio button to choose your preferred month, week, weekday, and start time for infrastructure maintenance.

b. Under **Maintenance months**, specify at least one month for each quarter during which Exadata infrastructure maintenance will take place. You can select more than one month per quarter. If you specify a long lead time for advanced notification (for example, 4 weeks), then you may want to specify two or three months per quarter during which maintenance runs can occur. This will ensure that your maintenance updates are applied in a timely manner after accounting for your required lead time. Lead time is discussed in the following steps:

> **Note:**
>
> For Exadata infrastructure resources in government regions, Oracle performs maintenance operations monthly. Enable maintenance operations for all months if your infrastructure is in a government region.

c. Under **Week of the month**, specify which week of the month maintenance will take place. Weeks start on the 1st, 8th, 15th, and 22nd days of the month, and have a duration of seven days. Weeks start and end based on calendar dates, not days of the week. Maintenance cannot be scheduled for the fifth week of months that contain more than 28 days.

d. *(Optional)* Under **Day of the week**, specify the day of the week on which the maintenance will occur. If you do not specify a day of the week, then Oracle will run the maintenance update on a weekend day to minimize disruption.

e. *(Optional)* Under **Start hour**, specify the hour during which the maintenance run will begin. If you do not specify a start hour, then Oracle will choose the least disruptive time to run the maintenance update.

f. Under **Lead Time**, specify the number of weeks ahead of the maintenance event you would like to receive a notification message. Your lead time ensures that a newly

released maintenance update is scheduled to account for your required period of advanced notification.

g. Click **Save Changes**.

11. In the **Provide maintenance details** : **Provide up to 10 unique maintenance contact email addresses**. Click **Add Contact**.

    In the **Contact Email** field, provide the email ID of a desired contact.

    > **Note:**
    >
    > At leaset one Contact is required.

    Click **Add Contact** to add another contact.

12. Click **Show Advanced Options** to specify advanced options for the initial database.

    In the **Tags** tab, you can add tags to the database. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see *Resource Tags* . If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

13. Click **Create Exadata Infrastructure**. The Cloud Exadata infrastructure appears in the Exadata Infrastructure list with a status of Provisioning. The infrastructure's icon changes from yellow to green (or red to indicate errors).

WHAT NEXT?

After the Cloud Exadata infrastructure resource is successfully provisioned and in the Available status, you can create a cloud VM cluster as described in *To create a cloud VM cluster resource* on your infrastructure. You must provision both an infrastructure resource and a VM cluster before you can create your first database in the new Exadata Cloud Infrastructure instance.

Related Topics

-
- Resource Tags

To create a cloud VM cluster resource

Create a VM cluster in an Exadata Cloud Infrastructure instance.

> **Note:**
>
> To create a cloud VM cluster in an Exadata Cloud Infrastructure instance, you must have first created a Cloud Exadata infrastructure resource.

> **Note:**
>
> Multi-VM enabled Infrastructure will support creating multiple VM Clusters. Infrastructures created before the feature Create and Manage Multiple Virtual Machines per Exadata System (MultiVM) and VM Cluster Node Subsetting was released only support creating a single cloud VM cluster.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.

   > **Note:**
   >
   > Multiple VM clusters may be created only in a Multi-VM enabled Infrastructure.

3. Click **Create Exadata VM Cluster**.
   The **Create Exadata VM Cluster** page is displayed. Provide the required information to configure the VM cluster.
4. **Compartment:** Select a compartment for the VM cluster resource.
5. **Display name:** Enter a user-friendly display name for the VM cluster. The name doesn't need to be unique. An Oracle Cloud Identifier (OCID) will uniquely identify the DB system. Avoid entering confidential information.
6. **Select Exadata infrastructure:** Select the infrastructure resource that will contain the VM cluster. You must choose an infrastructure resource that has enough resources to create a new VM cluster. Click **Change Compartment** and pick a different compartment from the one you are working in to view infrastructure resources in other compartments.

   > **Note:**
   >
   > Multiple VM clusters may be created only in a Multi-VM enabled Infrastructure

7. **Choose the Oracle Grid Infrastructure version:** From the list, choose the Oracle Grid Infrastructure release (19c and 23ai) that you want to install on the VM cluster.
   The Oracle Grid Infrastructure release determines the Oracle Database releases that can be supported on the VM cluster. You cannot run an Oracle Database release that is later than the Oracle Grid Infrastructure software release.

   > **Note:**
   >
   > Minimum requirements for provisioning a VM Cluster with Grid Infrastructure 23ai:
   >
   > - Exadata Guest VM running Exadata System Software 23.1.8
   > - Exadata Infrastructure running Exadata System Software 23.1.x

8. **Choose an Exadata image version:**

- **Exadata infrastructure with Oracle Linux 7 and Exadata image version 22.1.10.0.0.230422:**
  - The **Change image** button is not enabled.
  - The Oracle Grid Infrastructure version defaults to 19.0.0.0.0.
  - The Exadata guest version will be the same as that of the host OS.
- **Exadata infrastructure with Oracle Linux 8 and Exadata image version 23.1.3.0.0.230613:**
  - The Exadata guest version defaults to the latest (23.1.3.0).
  - The Oracle Grid Infrastructure version defaults to 19.0.0.0.0
  - The **Change image** button is enabled.
  - Click **Change image**.
    The resulting Change image panel displays the list of available major versions of Exadata image (23.1.3.0 and 22.1.3.0).

    The most recent release for each major version is indicated by "(latest)".
  - Slide **Display all available versions**.
    Six past versions including the latest versions of Exadata images 23.1.3.0 and 22.1.3.0 are displayed.
  - Choose a version.
  - Click **Save Changes**.

9. **Configure the VM cluster:** Specify the DB servers to used for new VM cluster (by default all DB Servers are selected). Click **Change DB Servers** to select from the available DB servers. In the **Resource allocation per VM** pane:
   - Specify the number of OCPU cores you want to allocate to each of the VM cluster's virtual machine compute nodes. Minimum is 2 OCPU per VM . The read-only **Requested OCPU count for the Exadata VM cluster** field displays the total number of OCPU cores you are allocating.
   - Specify the **Memory per VM** to allocate to each VM. The minimum per VM is 30 GB.
   - Specify the **Local Storage per VM** to allocate local storage to each VM. The minimum per VM is 60 GB.
     Each time when you create a new VM cluster, the space remaining out of the total available space is utilized for the new VM cluster.

     In addition to `/u02`, you can specify the size of additional local file systems.

     For more information and instructions to specify the size for each individual VM, see Introduction to Scale Up or Scale Down Operations.
     - Click **Show additional local file systems configuration options**.
     - Specify the size of `/`, `/u01`, `/tmp`, `/var`, `/var/log`, `/var/log/audit`, and `/home` file systems as needed.

> **Note:**
>
> * You can only expand these file systems and cannot reduce the size once expanded.
> * Due to backup partitions and mirroring, the `/` and `/var` file systems will consume twice the space they were allocated, which is indicated in the read-only **Total allocated storage for / (GB) due to mirroring** and **Total allocated storage for /tmp (GB) due to mirroring** fields.

  - After creating the VM Cluster, check the **Exadata Resources** section on the **Exadata Infrastructure Details** page to check the file size allocated to the local storage (`/u02`) and local storage (additional file systems).

10. **Configure Exadata storage:** Specify the following:
    - **Specify the usable Exadata storage TB**. Specify the storage in multiples of 1 TB. Minimum: 2 TB
    - **Allocate storage for Exadata sparse snapshots:** Select this configuration option if you intend to use snapshot functionality within your VM cluster. If you select this option, the SPARSE disk group is created, which enables you to use VM cluster snapshot functionality for PDB sparse cloning. If you do not select this option, the SPARSE disk group is not created and snapshot functionality will not be available on any database deployments that are created in the environment.

      > **Note:**
      >
      > The storage configuration option for sparse snapshots cannot be changed after VM cluster creation.

    - **Allocate storage for local backups:** Select this option if you intend to perform database backups to the local Exadata storage within your Exadata Cloud Infrastructure instance. If you select this option, more space is allocated to the RECO disk group, which is used to store backups on Exadata storage. If you do not select this option, more space is allocated to the DATA disk group, which enables you to store more information in your databases.

      > **Note:**
      >
      > The storage configuration option for local backups cannot be changed after VM cluster creation.

11. **Add SSH key:** Add the public key portion of each key pair you want to use for SSH access to the DB system:
    - **Generate SSH key pair** (Default option) Select this radio button to generate an SSH keypair. Then in the dialog below click **Save private key** to download the key, and optionally click **Save public key** to download the key.
    - **Upload SSH key files:** Select this radio button to browse or drag and drop .pub files.
    - **Paste SSH keys:** Select this radio button to paste in individual public keys. To paste multiple keys, click **+ Another SSH Key**, and supply a single key for each entry.

12. **Configure the network settings:** Specify the following:

> **Note:**
>
> IP addresses (100.64.0.0/10) are used for Exadata Cloud Infrastructure X8M interconnect.

- **Virtual cloud network:** The VCN in which you want to create the VM cluster. Click **Change Compartment** to select a VCN in a different compartment.
- **Client subnet:** The subnet to which the VM cluster should attach. Click **Change Compartment** to select a subnet in a different compartment.
  Do not use a subnet that overlaps with 192.168.16.16/28, which is used by the Oracle Clusterware private interconnect on the database instance. Specifying an overlapping subnet causes the private interconnect to malfunction.
- **Backup subnet:** The subnet to use for the backup network, which is typically used to transport backup information to and from the **Backup Destination**, and for Data Guard replication. Click **Change Compartment** to select a subnet in a different compartment, if applicable.
  Do not use a subnet that overlaps with 192.168.128.0/20. This restriction applies to both the client subnet and backup subnet.

  If you plan to back up databases to Object Storage or Autonomous Recovery service, see the network prerequisites in Managing Exadata Database Backups.

  > **Note:**
  >
  > In case Autonomous Recovery Service is used, a new dedicated subnet is highly recommended. Review the network requirements and configurations required to backup your Oracle Cloud databases to Recovery Service. See, Configuring Network Resources for Recovery Service.

- **Network Security Groups:** Optionally, you can specify one or more network security groups (NSGs) for both the client and backup networks. NSGs function as virtual firewalls, allowing you to apply a set of ingress and egress *security rules* to your Exadata Cloud Infrastructure VM cluster. A maximum of five NSGs can be specified. For more information, see *Network Security Groups* and *Network Setup for Exadata Cloud Infrastructure Instances*.

  Note that if you choose a subnet with a *security list*, the security rules for the VM cluster will be a union of the rules in the security list and the NSGs.

  To use network security groups:

  – Check the **Use network security groups to control traffic** check box. This box appears under both the selector for the client subnet and the backup subnet. You can apply NSGs to either the client or the backup network, or to both networks. Note that you must have a virtual cloud network selected to be able to assign NSGs to a network.

  – Specify the NSG to use with the network. You might need to use more than one NSG. If you're not sure, contact your network administrator.

  – To use additional NSGs with the network, click **+;Another Network Security Group**.

- *To use private DNS Service*

  > **Note:**
  >
  > A Private DNS must be configured before it can be selected. See *Configure Private DNS*

  - Check the **Use private DNS Service** check box.
  - Select a private view. Click **Change Compartment** to select a private view in a different compartment.
  - Select a private zone. Click **Change Compartment** to select a private zone in a different compartment.

- **Hostname prefix:** Your choice of hostname for the Exadata DB system. The host name must begin with an alphabetic character and can contain only alphanumeric characters and hyphens (-). The maximum number of characters allowed for an Exadata DB system is 12.

  > **Caution:**
  >
  > The hostname must be unique within the subnet. If it is not unique, the VM cluster will fail to provision.

- **Host domain name:** The domain name for the VM cluster. If the selected subnet uses the Oracle-provided Internet and VCN Resolver for DNS name resolution, this field displays the domain name for the subnet and it can't be changed. Otherwise, you can provide your choice of the domain name. Hyphens (-) are not permitted.
  If you plan to store database backups in Object Storage or Autonomous Recovery service, Oracle recommends that you use a VCN Resolver for DNS name resolution for the client subnet because it automatically resolves the Swift endpoints used for backups.
- **Host and domain URL:** This read-only field combines the host and domain names to display the fully qualified domain name (FQDN) for the database. The maximum length is 63 characters.

13. **Choose a license type:** The type of license you want to use for the VM cluster. Your choice affects metering for billing.
    - **License Included** means the cost of the cloud service includes a license for the Database service.
    - **Bring Your Own License (BYOL)** means you are an Oracle Database customer with an Unlimited License Agreement or Non-Unlimited License Agreement and want to use your license with Oracle Cloud Infrastructure. This removes the need for separate on-premises licenses and cloud licenses.
14. **Diagnostics Collection:** By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.

> **Note:**
>
> You are opting in with the understanding that the above list of events (or metrics, log files) can change in the future. You can opt out of this feature at any time.

- **Enable Diagnostic Events**: Allow Oracle to collect and publish critical, warning, error, and information events to me.
- **Enable Health Monitoring**: Allow Oracle to collect health metrics/events such as Oracle Database up/down, disk space usage, and so on, and share them with Oracle Cloud operations. You will also receive notification of some events.
- **Enable Incident Logs and Trace Collection**: Allow Oracle to collect incident logs and traces to enable fault diagnosis and issue resolution.

> **Note:**
>
> You are opting in with the understanding that the above list of events (or metrics, log files) can change in the future. You can opt-out of this feature at any time.

All three checkboxes are selected by default. You can leave the default settings as is or clear the checkboxes as needed. You can view the Diagnostic Collection settings on the **VM Cluster Details** page under **General Information >> Diagnostics Collection**.

- **Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (all three options).
- **Disabled:** When you choose not to collect diagnostics, health metrics, incident logs, and trace files (all three options).
- **Partially Enabled**: When you choose to collect diagnostics, health metrics, incident logs, and trace files (one or two options).

15. Click **Show Advanced Options** to specify advanced options for the VM cluster:

- **Time zone:** This option is located in the **Management** tab. The default time zone for the DB system is UTC, but you can specify a different time zone. The time zone options are those supported in both the Java.util.TimeZone class and the Oracle Linux operating system. For more information, see *DB System Time Zone* .

> **Note:**
>
> If you want to set a time zone other than UTC or the browser-detected time zone, and if you do not see the time zone you want, try selecting the **Select another time zone**, option, then selecting "Miscellaneous" in the **Region or country** list and searching the additional **Time zone** selections.

- **SCAN Listener Port**: This option is located in the **Network** tab. You can assign a SCAN listener port (TCP/IP) in the range between 1024 and 8999. The default is 1521

> **Note:**
>
> Manually changing the SCAN listener port of a VM cluster after provisioning using the backend software is not supported. This change can cause Data Guard provisioning to fail.

.

- **Tags**: If you have permissions to create a resource, then you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.

16. Click **Create Exadata VM Cluster**.

WHAT NEXT?

After your VM cluster is successfully created and in the Available state, you can view the VM Cluster Details page by clicking the name of the VM cluster in the list of clusters. From the VM Cluster Details page, you can *create your first database* in the cluster by clicking **Create Database**.

Related Topics

- Network Security Groups
- Network Setup for Exadata Cloud Infrastructure Instances
  This topic describes the recommended configuration for the VCN and several related requirements for the Exadata Cloud Infrastructure instance.
- Security Lists
- Configure Private DNS
  Prerequistes needed to use Private DNS
- DB System Time Zone
- Resource Tags
- To create a database in an existing Exadata Cloud Infrastructure instance
  This topic covers creating your first or subsequent databases.
- Getting Started with Events
- Overview of Database Service Events
- Overview of Automatic Diagnostic Collection
  By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.

Configuring Network Resources for Recovery Service

Review the network requirements and configurations required to backup your Oracle Cloud databases to Recovery Service.

- About Using a Private Subnet for Recovery Service
  Recovery Service uses a private subnet inside a virtual cloud network (VCN) where your database resides. The private subnet defines the network path for backups between your database and Recovery Service.

- Networking Service Permissions to Configure a Private Subnet
  Review the policies that provide permissions to create and manage the networking components required to enable Recovery Service.
- Subnet Size Requirements and Security Rules for Recovery Service Subnet
  In the database VCN, include a security list with ingress rules defined to allow backup traffic between a database and Recovery Service. You must associate the security list with the private subnet used by Recovery Service.
- Creating a Recovery Service Subnet in the Database VCN
  In the OCI Console, configure a private subnet for Recovery Service in your database VCN. You must then register the Recovery Service subnet.

About Using a Private Subnet for Recovery Service

Recovery Service uses a private subnet inside a virtual cloud network (VCN) where your database resides. The private subnet defines the network path for backups between your database and Recovery Service.

Oracle recommends that your database VCN must have a single private subnet dedicated for backups to Recovery Service. Your Oracle Cloud database can reside in the same private subnet used by Recovery Service, or in a different subnet within the same VCN.

Use a private subnet with a minimum size of /24 (256 IP addresses). You can either create a subnet, or use a preexisting subnet in your database VCN.

> **Note:**
>
> Select an IPv4-only subnet for Recovery Service in your database VCN. Do not select an IPv6-enabled subnet as Oracle does not support using an IPv6-enabled subnet for Recovery Service operations. See Creating a Subnet to learn more.

Associate security lists with the private subnet. The security list must include stateful ingress rules to allow destination ports 8005 and 2484.

You must register the private subnet as a Recovery Service subnet to control backup traffic between your database and Recovery Service.

> **Note:**
>
> Oracle recommends using a private subnet for your backups, but it is possible to use a public subnet.

Networking Service Permissions to Configure a Private Subnet

Review the policies that provide permissions to create and manage the networking components required to enable Recovery Service.

Table 4-6 Networking Service Permissions Required to Create Subnets, Security Lists, Service Gateway, and Route Tables

Operation	Required IAM Policies
Configure a private subnet in a database VCN	• `use vcns` for the compartment which the VCN is in • `use subnets` for the compartment which the VCN is in • `manage private-ips` for the compartment which the VCN is in • `manage vnics` for the compartment which the VCN is in • manage vnics for the compartment which the database is provisioned or is to be provisioned in

Alternatively, you can create a policy that allows a specified group with broader access to networking components.

For example, use this policy to allow a `NetworkAdmin` group to manage all networks in any compartment in a tenancy.

Example 4-1 Policy for Network Administrators

```
Allow group NetworkAdmin to manage virtual-network-family in tenancy
```

Subnet Size Requirements and Security Rules for Recovery Service Subnet

In the database VCN, include a security list with ingress rules defined to allow backup traffic between a database and Recovery Service. You must associate the security list with the private subnet used by Recovery Service.

> **Note:**
>
> Select an IPv4-only subnet for Recovery Service in your database VCN. Do not select an IPv6-enabled subnet as Oracle does not support using an IPv6-enabled subnet for Recovery Service operations. See Creating a Subnet to learn more.

Table 4-7 Subnet size requirements and ingress rules for a private subnet used by Recovery Service

Item	Requirements
Minimum subnet size	/24 (256 IP addresses)
General ingress rule 1: Allow HTTPS traffic from Anywhere	This rule allows backup traffic from your Oracle Cloud Infrastructure Database to Recovery Service. • **Stateless**: No (all rules must be stateful) • **Source Type**: CIDR • **Source CIDR**: CIDR of the VCN where the database resides • **IP Protocol**: TCP • **Source Port Range**: All • **Destination Port Range**: 8005

Table 4-7 (Cont.) Subnet size requirements and ingress rules for a private subnet used by Recovery Service

Item	Requirements
General ingress rule 2: Allows SQLNet Traffic from Anywhere	This rule allows recovery catalog connections and real-time data protection from your Oracle Cloud Infrastructure Database to Recovery Service. • **Stateless**: No (all rules must be stateful) • **Source Type**: CIDR • **Source CIDR**: CIDR of the VCN where the database resides • **IP Protocol**: TCP • **Source Port Range**: All • **Destination Port Range**: 2484

> **Note:**
>
> If your database VCN restricts network traffic between subnets, then ensure to add an egress rule for ports 2484 and 8005 from the database subnet to the Recovery Service subnet that you create.

Creating a Recovery Service Subnet in the Database VCN

In the OCI Console, configure a private subnet for Recovery Service in your database VCN. You must then register the Recovery Service subnet.

1. In the navigation menu, select **Networking**, and then select **Virtual cloud networks** to display the Virtual Cloud Networks page.
2. Select the VCN in which your database resides.
3. Under **Resources**, select **Security Lists**.
4. Select the security list that is used for the VCN, and add two ingress rules to allow destination ports 8005 and 2484.
5. Click **Add Ingress Rule**, and add these details to set up a rule that allows HTTPS traffic from anywhere:
   a. Source Type: `CIDR`
   b. Source CIDR: Specify the CIDR of the VCN where the database resides.
   c. IP Protocol: `TCP`.
   d. Source Port Range: `All`
   e. Destination Port Range: `8005`.
   f. Description: Specify an optional description of the ingress rule to help manage the security rules.

   See: Subnet Size Requirements and Security Rules for Recovery Service Subnet.
6. Click **Add Ingress Rule**, and add these details to set up a rule that allows SQLNet traffic from anywhere:
   a. Source Type: `CIDR`

b. Source CIDR: Specify the CIDR of the VCN where the database resides.

c. IP Protocol: `TCP`.

d. Source Port Range: `All`

e. Destination Port Range: `2484`.

f. Description: Specify an optional description of the ingress rule to help manage the security rules.

See: Subnet Size Requirements and Security Rules for Recovery Service Subnet.

7. In the Virtual Cloud Networks Details page, click Create Subnet. Create a private subnet with a minimum subnet size of /24 (256 IP addresses). See, Overview of VCN and Subnets.
Alternatively, select a suitable private subnet that already exists in the VCN.

> **Note:**
>
> Select an IPv4-only subnet for Recovery Service in your database VCN. Do not select an IPv6-enabled subnet as Oracle does not support using an IPv6-enabled subnet for Recovery Service operations. See Creating a Subnet to learn more.

8. Associate the security list with the private subnet. The security list must include ingress rules to allow destination ports 8005 and 2484.
See: Security Lists.

9. Register the private subnet in Recovery Service. See: Register Recovery Service Subnets.
Oracle recommends that you register a single Recovery Service subnet per VCN.

> **Note:**
>
> If your database VCN restricts network traffic between subnets, then ensure to add an egress rule for ports 2484 and 8005 from the database subnet to the Recovery Service subnet that you create.

For additional configuration details, refer the relevant database service documentation.

Connecting to an Exadata Cloud Infrastructure Instance

This topic explains how to connect to an Exadata Cloud Infrastructure instance using SSH or SQL Developer.

How you connect depends on how your cloud network is set up. You can find information on various networking scenarios in Networking Overview, but for specific recommendations on how you should connect to a database in the cloud, contact your network security administrator.

> **Note:**
>
> Exadata Cloud Infrastructure servers cannot be joined to Active Directory domains, and the service does not support the use of Active Directory for user authentication and authorization.

- Prerequisites
  List of the requirements for SSH access to a compute node in an Exadata Cloud Infrastructure instance.
- SCAN Listener Port Setting
  When creating a cloud VM cluster, you can optionally designate a different SCAN listener port number.
- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- Using Oracle Net Services to Connect to a Database
  Oracle Database Exadata Cloud Infrastructure supports remote database access by using Oracle Net Services.

Prerequisites

List of the requirements for SSH access to a compute node in an Exadata Cloud Infrastructure instance.

You'll need the following:

- The full path to the file that contains the private key associated with the public key used when the system was launched.
- The public or private IP address of the Exadata Cloud Infrastructure instance.

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

  – *Cloud VM clusters (new resource model):* On the **Exadata VM Cluster Details** page, click Virtual Machines in the **Resources** list.

  – *DB systems:* On the **DB System Details** page, click **Nodes** in the **Resources** list.

  The values are displayed in the **Public IP Address** and **Private IP Address & DNS Name** columns of the table displaying the **Virtual Machines** or **Nodes** of the Exadata Cloud Infrastructure instance.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

SCAN Listener Port Setting

When creating a cloud VM cluster, you can optionally designate a different SCAN listener port number.

The default SCAN listener port for cloud VM clusters is 1521. When using the console To create a cloud VM cluster resource, you can optionally designate a different SCAN listener port number. In the OCI Console, this option appears under **Advanced Options** when creating the cluster.

> **Note:**
>
> Manually changing the SCAN listener port of a VM cluster after provisioning using the backend software is not supported. This change can cause Data Guard provisioning to fail.

Connecting to a Virtual Machine with SSH

You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

Most Unix-style systems (including Linux, Oracle Solaris, and macOS) include an SSH client. For Microsoft Windows systems, you can download a free SSH client called PuTTY from the following site: "http://www.putty.org".

- Connecting from a Unix-Style System
  To access a virtual machine on an Oracle ExaDB-D system from a Unix-style system using SSH, use this procedure.
- Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY
  Learn to access a virtual machine from a Microsoft Windows system using PuTTY.
- Accessing a Database After You Connect to the Virtual Machine
  After you connect to a virtual machine, you can use the following series of commands to identify a database and connect to it.

Related Topics

- http://www.putty.org/

Connecting from a Unix-Style System

To access a virtual machine on an Oracle ExaDB-D system from a Unix-style system using SSH, use this procedure.

- Enter the following SSH command to access the virtual machine:

  ```
  ssh -i private-key user@node
  ```

  In the preceding syntax:

  - *private-key* is the full path and name of the file that contains the SSH private key that corresponds to a public key that is registered in the system.
  - `user` is the operating system user that you want to use to connect:
    - To perform operations as the Oracle Database software owner, connect as as `opc` and `su oracle`. The `oracle` user does not have `root` user access to the virtual machine.
    - To perform operations that require `root` access to the virtual machine, such as patching, connect as `opc`. The `opc` user can use the `sudo -s` command to gain `root` access to the virtual machine.
  - *node* is the host name or IP address for the virtual machine that you want to access.

Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY

Learn to access a virtual machine from a Microsoft Windows system using PuTTY.

Before you begin

Before you use the PuTTY program to connect to a virtual machine, you need the following:

- The IP address of the virtual machine
- The SSH private key file that matches the public key associated with the deployment. This private key file must be in the PuTTY `.ppk` format. If the private key file was originally created on the Linux platform, you can use the PuTTYgen program to convert it to the `.ppk` format.

To connect to a virtual machine using the PuTTY program on Windows:

1. Download and install PuTTY.

   To download PuTTY, go to http://www.putty.org/ and click the **You can download PuTTY here** link.
2. Run the PuTTY program (`putty.exe`).

   The PuTTY Configuration window is displayed, showing the **Session** panel.
3. In the **Host Name (or IP address)** field, enter the host name or IP address of the virtual machine that you want to access.
4. Confirm that the **Connection type** option is set to **SSH**.
5. In the **Category** tree, expand **Connection** if necessary and then click **Data**.

   The **Data** panel is displayed.
6. In the **Auto-login username** field, enter the operating system user you want to connect as:

   - Connect as the user `opc` to perform operations that require `root` or `oracle` access to the virtual machine, such as backing up or patching; this user can use the `sudo` command to gain `root` or `oracle` access to the VM.
7. Confirm that the **When username is not specified** option is set to **Prompt**.
8. In the **Category** tree, expand **SSH** and then click **Auth**.

   The **Auth** panel is displayed.
9. Click the **Browse** button next to the **Private key file for authentication** field. Then, in the **Select private key file** window, navigate to and open the private key file that matches the public key that is associated with the deployment.
10. In the **Category** tree, click **Session**.

    The **Session** panel is displayed.
11. In the **Saved Sessions** field, enter a name for the connection configuration. Then, click **Save**.
12. Click **Open** to open the connection.

    The PuTTY Configuration window closes and the PuTTY terminal window displays.

    If this is the first time you are connecting to the VM, the PuTTY Security Alert window is displayed, prompting you to confirm the public key. Click **Yes** to continue connecting.

Accessing a Database After You Connect to the Virtual Machine

After you connect to a virtual machine, you can use the following series of commands to identify a database and connect to it.

1. SSH in as the `opc` user.
2. `sudo su oracle`
3. Use the `srvctl` utility located under the Oracle Grid Infrastructure home directory to list the databases on the system. For example:

```
/u01/app/12.2.0.1/grid/bin/srvctl config database -v
nc122   /u02/app/oracle/product/12.2.0/dbhome_6 12.2.0.1.0
s12c    /u02/app/oracle/product/12.2.0/dbhome_2 12.2.0.1.0
```

4. Identify the database instances for the database that you want to access. For example:

```
/u01/app/12.2.0.1/grid/bin/srvctl status database -d s12c
Instance s12c1 is running on node node01
Instance s12c2 is running on node node02
```

5. Configure the environment settings for the database that you want to access. For example:

```
. oraenv
ORACLE_SID = [oracle] ? s12c
The Oracle base has been set to /u02/app/oracle


export ORACLE_SID=s12c1
```

6. You can use the `svrctl` command to display more detailed information about the database. For example:

```
srvctl config database -d s12c
Database unique name: s12c
Database name:
Oracle home: /u02/app/oracle/product/12.2.0/dbhome_2
Oracle user: oracle
Spfile: +DATAC4/s12c/spfiles12c.ora
Password file: +DATAC4/s12c/PASSWORD/passwd
Domain: example.com
Start options: open
Stop options: immediate
Database role: PRIMARY
Management policy: AUTOMATIC
Server pools:
Disk Groups: DATAC4
Mount point paths:
Services:
Type: RAC
Start concurrency:
Stop concurrency:
OSDBA group: dba
OSOPER group: racoper
```

```
Database instances: s12c1,s12c2
Configured nodes: node01,node02
CSS critical: no
CPU count: 0
Memory target: 0
Maximum memory: 0
Default network number for database services:
Database is administrator managed
```

7. You can access the database by using SQL*Plus. For example:

```
sqlplus / as sysdba

SQL*Plus: Release 12.2.0.1.0 Production ...

Copyright (c) 1982, 2016, Oracle.  All rights reserved.

Connected to:
Oracle Database 12c EE Extreme Perf Release 12.2.0.1.0 - 64bit Production
```

Using Oracle Net Services to Connect to a Database

Oracle Database Exadata Cloud Infrastructure supports remote database access by using Oracle Net Services.

Because Exadata Cloud Infrastructure uses Oracle Grid Infrastructure, you can make Oracle Net Services connections by using **Single Client Access Name** (SCAN) connections. SCAN is a feature that provides a consistent mechanism for clients to access the Oracle Database instances running in a cluster.

By default, the SCAN is associated with three virtual IP addresses (VIPs). Each SCAN VIP is also associated with a SCAN listener that provides a connection endpoint for Oracle Database connections using Oracle Net Services. To maximize availability, Oracle Grid Infrastructure distributes the SCAN VIPs and SCAN listeners across the available cluster nodes. In addition, if there is a node shutdown or failure, then the SCAN VIPs and SCAN listeners are automatically migrated to a surviving node. By using SCAN connections, you enhance the ability of Oracle Database clients to have a reliable set of connection endpoints that can service all of the databases running in the cluster.

The SCAN listeners are in addition to the Oracle Net Listeners that run on every node in the cluster, which are also known as the node listeners. When an Oracle Net Services connection comes through a SCAN connection, the SCAN listener routes the connection to one of the node listeners, and plays no further part in the connection. A combination of factors, including listener availability, database instance placement, and workload distribution, determines which node listener receives each connection.

> **Note:**
>
> This documentation provides basic requirements for connecting to your Exadata Cloud Infrastructure databases by using Oracle Net Services.

- Prerequisites for Connecting to a Database with Oracle Net Services
  Review the prerequisites to connect to an Oracle Database instance on Oracle ExaDB-D using Oracle Net Services.
- Connecting to a Database with SQL Developer
  You can connect to a database with SQL Developer by using one of the following methods:
- Connecting to a Database Using SCAN
  To create an Oracle Net Services connection by using the SCAN listeners, you can choose between two approaches.
- Connecting to a Database Using a Node Listener
  To connect to an Oracle Database instance on Exadata Cloud Infrastructure with a connect descriptor that bypasses the SCAN listeners, use this procedure to route your connection directly to a node listener.

Prerequisites for Connecting to a Database with Oracle Net Services

Review the prerequisites to connect to an Oracle Database instance on Oracle ExaDB-D using Oracle Net Services.

To connect to an Oracle Database on Exadata Cloud Infrastructure with Oracle Net Services, you need the following:

- The IP addresses for your SCAN VIPs, or the hostname or IP address for a virtual machine that hosts the database that you want to access.
- The database identifier: Either the database system identifier (SID), or a service name.

Connecting to a Database with SQL Developer

You can connect to a database with SQL Developer by using one of the following methods:

- Create a temporary SSH tunnel from your computer to the database. This method provides access only for the duration of the tunnel. (When you are done using the database, be sure to close the SSH tunnel by exiting the SSH session.)
- Open the port used as the Oracle SCAN listener by updating the security list used for the cloud VM cluster or DB system resource in the Exadata Cloud Service instance. The default SCAN listener port is 1521. This method provides more durable access to the database. For more information, see Updating the Security List.

After you've created an SSH tunnel or opened the SCAN listener port as described above, you can connect to an Exadata Cloud Infrastructure instance using SCAN IP addresses or public IP addresses, depending on how your network is set up and where you are connecting from. You can find the IP addresses in the Console, in the **Database** details page.

- To connect using SCAN IP addresses
  You can connect to the database using the SCAN IP addresses if your client is on-premises and you are connecting using a FastConnect or Site-to-Site VPN connection.
- To connect using public IP addresses
  You can use the node's public IP address to connect to the database if the client and database are in different VCNs, or if the database is on a VCN that has an internet gateway.

To connect using SCAN IP addresses

You can connect to the database using the SCAN IP addresses if your client is on-premises and you are connecting using a FastConnect or Site-to-Site VPN connection.

You have the following options:

- Use the private SCAN IP addresses, as shown in the following `tnsnames.ora` example:

```
testdb=
  (DESCRIPTION =
    (ADDRESS_LIST=
      (ADDRESS = (PROTOCOL = TCP)(HOST = <scanIP1>)(PORT = 1521))
      (ADDRESS = (PROTOCOL = TCP)(HOST = <scanIP2>)(PORT = 1521)))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = <dbservice.subnetname.dbvcn.oraclevcn.com>)
    )
  )
```

- Define an external SCAN name in your on-premises DNS server. Your application can resolve this external SCAN name to the DB System's private SCAN IP addresses, and then the application can use a connection string that includes the external SCAN name. In the following `tnsnames.ora` example, `extscanname.example.com` is defined in the on-premises DNS server.

```
testdb =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = <extscanname.example.com>)(PORT =
1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = <dbservice.subnetname.dbvcn.oraclevcn.com>)
    )
  )
```

To connect using public IP addresses

You can use the node's public IP address to connect to the database if the client and database are in different VCNs, or if the database is on a VCN that has an internet gateway.

However, there are important implications to consider:

- When the client uses the public IP address, the client bypasses the SCAN listener and reaches the node listener, so server side load balancing is not available.
- When the client uses the public IP address, it cannot take advantage of the VIP failover feature. If a node becomes unavailable, new connection attempts to the node will hang until a TCP/IP timeout occurs. You can set client side sqlnet parameters to limit the TCP/IP timeout.

The following `tnsnames.ora` example shows a connection string that includes the CONNECT_TIMEOUT parameter to avoid TCP/IP timeouts.

```
test=
  (DESCRIPTION =
    (CONNECT_TIMEOUT=60)
    (ADDRESS_LIST=
      (ADDRESS = (PROTOCOL = TCP)(HOST = <publicIP1>)(PORT = 1521))
      (ADDRESS = (PROTOCOL = TCP)(HOST = <publicIP2>)(PORT = 1521))
    )
    (CONNECT_DATA =
```

```
    (SERVER = DEDICATED)
    (SERVICE_NAME = <dbservice.subnetname.dbvcn.oraclevcn.com>)
  )
)
```

Connecting to a Database Using SCAN

To create an Oracle Net Services connection by using the SCAN listeners, you can choose between two approaches.

- Connecting to a Database Using a Connect Descriptor that References All of the SCAN VIPs
  You can set up a connect descriptor for Oracle Exadata Database Service on Dedicated Infrastructure System using multiple SCAN listeners.
- Connecting to a Database Use a Connect Descriptor that References a Custom SCAN Name
  You can set up a connect descriptor for Oracle Exadata Database Service on Dedicated Infrastructure System using a custom SCAN name.

Connecting to a Database Using a Connect Descriptor that References All of the SCAN VIPs

You can set up a connect descriptor for Oracle Exadata Database Service on Dedicated Infrastructure System using multiple SCAN listeners.

This approach requires you to supply all of the single client access name (SCAN) virtual IP (VIP) addresses, and enables Oracle Net Services to connect to an available SCAN listener.

- Use the following template to define a Net Services alias, which is typically used to provide a convenient name for the connect descriptor:

```
alias-name = (DESCRIPTION=
  (ADDRESS_LIST=
    (ADDRESS=(PROTOCOL=tcp)(HOST=SCAN-VIP-1)(PORT=1521))
    (ADDRESS=(PROTOCOL=tcp)(HOST=SCAN-VIP-2)(PORT=1521))
    (ADDRESS=(PROTOCOL=tcp)(HOST=SCAN-VIP-3)(PORT=1521)))
  (CONNECT_DATA=(sid-or-service-entry)))
```

  Where:

  alias-name is the name you use to identify the alias.

  SCAN-VIP-[1-3] are the IP addresses for the SCAN VIPs.

  sid-or-service-entry identifies the database SID or service name using one of the following formats:

  - `SID=sid-name`. For example: `SID=S12C1`.
  - `SERVICE_NAME=service-name`. For example: `SERVICE_NAME=PDB1.example.yourcloud.com`.

> **Note:**
>
> By default, Oracle Net Services randomly selects one of the addresses in the address list to balance the load between the SCAN listeners.

Connecting to a Database Use a Connect Descriptor that References a Custom SCAN Name

You can set up a connect descriptor for Oracle Exadata Database Service on Dedicated Infrastructure System using a custom SCAN name.

Using this approach, you define a custom single client access name (SCAN) name in your domain name server (DNS), which resolves to the three SCAN virtual IP addresses (VIPs).

- Use the following template to define a Net Services alias that references the custom SCAN name:

```
alias-name = (DESCRIPTION=
  (ADDRESS_LIST=(ADDRESS=(PROTOCOL=tcp)(HOST=scan-name)(PORT=1521)))
  (CONNECT_DATA=(sid-or-service-entry)))
```

Where:

alias-name is the name you use to identify the alias.

scan-name is the custom SCAN name.

sid-or-service-entry identifies the database SID or service name using one of the following formats:

- `SID=`*sid-name*. For example: `SID=S12C1`.
- `SERVICE_NAME=`*service-name*. For example: `SERVICE_NAME=PDB1.example.yourcloud.com`.

Alternatively, you can use the easy connect method to specify a connect descriptor with the following format:

```
scan-name:1521/sid-or-service-entry
```

For example:

```
exa1scan.example.com:1521/S12C1
```

Or

```
exa1scan.example.com:1521/PDB1.example.yourcloud.com
```

Connecting to a Database Using a Node Listener

To connect to an Oracle Database instance on Exadata Cloud Infrastructure with a connect descriptor that bypasses the SCAN listeners, use this procedure to route your connection directly to a node listener.

By using this method, you give up the high-availability and load-balancing provided by SCAN. However, this method may be desirable if you want to direct connections to a specific node or network interface. For example, you might want to ensure that connections from a program that performs bulk data loading use the backup network.

Using this approach, you direct your connection using the hostname or IP address of the node.

Example 4-2 Defining a Net Service Alias That Directly References the Node

```
alias-name = (DESCRIPTION=
  (CONNECT_TIMEOUT=timeout)
  (ADDRESS_LIST=(ADDRESS=(PROTOCOL=tcp)(HOST=node)(PORT=1521)))
  (CONNECT_DATA=(sid-or-service-entry)))
```

Where:

alias-name is the name you use to identify the alias.

timeout specifies a timeout period (in seconds), which enables you to terminate a connection attempt without having to wait for a TCP timeout. The (CONNECT_TIMEOUT=*timeout*) parameter is optional.

node is the hostname or IP address for the virtual machine that you want to use.

sid-or-service-entry identifies the database SID or service name using one of the following formats:

- SID=*sid-name*. For example, SID=S12C1.
- SERVICE_NAME=*service-name*. For example, SERVICE_NAME=PDB1.example.oraclecloudatcust.com.

Alternatively, you can use the easy connect method to specify a connect descriptor with the following format:

```
node:1521/sid-or-service-entry
```

For example:

```
exa1node01.example.com:1521/S12C1
```

Or

```
exa1node01.example.com:1521/PDB1.example.oraclecloudatcust.com
```

Best Practices for Exadata Cloud Infrastructure Instances

Oracle recommends that you follow these best practice guidelines to ensure the manageability of your Exadata Cloud Infrastructure instance:

When followed best practice guidelines can prevent problems that could affect the manageability and performance of yourExadata Cloud Infrastructure instance:

- Wherever possible, use the Oracle-supplied cloud interfaces such as the Oracle Cloud InfrastructureConsole, API, or CLI, or cloud-specific tools such as `dbaascli` to perform lifecycle management and administrative operations on your Exadata Cloud Infrastructure instance. For example, use the OCI console, API, CLI, or dbaascli to apply Oracle Database patches instead of manually running `opatch`. In addition, if an operation can be performed by using the Console as well as a command line utility, Oracle recommends that you use the Console. For example, use the Console instead of using `dbaascli` to create databases.
- Do not change the compute node OS users or manually manipulate SSH key settings associated with your Exadata DB system.
- Apply *only* patches that are available through the Database service. Do *not* apply patches from any other source unless you are directed to do so by Oracle Support.
- Apply the quarterly patches regularly, every quarter if possible.
- Do not change the ports for Oracle Net Listener.

Moving to Oracle Cloud Using Zero Downtime Migration

Oracle now offers the Zero Downtime Migration service, a quick and easy way to move on-premises databases to Oracle Cloud Infrastructure.

Zero Downtime Migration leverages Oracle Active Data Guard to create a standby instance of your database in an Oracle Cloud Infrastructure system. You switch over only when you are ready, and your source database remains available as a standby. Use the Zero Downtime Migration service to migrate databases individually or at the fleet level. See *Move to Oracle Cloud Using Zero Downtime Migration* for more information.

Related Topics

- Move to Oracle Cloud Using Zero Downtime Migration

5

How-to Guides

A collection of tasks and procedures for managing Exadata Database Service on Dedicated Infrastructure.

- Manage Database Security with Oracle Data Safe
- Connecting to an Exadata Cloud Infrastructure Instance
  This topic explains how to connect to an Exadata Cloud Infrastructure instance using SSH or SQL Developer.
- Manage Exadata Cloud Infrastructure
  Use the provided tools to manage the Infrastructure.
- Cloud Infrastructure Maintenance Updates
  Oracle performs the updates to all of the Oracle-managed infrastructure components on Exadata Cloud Infrastructure.
- Manage VM Clusters
  Learn how to manage your VM clusters on Exadata Cloud Infrastructure.
- Manage Software Images
- Create Oracle Database Homes on an Exadata Cloud Infrastructure System
  Learn to create Oracle Database Homes on Exadata Cloud Infrastructure.
- Managing Oracle Database Homes on an Exadata Cloud Infrastructure Instance
  You can delete or view information about Oracle Database Homes (referred to as "Database Homes" in Oracle Cloud Infrastructure) by using the Oracle Cloud Infrastructure Console, the API, or the CLI.
- Manage Databases on Exadata Cloud Infrastructure
- Manage Database Backup and Recovery on Oracle Exadata Database Service on Dedicated Infrastructure
  Learn how to work with the backup and recovery facilities provided by Oracle Exadata Database Service on Dedicated Infrastructure.
- Patch and Update an Exadata Cloud Infrastructure System
- Interim Software Updates
  For authorized environments, learn how to download interim software updates.
- Use Oracle Data Guard with Exadata Cloud Infrastructure
  Learn to configure and manage Data Guard associations in your VM cluster.
- Configure Oracle Database Features for Exadata Cloud Infrastructure
  This topic describes how to configure Oracle Multitenant, tablespace encryption, and Huge Pages for use with your Exadata Cloud Infrastructure instance.
- Managing Exadata Cloud Infrastructure I/O Resource Management (IORM)
- Migrate to Exadata Cloud Infrastructure
  For general guidance on methods and tools to migrate databases to Oracle Cloud Infrastructure database services, including Exadata Cloud Infrastructure see "Migrating Databases to the Cloud".

- Switch an Exadata DB System to the New Resource Model and APIs
  If you have existing Exadata DB systems in Oracle Cloud Infrastructure, you can switch them to the new resource model and APIs.
- Connect Identity and Access Management (IAM) Users to Oracle Exadata Database Service on Dedicated Infrastructure
  You can configure Oracle Exadata Database Service on Dedicated Infrastructure to use Oracle Cloud Infrastructure Identity and Access Management (IAM) authentication and authorization to allow IAM users to access an Oracle Database with IAM credentials.
- Authenticating and Authorizing Microsoft Azure Active Directory Users for Oracle Databases
  An Oracle Database instance can be configured for Microsoft Azure AD users to connect using Azure `OAuth2` access tokens.

Interim Software Updates

For authorized environments, learn how to download interim software updates.

This feature enables cloud-only customers to download one-off patches from the OCI console and API. There is no option to apply the downloaded patch via console and API. To apply these patches, customers must log in to their VM and run the patch apply utility.

> **Note:**
>
> To be able to download interim software update, you should at least have an ExaDB-D infrastructure provisioned.

Downloading one-off patches does not replace Database Software Image (DSI) creation. Customers must continue to use Database Software Images (DSI) to build and deploy their customized images.

- Create an Interim Software Update
- Download an Interim Software Update
- Delete an Interim Software Update
- Move an Interim Software Update Resource to Another Compartment
- Using the API to Manage Interim Software Updates

Create an Interim Software Update

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata on Oracle Public Cloud**.
2. Under **Resources**, click **Interim software updates**.

   Interim software update page is displayed.
3. Click **Create interim software update**.

   Create interim software update panel is displayed.
4. Enter the following details in the panel:
   a. **Name**: Descriptive name for the patch download path.

b. **Compartment**: Select a compartment where you want to create the patch resource.
c. **Database version**: Choose the Database version for your image.
d. **Release Update**: Choose any supported Oracle Database release update (RU).
e. **Interim software update number**: Optionally, enter an interim patch number.
f. **Tag**: Apply a tag.

5. Click **Create**.

Download an Interim Software Update

The patch download path is valid for four days. Download the patch within the specified timeframe.

1. On the Update details page, click **Download**.

   The system starts downloading the patch.

2. You can also download a patch from the Interim Software Updates page.

   - Click the Actions button (three dots) for the patch you're interested in, and select **Download**.

> **Note:**
>
> You can only download the patches that are in **Available** state.

Interim Software Updates Lifecycle States:

- **Available**: Patch has been created successfully and the time-to-live (TTL) has not expired.
- **Creating**: The patch creation process is in progress.
- **Expired**: The lifetime of the patch download link has expired, which means you cannot download it.
- **Failed**: The patch create failed due to some error.
- **Terminating**: The patch deletion process is in progress.
- **Terminated**: The patch has been deleted.

Delete an Interim Software Update

Be discrete in deleting interim software updates. However, you can delete the interim software updates that have expired to free up space in the Object Store.

1. On the Update details page, click **Delete**.
2. In the resulting dialog, enter the name of the patch to confirm and then click **Delete**.
3. You can also delete a patch from the Interim Software Updates page.

   - Click the Actions button (three dots) for the patch you're interested in, and select **Delete**.

Move an Interim Software Update Resource to Another Compartment

1. On the Update details page, click **Move Resource**.
2. In the resulting dialog, choose a new compartment, and click **Move Resource**.
3. You can also move a patch resource from the Interim Software Updates page.
   - Click the Actions button (three dots) for the patch you're interested in, and select **Move Resource**.

Using the API to Manage Interim Software Updates

ExaDB-C@C and ExaDB-D use the same API to manage interim software updates.

For information about using the API and signing requests, see *REST APIs* and *Security Credentials*. For information about SDKs, see *Software Development Kits and Command Line Interface*.

Use these API operations to manage interim software updates:

- `CreateOneoffPatch`
- `DeleteOneoffPatch`
- `DownloadOneoffPatch`
- `UpdateOneoffPatch`
- `ListOneoffPatches`
- `GetOneoffPatch`
- `ChangeOneoffPatchCompartment`

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- OneoffPatch Reference

Manage Database Security with Oracle Data Safe

About Oracle Data Safe

Your corporate policy requires that you monitor your databases and retain audit records. Your developers are asking for copies of production data for that new application, and you're wondering what kinds of sensitive information it will contain. Meanwhile, you need to make sure that recent maintenance activities haven't left critical security configuration gaps on your

production databases and that staff changes haven't left dormant user accounts on the databases. Oracle Data Safe assists you with these tasks and is included with your Exadata Database Service*.

Oracle Data Safe is a unified control center, that helps you to manage the day-to-day security and compliance requirements of Oracle Databases no matter if they are running in the Oracle Cloud Infrastructure, at Cloud@Customer, on-premises or in any other cloud.

Data Safe supports you to evaluate security controls, assess user security, monitor user activity, and address data security compliance requirements for your database by evaluating the sensitivity of your data as well as masking sensitive data for non-production databases.

Data Safe provides the following features:

- **Security Assessment**: Configuration errors and configuration drift are significant contributors to data breaches. Use security assessment to evaluate your database's configuration and compare it to Oracle and industry best practices. Security assessment reports on areas of risk and notifies you when configurations change.
- **User Assessment**: Many breaches start with a compromised user account. User Assessment helps you spot the riskiest database accounts - those accounts which, if compromised, could cause the most damage - and take proactive steps to secure them. User Assessment Baselines make it easy to know when new accounts are added, or an account's privileges are modified. Use OCI events to receive proactive notifications when a database deviates from its baseline.
- **Activity Auditing**: Understanding and reporting on user activity, data access, and changes to database structures supports regulatory compliance requirements and can aid in post-incident investigations. Activity auditing collects audit records from databases and helps you manage audit policies. Audit insights make it easy to identify inefficient audit policies, while alerts based on audit data proactively notify you of risky activity.
- **Sensitive Data Discovery**: Knowing what sensitive data is managed in your applications is critical for security and privacy. Data discovery scans your database for over 150 different types of sensitive data, helping you understand what types and how much sensitive data you are storing. Use these reports to formulate audit policies, develop data masking templates, and create effective access control policies.
- **Data Masking**: Minimizing the amount of sensitive data your organization maintains helps you meet compliance requirements and satisfy data privacy regulations. Data masking helps you remove risk from your non-production databases by replacing sensitive information with masked data. With reusable masking templates, over 50 included masking formats, and the ability to easily create custom formats for your organization's unique requirements, data masking can streamline your application development and testing operations.

**Includes 1 million audit records per database per month if using the audit collection for Activity Auditing*

Get Started

To get started you just need to register your database with Oracle Data Safe:

- Pre-requisite: Obtain the necessary Identity and Access Management (IAM) permissions to register your target database in Data Safe: Permissions to Register an Oracle Cloud Database with Oracle Data Safe
- Connecting your database to Data Safe

  - If your database is running in a private virtual cloud network (VCN), you can connect it to Data Safe via a **Data Safe private endpoint**.

    The private endpoint essentially represents the Oracle Data Safe service in your VCN with a private IP address in a subnet of your choice.

    You can create the private endpoint in the VCN of your database either before the registration or during the registration process. You can find more details on how to create the private endpoint under Create an Oracle Data Safe Private Endpoint.

- Register your database in Data Safe

Using Oracle Data Safe

Once your database is registered in Data Safe, you can leverage all features.

Security Assessment

Security Assessments are automatically scheduled once a week in Data Safe and provide an overall picture of your database security posture. It analyzes your database configurations, users and user entitlements, as well as security policies to uncover security risks and improve the security posture of Oracle Databases within your organization. A security assessment provides findings with recommendations for remediation activities that follow best practices to reduce or mitigate risk.

Start by reviewing the security assessment report for your database: View the latest assessment for a target database

You can find more details on Security Assessment under Security Assessment Overview.

User Assessment

User Assessments are automatically scheduled once a week in Data Safe and help you to identify highly privileged user accounts that could pose a threat if misused or compromised. User Assessment reviews information about your users in the data dictionaries on your target databases and then calculates a potential risk for each user, based on system privileges and role grants.

Start by reviewing the user assessment report for your database: View the latest user assessment for a target database

You can find more details on User Assessment under User Assessment Overview.

Data Discovery

Data Discovery searches for sensitive columns in your database. It comes with over 150 pre-defined sensitive types and you can also create your own sensitive types. You tell Data Discovery if you want to scan your entire database or just certain schemas and what type of sensitive information to look for, and it finds the sensitive columns that meet your criteria and stores them in a sensitive data model (SDM).

Start by discovering sensitive data in your database: Create Sensitive Data Models

You can find more details on Data Discovery under Data Discovery Overview.

Data Masking

Data masking, also known as static data masking helps you to replace sensitive or confidential information in your non-production databases with realistic and fully functional data with similar

characteristics as the original data. Data Safe comes with pre-defined masking formats for each of the pre-defined sensitive types that can also be leveraged for your own sensitive types.

Once you know where sensitive data is stored in your database (for instance after running Data Discovery in Data Safe), you can start by creating a masking policy: Create Masking Policies

After you created a masking policy and copied your production database, you can mask your non-production copy: Mask Sensitive Data on a Target Database

You can find more details on Data Masking under Data Masking Overview.

Activity Auditing

Activity Auditing in Oracle Data Safe helps to ensure accountability and improve regulatory compliance. With Activity Auditing, you can collect and retain audit records per industry and regulatory compliance requirements and monitor user activities on Oracle databases with pre-defined reports and alerts. For example, you can audit access to sensitive data, security-relevant events, administrator and user activities, activities recommended by compliance regulations like the Center for Internet Security (CIS), and activities defined by your own organization.

If you are using the audit collection in Data Safe, up to 1 million audit records per target database per month are included for your Cloud@Customer database.

To use activity auditing, start the audit trail for your target database in Data Safe: Start an Audit Trail

Once the audit trail is started, you can monitor and analyze your audit data with pre-defined audit reports: View a Predefined or Custom Audit Report

You can find more details on Activity Auditing under Activity Auditing Overview.

Connecting to an Exadata Cloud Infrastructure Instance

This topic explains how to connect to an Exadata Cloud Infrastructure instance using SSH or SQL Developer.

How you connect depends on how your cloud network is set up. You can find information on various networking scenarios in Networking Overview, but for specific recommendations on how you should connect to a database in the cloud, contact your network security administrator.

> **Note:**
>
> Exadata Cloud Infrastructure servers cannot be joined to Active Directory domains, and the service does not support the use of Active Directory for user authentication and authorization.

- Prerequisites
  List of the requirements for SSH access to a compute node in an Exadata Cloud Infrastructure instance.
- About Connecting to a Compute Node with SSH
  You can connect to the compute nodes in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

- Connect to the Exadata Cloud Infrastructure Service
  Learn how to connect to an Exadata Cloud Infrastructure system using SSH, and how to connect to an Exadata Cloud Infrastructure database using Oracle Net Services (SQL*Net).

Prerequisites

List of the requirements for SSH access to a compute node in an Exadata Cloud Infrastructure instance.

You'll need the following:

- The full path to the file that contains the private key associated with the public key used when the system was launched.
- The public or private IP address of the Exadata Cloud Infrastructure instance.

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

  – *Cloud VM clusters (new resource model):* On the **Exadata VM Cluster Details** page, click Virtual Machines in the **Resources** list.

  – *DB systems:* On the **DB System Details** page, click **Nodes** in the **Resources** list.

  The values are displayed in the **Public IP Address** and **Private IP Address & DNS Name** columns of the table displaying the **Virtual Machines** or **Nodes** of the Exadata Cloud Infrastructure instance.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

About Connecting to a Compute Node with SSH

You can connect to the compute nodes in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

Most Unix-style systems (including Linux, Oracle Solaris, and Apple MacOS) include an SSH client. For Microsoft Windows, you can download a free SSH client called PuTTY from the following address: http://www.putty.org

- Connecting from a Unix-Style System
  To access a virtual machine on an Oracle ExaDB-D system from a Unix-style system using SSH, use this procedure.
- Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY
  Learn to access a virtual machine from a Microsoft Windows system using PuTTY.
- To access a database after you connect to the compute node
  To connect to the database, you set environment information for the database.

Connecting from a Unix-Style System

To access a virtual machine on an Oracle ExaDB-D system from a Unix-style system using SSH, use this procedure.

- Enter the following SSH command to access the virtual machine:

  ```
  ssh -i private-key user@node
  ```

  In the preceding syntax:

  - *private-key* is the full path and name of the file that contains the SSH private key that corresponds to a public key that is registered in the system.
  - `user` is the operating system user that you want to use to connect:
    - To perform operations as the Oracle Database software owner, connect as as `opc` and `su oracle`. The `oracle` user does not have `root` user access to the virtual machine.
    - To perform operations that require `root` access to the virtual machine, such as patching, connect as `opc`. The `opc` user can use the `sudo -s` command to gain `root` access to the virtual machine.
  - *node* is the host name or IP address for the virtual machine that you want to access.

Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY

Learn to access a virtual machine from a Microsoft Windows system using PuTTY.

Before you begin

Before you use the PuTTY program to connect to a virtual machine, you need the following:

- The IP address of the virtual machine
- The SSH private key file that matches the public key associated with the deployment. This private key file must be in the PuTTY `.ppk` format. If the private key file was originally created on the Linux platform, you can use the PuTTYgen program to convert it to the `.ppk` format.

To connect to a virtual machine using the PuTTY program on Windows:

1. Download and install PuTTY.

   To download PuTTY, go to http://www.putty.org/ and click the **You can download PuTTY here** link.
2. Run the PuTTY program (`putty.exe`).

   The PuTTY Configuration window is displayed, showing the **Session** panel.
3. In the **Host Name (or IP address)** field, enter the host name or IP address of the virtual machine that you want to access.
4. Confirm that the **Connection type** option is set to **SSH**.
5. In the **Category** tree, expand **Connection** if necessary and then click **Data**.

   The **Data** panel is displayed.

6. In the **Auto-login username** field, enter the operating system user you want to connect as:

   - Connect as the user `opc` to perform operations that require `root` or `oracle` access to the virtual machine, such as backing up or patching; this user can use the `sudo` command to gain `root` or `oracle` access to the VM.

7. Confirm that the **When username is not specified** option is set to **Prompt**.
8. In the **Category** tree, expand **SSH** and then click **Auth**.

   The **Auth** panel is displayed.

9. Click the **Browse** button next to the **Private key file for authentication** field. Then, in the **Select private key file** window, navigate to and open the private key file that matches the public key that is associated with the deployment.
10. In the **Category** tree, click **Session**.

    The **Session** panel is displayed.

11. In the **Saved Sessions** field, enter a name for the connection configuration. Then, click **Save**.
12. Click **Open** to open the connection.

    The PuTTY Configuration window closes and the PuTTY terminal window displays.

    If this is the first time you are connecting to the VM, the PuTTY Security Alert window is displayed, prompting you to confirm the public key. Click **Yes** to continue connecting.

To access a database after you connect to the compute node

To connect to the database, you set environment information for the database.

1. Log in as opc and then use sudo to connect as the oracle user.

```
login as: opc

[opc@<host_name> ~]$ sudo su - oracle
```

2. Source the database's `.env` file to set the environment.

```
[oracle@<host_name>]# . <database_name>.env
```

   In the following example, the host name is "ed1db01" and the database name is "cdb01".

```
[oracle@ed1db01]# . cdb01.env
ORACLE_SID = [root] ? +ASM1
The Oracle base has been set to /u01/app/grid
```

Connect to the Exadata Cloud Infrastructure Service

Learn how to connect to an Exadata Cloud Infrastructure system using SSH, and how to connect to an Exadata Cloud Infrastructure database using Oracle Net Services (SQL*Net).

- Connecting to a Database with SQL Developer
  You can connect to a database with SQL Developer by using one of the following methods:

- Connecting to a Database with Oracle Net Services
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system using Oracle Net Services.

Connecting to a Database with SQL Developer

You can connect to a database with SQL Developer by using one of the following methods:

- Create a temporary SSH tunnel from your computer to the database. This method provides access only for the duration of the tunnel. (When you are done using the database, be sure to close the SSH tunnel by exiting the SSH session.)
- Open the port used as the Oracle SCAN listener by updating the security list used for the cloud VM cluster or DB system resource in the Exadata Cloud Service instance. The default SCAN listener port is 1521. This method provides more durable access to the database. For more information, see Updating the Security List.

After you've created an SSH tunnel or opened the SCAN listener port as described above, you can connect to an Exadata Cloud Infrastructure instance using SCAN IP addresses or public IP addresses, depending on how your network is set up and where you are connecting from. You can find the IP addresses in the Console, in the **Database** details page.

Connecting to a Database with Oracle Net Services

You can connect to the virtual machines in an Exadata Cloud Infrastructure system using Oracle Net Services.

- Using Oracle Net Services to Connect to a Database
  Oracle Database Exadata Cloud Infrastructure supports remote database access by using Oracle Net Services.
- Prerequisites for Connecting to a Database with Oracle Net Services
  Review the prerequisites to connect to an Oracle Database instance on Oracle ExaDB-D using Oracle Net Services.
- Connecting to a Database Using SCAN
  To create an Oracle Net Services connection by using the SCAN listeners, you can choose between two approaches.
- Connecting to a Database Using a Node Listener
  To connect to an Oracle Database instance on Exadata Cloud Infrastructure with a connect descriptor that bypasses the SCAN listeners, use this procedure to route your connection directly to a node listener.

Using Oracle Net Services to Connect to a Database

Oracle Database Exadata Cloud Infrastructure supports remote database access by using Oracle Net Services.

Because Exadata Cloud Infrastructure uses Oracle Grid Infrastructure, you can make Oracle Net Services connections by using **Single Client Access Name** (SCAN) connections. SCAN is a feature that provides a consistent mechanism for clients to access the Oracle Database instances running in a cluster.

By default, the SCAN is associated with three virtual IP addresses (VIPs). Each SCAN VIP is also associated with a SCAN listener that provides a connection endpoint for Oracle Database connections using Oracle Net Services. To maximize availability, Oracle Grid Infrastructure distributes the SCAN VIPs and SCAN listeners across the available cluster nodes. In addition, if there is a node shutdown or failure, then the SCAN VIPs and SCAN listeners are

automatically migrated to a surviving node. By using SCAN connections, you enhance the ability of Oracle Database clients to have a reliable set of connection endpoints that can service all of the databases running in the cluster.

The SCAN listeners are in addition to the Oracle Net Listeners that run on every node in the cluster, which are also known as the node listeners. When an Oracle Net Services connection comes through a SCAN connection, the SCAN listener routes the connection to one of the node listeners, and plays no further part in the connection. A combination of factors, including listener availability, database instance placement, and workload distribution, determines which node listener receives each connection.

> **Note:**
>
> This documentation provides basic requirements for connecting to your Exadata Cloud Infrastructure databases by using Oracle Net Services.

Prerequisites for Connecting to a Database with Oracle Net Services

Review the prerequisites to connect to an Oracle Database instance on Oracle ExaDB-D using Oracle Net Services.

To connect to an Oracle Database on Exadata Cloud Infrastructure with Oracle Net Services, you need the following:

- The IP addresses for your SCAN VIPs, or the hostname or IP address for a virtual machine that hosts the database that you want to access.
- The database identifier: Either the database system identifier (SID), or a service name.

Connecting to a Database Using SCAN

To create an Oracle Net Services connection by using the SCAN listeners, you can choose between two approaches.

- Identifying IP Addresses Using the SDK or CLI
  You can use the SDK or the OCI CLI to identify the IP addresses of Exadata Cloud Infrastructure compute nodes. You can then use the IP addresses to connect to your system.
- Connecting to a Database Using a Connect Descriptor that References All of the SCAN VIPs
  You can set up a connect descriptor for Oracle Exadata Database Service on Dedicated Infrastructure System using multiple SCAN listeners.
- Connecting to a Database Use a Connect Descriptor that References a Custom SCAN Name
  You can set up a connect descriptor for Oracle Exadata Database Service on Dedicated Infrastructure System using a custom SCAN name.

Identifying IP Addresses Using the SDK or CLI

You can use the SDK or the OCI CLI to identify the IP addresses of Exadata Cloud Infrastructure compute nodes. You can then use the IP addresses to connect to your system.

NOT_SUPPORTED

1. Use the GetDbNode API to return the details of the Exadata Cloud InfrastructuredbNode. Note the OCIDs returned for the `hostIpId` and `backupIpId` parameters of the dbNode.
2. With the OCIDs found in the `hostIpId` and `backupIpId` parameters, you can use the GetPrivateIp API to get the private IP addresses used by the client and backup subnets. For public subnet IP addresses, use the GetPublicIpByPrivateIpId API.

Connecting to a Database Using a Connect Descriptor that References All of the SCAN VIPs

You can set up a connect descriptor for Oracle Exadata Database Service on Dedicated Infrastructure System using multiple SCAN listeners.

This approach requires you to supply all of the single client access name (SCAN) virtual IP (VIP) addresses, and enables Oracle Net Services to connect to an available SCAN listener.

- Use the following template to define a Net Services alias, which is typically used to provide a convenient name for the connect descriptor:

```
alias-name = (DESCRIPTION=
  (ADDRESS_LIST=
    (ADDRESS=(PROTOCOL=tcp)(HOST=SCAN-VIP-1)(PORT=1521))
    (ADDRESS=(PROTOCOL=tcp)(HOST=SCAN-VIP-2)(PORT=1521))
    (ADDRESS=(PROTOCOL=tcp)(HOST=SCAN-VIP-3)(PORT=1521)))
  (CONNECT_DATA=(sid-or-service-entry)))
```

Where:

alias-name is the name you use to identify the alias.

SCAN-VIP-[1-3] are the IP addresses for the SCAN VIPs.

sid-or-service-entry identifies the database SID or service name using one of the following formats:

- `SID=`*sid-name*. For example: `SID=S12C1`.
- `SERVICE_NAME=`*service-name*. For example: `SERVICE_NAME=PDB1.example.yourcloud.com`.

> **Note:**
>
> By default, Oracle Net Services randomly selects one of the addresses in the address list to balance the load between the SCAN listeners.

Connecting to a Database Use a Connect Descriptor that References a Custom SCAN Name

You can set up a connect descriptor for Oracle Exadata Database Service on Dedicated Infrastructure System using a custom SCAN name.

Using this approach, you define a custom single client access name (SCAN) name in your domain name server (DNS), which resolves to the three SCAN virtual IP addresses (VIPs).

- Use the following template to define a Net Services alias that references the custom SCAN name:

  ```
  alias-name = (DESCRIPTION=
    (ADDRESS_LIST=(ADDRESS=(PROTOCOL=tcp)(HOST=scan-name)(PORT=1521)))
    (CONNECT_DATA=(sid-or-service-entry)))
  ```

  Where:

  alias-name is the name you use to identify the alias.

  scan-name is the custom SCAN name.

  sid-or-service-entry identifies the database SID or service name using one of the following formats:

  - `SID=`*sid-name*. For example: `SID=S12C1`.
  - `SERVICE_NAME=`*service-name*. For example: `SERVICE_NAME=PDB1.example.yourcloud.com`.

  Alternatively, you can use the easy connect method to specify a connect descriptor with the following format:

  ```
  scan-name:1521/sid-or-service-entry
  ```

  For example:

  ```
  exa1scan.example.com:1521/S12C1
  ```

  Or

  ```
  exa1scan.example.com:1521/PDB1.example.yourcloud.com
  ```

Connecting to a Database Using a Node Listener

To connect to an Oracle Database instance on Exadata Cloud Infrastructure with a connect descriptor that bypasses the SCAN listeners, use this procedure to route your connection directly to a node listener.

By using this method, you give up the high-availability and load-balancing provided by SCAN. However, this method may be desirable if you want to direct connections to a specific node or network interface. For example, you might want to ensure that connections from a program that performs bulk data loading use the backup network.

Using this approach, you direct your connection using the hostname or IP address of the node.

Example 5-1 Defining a Net Service Alias That Directly References the Node

```
alias-name = (DESCRIPTION=
  (CONNECT_TIMEOUT=timeout)
  (ADDRESS_LIST=(ADDRESS=(PROTOCOL=tcp)(HOST=node)(PORT=1521)))
  (CONNECT_DATA=(sid-or-service-entry)))
```

Where:

alias-name is the name you use to identify the alias.

timeout specifies a timeout period (in seconds), which enables you to terminate a connection attempt without having to wait for a TCP timeout. The (CONNECT_TIMEOUT=*timeout*) parameter is optional.

node is the hostname or IP address for the virtual machine that you want to use.

sid-or-service-entry identifies the database SID or service name using one of the following formats:

- SID=*sid-name*. For example, SID=S12C1.
- SERVICE_NAME=*service-name*. For example, SERVICE_NAME=PDB1.example.oraclecloudatcust.com.

Alternatively, you can use the easy connect method to specify a connect descriptor with the following format:

node:1521/*sid-or-service-entry*

For example:

exa1node01.example.com:1521/S12C1

Or

exa1node01.example.com:1521/PDB1.example.oraclecloudatcust.com

Manage Exadata Cloud Infrastructure

Use the provided tools to manage the Infrastructure.

- Using the Console to Provision Exadata Cloud Infrastructure
  Learn how to provision an Exadata Cloud Infrastructure system.
- Using the API to Create Infrastructure Components
- Using the API to Manage Exadata Cloud Infrastructure Instance

Using the Console to Provision Exadata Cloud Infrastructure

Learn how to provision an Exadata Cloud Infrastructure system.

- Lifecycle Management Operations

- Network Management Operations
- Management Tasks for the Oracle Cloud Infrastructure Platform
- Oracle Database License Management Tasks
- Scaling Resources within an Exadata Infrastructure Instance
  If an Exadata Cloud Infrastructure instance requires more resources, you can scale up the number of DB servers, or storage servers.

Lifecycle Management Operations

- To check the status of a Cloud Exadata infrastructure resource
- To change the infrastructure display name
- To check the status of a cloud VM cluster
- To check the status of an Exadata DB system
- To start, stop, or reboot an Exadata Cloud Infrastructure cloud VM cluster or DB system
- To terminate Exadata Cloud Infrastructure infrastructure-level resources
- Using the Console to View a List of DB Servers on an Exadata Infrastructure
  To view a list of database server hosts on an Oracle Exadata Database Service on Dedicated Infrastructure system, use this procedure.

To check the status of a Cloud Exadata infrastructure resource

> **Note:**
>
> This topic only applies to Exadata Cloud Infrastructure instances using the new Exadata Cloud Infrastructure instance resource model.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata Infrastructure** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
4. In the list of Cloud Exadata infrastructure resources, click the name of the infrastructure you're interested in and check its icon. The icon text indicates the status of the system. The following lifecycle states apply to the Cloud Exadata infrastructure
   - **Provisioning:** Reources are being reserved for the Cloud Exadata infrastructure resource. Provisioning can take several minutes. The resource is not ready to be used.
   - **Available:** The Cloud Exadata infrastructure was successfully provisioned. You can create a cloud VM cluster on the resource to complete the infrastructure provisioning.
   - **Updating:** The Cloud Exadata infrastructure is being updated. The resource goes into the updating state during management tasks. For example, when moving the resource to another compartment, or creating a cloud VM cluster on the resource.
   - **Maintenance in Progress:** A maintenance update is currently being performed on the infrastructure resource. See Maintaining an Exadata Cloud Service Instance for details on infrastructure maintenance scheduling and impacts.

- **Terminating:** The Cloud Exadata infrastructure is being deleted by the terminate action in the Console or API.
- **Terminated:**The Cloud Exadata infrastructure has been deleted and is no longer available.
- **Failed:** An error condition prevented the provisioning or continued operation of the Cloud Exadata infrastructure.

To change the infrastructure display name

> **Note:**
>
> This topic only applies to Exadata Cloud Infrastructure instances using the new Exadata Cloud Service instance resource model.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata Infrastructure** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
4. In the list of Cloud Exadata infrastructure resources, click the name of the infrastructure you're interested in
5. On rthe **Infrastructure Details** page, click **Update Display Name** .
6. In the **Update Display Name** dialog, Enter the **New display name**, and the **current display name** as instructed.
7. Click **Update Display Name**.

To check the status of a cloud VM cluster

> **Note:**
>
> This topic only applies to Exadata Cloud Infrastructure instances using the new Exadata Cloud Infrastructure instance resource model.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
4. In the list of cloud VM clusters, find the cluster you're interested in and check its icon. The icon text indicates the status of the system. The following lifecycle states apply to the cloud VM cluster:
   - **Provisioning:** Resources are being reserved for the Cloud Exadata infrastructure resource. Provisioning can take several minutes. The resource is not ready to use yet.
   - **Available:** The Cloud Exadata infrastructure was successfully provisioned. You can create a cloud VM cluster on the resource to complete the infrastructure provisioning.

- **Updating:** The Cloud Exadata infrastructure is being updated. The resource goes into the updating state during management tasks. For example, when moving the resource to another compartment, or creating a cloud VM cluster on the resource.
- **Terminating:** The Cloud Exadata infrastructure is being deleted by the terminate action in the Console or API.
- **Terminated:**The Cloud Exadata infrastructure has been deleted and is no longer available.
- **Failed:** An error condition prevented the provisioning or continued operation of the Cloud Exadata infrastructure.

To view the status of a virtual machine (database node) in the cloud VM cluster, under Resources, click **Virtual Machines** to see the list of virtual machines. In addition to the states listed for a cloud VM cluster, a virtual machine's status can be one of the following:

- **Starting:** The database node is being powered on by the start or reboot action in the Console or API.
- **Stopping:** The database node is being powered off by the stop or reboot action in the Console or API.
- **Stopped:** The database node was powered off by the stop action in the Console or API.

To check the status of an Exadata DB system

> **Note:**
>
> This topic only applies to Exadata Cloud Infrastructure instances using the DB system resource model.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.

   A list of DB systems is displayed.
3. In the list of DB systems, find the system you're interested in and check its icon. The icon text indicates the status of the system. The following lifecycle states apply to the DB system resource:
   - **Provisioning:** Resources are being reserved for the DB system, the system is booting, and the initial database is being created. Provisioning can take several minutes. The system is not ready to use yet.
   - **Available:** The DB system was successfully provisioned. A few minutes after the system enters this state, you can SSH to it and begin using it.
   - **Terminating:** The DB system is being deleted by the terminate action in the Console or API.
   - **Terminated:** The DB system has been deleted and is no longer available.
   - **Failed:** An error condition prevented the provisioning or continued operation of the DB system.

   To view the status of a database node, under Resources, click **Nodes** to see the list of nodes. In addition to the states listed for a DB system, a node's status can be one of the following:

- **Starting:** The database node is being powered on by the start or reboot action in the Console or API.
- **Stopping:** The database node is being powered off by the stop or reboot action in the Console or API.
- **Stopped:** The database node was powered off by the stop action in the Console or API.

You can also check the status of DB systems and database nodes using the ListDbSystems or ListDbNodes API operations, which return the `lifecycleState` attribute.

To start, stop, or reboot an Exadata Cloud Infrastructure cloud VM cluster or DB system

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system you want to start, stop, or reboot:

   *Cloud VM clusters (*new resource model*):* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems: Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. Under **Resources**, click **Virtual Machines** (for cloud VM clusters) or **Nodes** (for DB systems) to display the compute nodes of the cloud service instance. Click Actions icon (three dots) for a node and then click one of the following actions:
   - **Start:**;Restarts a stopped node. After the node is restarted, the **Stop** action is enabled.
   - **Stop:** Shuts down the node. After the node is powered off, the **Start** action is enabled.
   - **Reboot:** Shuts down the node, and then restarts it.

> **Note:**
>
> - For billing purposes, the **Stop** state has no effect on the resources you consume. Billing continues for virtual machines or nodes that you stop, and related resources continue to apply against any relevant quotas. You must **Terminate** a cloud VM cluster or DB system to remove its resources from billing and quotas.
> - After you restart or reboot a node, the floating IP address might take several minutes to be updated and display in the Console.

To terminate Exadata Cloud Infrastructure infrastructure-level resources

This topic describes how to terminate a Cloud Exadata infrastructure, cloud VM cluster, or DB system resource in anExadata Cloud Infrastructure instance.

> **Note:**
>
> The database data is local to the cloud VM cluster or DB system hosting it and is lost when the system is terminated. Oracle recommends that you back up any data in the cloud VM cluster or DB system before terminating it.
>
> Terminating an Exadata Cloud Infrastructure resource permanently deletes it and any databases running on it. The data is deleted in compliance with the NIST SP-800-88r1 standard and implemented as an Exadata crypto erase using the hardware Instant Secure Erase (ISE) feature of the Exadata storage devices, as documented in the Exadata Database Machine Documentation at Securely Erasing Exadata Database Machine

1. Open the navigation menu. Click Oracle Database, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the Cloud Exadata infrastructure, cloud VM cluster or DB system you want to move:

   *Cloud Exadata infrastructure (*new resource model*):* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata Infrastructure**. In the list of infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.

   *Cloud VM clusters (*new resource model*):* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems: Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. For cloud VM clusters and DB systems, click **More Actions**, then **Terminate** on the resource details page. For Cloud Exadata infrastructure resources, click **Terminate** on the resource details page.
   Confirm when prompted.

   The resource's icon indicates Terminating.

   > **Note:**
   >
   > If you are terminating a Cloud Exadata infrastructure resource that contains a cloud VM cluster, you must check the box labelled **Also delete the VM cluster associated with this infrastructure** to confirm that you intend to delete the VM cluster.

At this point, you cannot connect to the system and any open connections are terminated.

Using the Console to View a List of DB Servers on an Exadata Infrastructure

To view a list of database server hosts on an Oracle Exadata Database Service on Dedicated Infrastructure system, use this procedure.

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata Infrastructure**.
3. In the list of Exadata Infrastructures, click the display name of the infrastructure you wish to view details.
4. Under **Resources**, click **DB Servers**.
5. In the list of DB Servers, click the name of the DB Server that you wish to view details.

   DB Server lists VMs from each cluster hosted on them along with resources allocated to them.

Network Management Operations

- To edit the network security groups (NSGs) for your client or backup network

To edit the network security groups (NSGs) for your client or backup network

Your client and backup networks can each use up to five network security groups (NSGs). Note that if you choose a subnet with a security list, the security rules for the cloud VM cluster or DB system will be a union of the rules in the security list and the NSGs. For more information, see Network Security Groups and Network Setup for Exadata Cloud Infrastructure Instances.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system you want to manage:
   Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems: Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. In the **Network** details, click the **Edit** link to the right of the **Client Network Security Groups** or **Backup Network Security Groups** field.
5. In the **Edit Network Security Groups** dialog, click **+ Another Network Security Group** to add an NSG to the network.

   To change an assigned NSG, click the drop-down menu displaying the NSG name, then select a different NSG.

   To remove an NSG from the network, click the **X**;icon to the right of the displayed NSG name.
6. Click **Save**.

Management Tasks for the Oracle Cloud Infrastructure Platform

- To view a work request for your Exadata Cloud Infrastructure resources
- To move an Exadata Cloud Infrastructure resource to another compartment
- To manage tags for your Exadata Cloud Infrastructure resources

- Managing Infrastructure Maintenance Contacts
  Learn to manage your Exadata infrastructure maintenance contacts.

To view a work request for your Exadata Cloud Infrastructure resources

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.

   A list of DB systems is displayed.
3. Find the Cloud Exadata infrastructure, cloud VM cluster, DB system or database resource you're interested in, and click the name.
4. In the **Resources** section, click **Work Requests**. The status of all work requests appears on the page.
5. To see the log messages, error messages, and resources that are associated with a specific work request, click the operation name. Then, select an option in the **More information** section.

   For associated resources, you can click the Actions icon (three dots) next to a resource to copy the resource's OCID.

Related Topics

- Work Requests

To move an Exadata Cloud Infrastructure resource to another compartment

> **Note:**
>
> - To move resources between compartments, resource users must have sufficient access permissions on the compartment that the resource is being moved to, as well as the current compartment. For more information about permissions for Database resources, see *Details for the Database Service*.
> - If your Exadata Cloud Infrastructure instance is in a security zone, the destination compartment must also be in a security zone. See the *Security Zone Policies* topic for a full list of policies that affect Database service resources.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the Cloud Exadata infrastructure, cloud VM cluster or DB system you want to move:

   Cloud Exadata infrastructure (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata Infrastructure**. In the list of infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.

   Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

DB systems: Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

4. Click **Move Resource**.
5. Select the new compartment.
6. Click **Move Resource**.

Related Topics

- Details for the Database Service
- Security Zone Policies
- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

To manage tags for your Exadata Cloud Infrastructure resources

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Find the Cloud Exadata infrastructure, cloud VM cluster, DB system or database resource you're interested in, and click the name.
4. Click the **Tags** tab to view or edit the existing tags. Or click **More Actions** and then **Apply Tags** to add new ones.

Related Topics

- Resource Tags

Managing Infrastructure Maintenance Contacts

Learn to manage your Exadata infrastructure maintenance contacts.

- To manage maintenance contacts in an Exadata Cloud Infrastructure
  Manage contacts for Exadata infrastructure maintenance notifications using the Console.

To manage maintenance contacts in an Exadata Cloud Infrastructure

Manage contacts for Exadata infrastructure maintenance notifications using the Console.

To prevent an Exadata infrastructure administrator from being overwhelmed by system update notifications, you can specify up to 10 email addresses of people to whom maintenance notifications are sent.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. In the **Oracle Exadata Database Service on Dedicated Infrastructure** section, click **Exadata Infrastructure** to display a list of Exadata infrastructures in the default compartment. You can select a different compartment from the **Compartment** drop-down located in the **List Scope** section.
3. In the list of Exadata infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.

4. In the **Maintenance** section, click **Manage** in the **Customer Contacts** field to display the Manage Contacts dialog.
5. Click the **Add Contacts** button to display a field in which to enter a valid email address. You can have up to 10 maintenance contacts for each Exadata infrastructure.
6. To edit an email address, in the Manage Contacts dialog, select the box preceding the email address you want to edit and click the **Edit** button.
7. To remove an email address from the list, in the Manage Contacts dialog, select the box preceding the email address you want to remove and click the **Remove** button.

Oracle Database License Management Tasks

- To manage your BYOL database licenses
  f you want to control the number of database licenses that you run at any given time, you can scale up or down the number of OCPUs on the instance. These additional licenses are metered separately.
- To change the license type of a cloud VM cluster or DB system

To manage your BYOL database licenses

f you want to control the number of database licenses that you run at any given time, you can scale up or down the number of OCPUs on the instance. These additional licenses are metered separately.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system you want to scale:

   Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems: Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. Click **Scale VM Cluster** (for cloud VM clusters) or **Scale CPU;Cores** (for DB systems) and then specify a new number of CPU cores. The text below the field indicates the acceptable values, based on the shape used when the DB system was launched.
5. Click **Update**.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

To change the license type of a cloud VM cluster or DB system

> **Note:**
>
> Updating the license type is not supported for systems running on the X6 shape. The feature is supported for X7 and higher shapes.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system you want to manage:
   a. *Cloud VM clusters (new resource model)* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

      DB systems Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. On the resource details page, click **Update License Type**.

   The dialog displays the options with your current license type selected.
5. Select the new license type.
6. Click **Save**.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

Scaling Resources within an Exadata Infrastructure Instance

If an Exadata Cloud Infrastructure instance requires more resources, you can scale up the number of DB servers, or storage servers.

There are three distinct scaling paths:

- Scaling an X8M or X9M based Multi-VM enabled Infrastructure by adding DB servers and Storage servers to an existing infrastructure. See *Add Resources To a Multi-VM Enabled Infrastructure*
- Scaling an X8M or X9M based that is NOT a Multi-VM enabled Infrastructure. See *Scaling Exadata X8M and X9M Compute and Storage*
- Scaling an X6, X7 and X8 Exadata infrastructure (fixed shape). See *Scaling X6, X7 and X8 Exadata DB System Configurations*
- Add Resources to a Multi-VM enabled Infrastructure
  Add DB servers or storage servers to an existing Multi-VM enabled Infrastructure
- Remove DB Servers from a Multi-VM enabled Infrastructure
  Remove DB servers from an existing Multi-VM enabled Infrastructure

- Scaling CPU cores within an Exadata Cloud Infrastructure instance
  If an Exadata Cloud Infrastructure instance requires more compute node processing power, you can scale up the number of enabled CPU cores symmetrically across all the nodes in the system as follows:
- Scaling X6, X7 and X8 Exadata DB System Configurations
  Scaling an Exadata X6, X7, or X8 Exadata Cloud Infrastructure instance by moving to a shape with more capacity enables you meet the needs of your growing workload.

Related Topics

- Add Resources to a Multi-VM enabled Infrastructure
  Add DB servers or storage servers to an existing Multi-VM enabled Infrastructure

Add Resources to a Multi-VM enabled Infrastructure

Add DB servers or storage servers to an existing Multi-VM enabled Infrastructure

You can scale an X8M or X9M Exadata Multi-VM enabled Infrastructure instance in the Console on the cloud Exadata infrastructure details page. After adding additional database or storage servers to your cloud Exadata infrastructure resource, you must add the increased capacity to the associated cloud VM cluster to utilize the newly-provisioned CPU or storage resources.

Adding DB servers, or Storage servers do not require any database downtime.

> **Note:**
>
> - Neither the Exadata X8M nor the X9M shapes support removing storage or database servers from an existing cloud infrastructure instance.

To add DB Servers to Multi-VM enabled Infrastructure

1. Navigate to **Oracle Cloud** menu and click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Select **Exadata Infrasrtucture** under **Oracle Exadata Database Service on Dedicated Infrastructure**
3. Select the desired Infrastructure in the desired compartment.
4. On the **Infrastructure Details** page click **Scale Infrastructure**.
5. In the **Scale Infrastructure** page, set the **Additional DB servers** to a value so that the total of DB servers is 8 or less.
6. Click **Scale Infrastructure**.

To add storage servers to Multi-VM enabled Infrastructure

1. Navigate to **Oracle Cloud** menu and click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Select **Exadata Infrastructure** under **Oracle Exadata Database Service on Dedicated Infrastructure**
3. Select the desired Infrastructure in the desired compartment.
4. On the **Infrastructure Details** page click **Scale Infrastructure**.

5. In the **Scale Infrastructure** page, set the **Additional storage servers** to a value so that the total of storage servers is 12 or less.

> **Note:**
>
> - This operation adds the storage servers to the infrastructure, but the storage capacity must be made available for VM Cluster consumption.
> - You will be able to scale down a storage server if the server has not been used to expand Exadata infrastructure storage.

6. Click **Scale Infrastructure**.
7. On the **Infrastructure Details** page, a banner directs you to **Add Storage Cappacity** to make the the storage capacity available for VM Cluster consumption.
8. Click **Add Storage Capacity**.
9. In the **Add Storage Capacity** page, click **Add Storage Capacity**.

Remove DB Servers from a Multi-VM enabled Infrastructure

Remove DB servers from an existing Multi-VM enabled Infrastructure

- Database servers will be removed if there are no VMs running on them.

> **Note:**
>
> You will not be able to choose the DB Server to remove. This functionality will automatically remove Database Servers in which there are no VMs.

For more information about removing a VM, see *Terminate a VM from a VM Cluster*.

Related Topics

- Terminate a VM from a VM Cluster
  To remove a virtual machine from a provisioned cluster, use this procedure.

Scaling CPU cores within an Exadata Cloud Infrastructure instance

If an Exadata Cloud Infrastructure instance requires more compute node processing power, you can scale up the number of enabled CPU cores symmetrically across all the nodes in the system as follows:

The options for each of the shapes are:

X8M or X9M flexible infrastructure systems: You can scale CPU cores in multiples of the number of database servers currently provisioned for the cloud VM cluster. For example, if you have 6 database servers provisioned, you can add CPU cores in multiples of 6. At the time of provisioning, X8M systems have as few as 2 database servers or up to 32 database servers.. For more information on adding compute and storage resources to an X8M or X9M system, see *Scaling Exadata X8M and X9M Compute and Storage* .

Non-X8M fixed-shape systems: For a base system or an X7 or X8 quarter rack, you can scale in multiples of 2 across the 2 database compute nodes. For an X7 or X8 half rack, you can scale in multiples of 4 across the 4 database compute nodes. For an X7 or X8 full rack, you can scale in multiples of 8 across the 8 database compute nodes.

For a non-metered service instances, you can temporarily modify the compute node processing power (bursting) or add compute node processing power on a more permanent basis. For a metered service instance, you can simply modify the number of enabled CPU cores.

You can provision an Exadata Cloud Infrastructure instance with zero CPU cores, or scale the service instance down to zero cores after you provision it. With zero cores, you are billed only for the infrastructure until you scale up the system. For detailed information about pricing, see Exadata Cloud Service Pricing.

> **Note:**
>
> OCPU scaling activities are done online with no downtime.

For information on CPU cores per configuration, see *Exadata Shape Configurations*. To learn how to scale a system, see *To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster or DB system* .

- Scaling Exadata X8M and X9M Compute and Storage
  The flexible X8M and X9M system model is designed to be easily scaled in place, with no need to migrate the database using a backup or Data Guard

Related Topics

- To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster or DB system

Scaling Exadata X8M and X9M Compute and Storage

The flexible X8M and X9M system model is designed to be easily scaled in place, with no need to migrate the database using a backup or Data Guard

You can scale an X8M or X9M Exadata cloud infrastructure instance in the Console on the cloud Exadata infrastructure details page. After adding additional database or storage servers to your cloud Exadata infrastructure resource, you must add the increased capacity to the associated cloud VM cluster to utilize the newly-provisioned CPU or storage resources. After adding additional database servers to a VM cluster, you can then allocate the new CPU cores as described in see *To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster or DB system*. After adding additional storage servers to your VM cluster, you do not need to take any further action to utilize the new storage.

> **Note:**
>
> - Neither the Exadata X8M nor the X9M shapes support removing storage or database servers from an existing cloud infrastructure instance.

- To add compute and storage resources to a flexible cloud Exadata infrastructure resource
  This task describes how to use the Oracle Cloud Infrastructure Console to scale a flexible cloud Exadata infrastructure resource.
- To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster or DB system

Related Topics

- To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster or DB system

- Overview of X8M and X9M Scalable Exadata Infrastructure
  Oracle Cloud Infrastructure scalable X8M and X9M Exadata cloud infrastructure model allows you to add additional database and storage servers after provisioning and create a system that matches your capacity needs.
- To add compute and storage resources to a flexible cloud Exadata infrastructure resource
  This task describes how to use the Oracle Cloud Infrastructure Console to scale a flexible cloud Exadata infrastructure resource.
- To add database server or storage server capacity to a cloud VM cluster
  This topic describes how to use the Oracle Cloud Infrastructure (OCI) Console to add the new capacity to your cloud VM cluster.

To add compute and storage resources to a flexible cloud Exadata infrastructure resource

This task describes how to use the Oracle Cloud Infrastructure Console to scale a flexible cloud Exadata infrastructure resource.

Currently, only Exadata X8M and X9M systems in Oracle Cloud Infrastructure have the ability to add database (compute) and storage servers to an existing service instance.

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata Infrastructure**.
3. In the list of cloud Exadata infrastructure resources, click the name of the resource you want to scale.
4. Click **Scale Infrastructure**.
5. Add either Database servers or Storage Servers by selecting the proper radio button
   a. *Adding database servers:* To add compute servers to the infrastructure resource, select the **Database Servers** radio button, then enter the number of servers you want to add in the **Database servers** field.
   b. *Adding storage servers:* To add storage servers to the infrastructure resource, select the **Storage Servers** radio button, then enter the number of servers you want to add in the **Storage servers** field.
6. Click **Scale**.

> **Note:**
>
> After scaling your infrastructure, you must add the new capacity to the cloud VM cluster before you can use the additional CPU and storage resources in the Exadata Cloud Infrastructure instance.

Related Topics

- To add database server or storage server capacity to a cloud VM cluster
  This topic describes how to use the Oracle Cloud Infrastructure (OCI) Console to add the new capacity to your cloud VM cluster.

To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster or DB system

> **Note:**
>
> For information on adding additional database (compute) and storage servers to **X8M or X9M cloud VM clusters**, see To add compute and storage resources to a flexible cloud Exadata infrastructure resource and To add database server or storage server capacity to a cloud VM cluster. Adding additional database servers to your X8M cloud VM cluster will increase the number of CPU cores available for scaling.

If an Exadata Cloud Infrastructure instance requires more compute node processing power, you can scale up (increase) the number of enabled CPU cores (OCPUs) in the instance.

You can also scale a cloud VM cluster or DB system (except for X6 systems) down to zero (0) CPU cores to temporarily stop the system and be charged only for the hardware infrastructure. For more information about scaling down, see Scaling Options. Oracle recommends that if you are not scaling down to a stopped system (0 cores), that you scale to at least 2 cores per node.

> **Note:**
>
> The minimum cores is 1 for X8 and older, minimum for X8M and newer is 2.

CPU cores must be scaled symmetrically across all nodes in the cloud VM cluster or DB system. Use multiples of two CPUs per database server. For example, if you have two database servers, a minimum of 2 CPU cores per server or a total of 4 CPU cores. The total number of CPU cores must not exceed the maximum limit for that shape and/or resources.

> **Tip:**
>
> OCPU scaling activities are done online with no downtime.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system you want to scale:
   Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems: Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. Click **Scale VM Cluster** (for cloud VM clusters) or **Scale CPU;Cores** (for DB systems) and then specify a new number of CPU cores. The text below the field indicates the acceptable values, based on the shape used when the DB system was launched.
5. Click **Update**.

> **Note:**
>
> If you scale an scale a cloud VM cluster or DB system (except for X6 systems) down to zero (0) CPU cores, the floating IP address of the nodes might take several minutes to be updated and display in the Console.

Scaling X6, X7 and X8 Exadata DB System Configurations

Scaling an Exadata X6, X7, or X8 Exadata Cloud Infrastructure instance by moving to a shape with more capacity enables you meet the needs of your growing workload.

This is useful when a database deployment requires:

- Processing power that is beyond the capacity of the current system configuration.
- Storage capacity that is beyond the capacity of the current system configuration.
- A performance boost that can be delivered by increasing the number of available compute nodes.
- A performance boost that can be delivered by increasing the number of available Exadata Storage Servers.

You can move you workloads to a larger fixed shape (X7 and X8 hardware shapes), or move to the flexible X8M shape that allows for easy expansion of compute and storage resources as your workloads grow.

To assist with moving your database deployments between Exadata Cloud Infrastructuree instances, you can restore a backup to a different service instance that has more capacity, or create a Data Guard association for your database in a service instance with more capacity, and then perform a switchover so that your new standby database assumes the primary role. To start the process, contact Oracle and request a service limit increase so that you can provision the larger service instance needed by your database.

Using the API to Create Infrastructure Components

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.

Use these API operations to create Exadata Cloud Infrastructure components.

APIs for the New Exadata Cloud Infrastructure Resource Model

The new Exadata resource model is compatible with all offered Exadata shape families (X7, X8, and X8M). See The Exadata Cloud Infrastructure Resource Model for more information.

> **Tip:**
>
> As of November 15, 2021 new Exadata Cloud Infrastructure instances may only be provisioned using the new resource model.

Cloud Exadata infrastructure resource:

- GetCloudExadataInfrastructure
- CreateCloudExadataInfrastructure

- ListCloudExadataInfrastructures

Cloud VM cluster resource:

- GetCloudVmCluster
- CreateCloudVmCluster
- ListCloudVmClusters
- ListCloudVmClusters
- CreateCloudVmCluster
- GetCloudVmClusterIormConfig
- UpdateCloudVmClusterIormConfig

Databases

- GetDatabase
- ListDatabases

Shapes and Database Versions

- ListDbSystemShapes
- ListDbVersions

Database Homes

- CreateDbHome
- GetDbHome
- ListDbHomes

NOT_SUPPORTED

> **Note:**
>
> The DB system APIs are deprecated forExadata Cloud Infrastructure. Oracle recommends converting existing Exadata DB systems to the new resource model as soon as possible. Converting to the new resource model does not involve system downtime. Learn more.

- GetDbSystem
- LaunchDbSystem
- ListDbSystems

Using the API to Manage Exadata Cloud Infrastructure Instance

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.

Use these API operations to manage Exadata Cloud Infrastructure instance components.

Cloud Exadata infrastructure resource (new resource model):

- ListCloudExadataInfrastructures
- GetCloudExadataInfrastructure
- ChangeCloudExadataInfrastructureCompartment
- UpdateCloudExadataInfrastructure
- DeleteCloudExadataInfrastructure

Cloud VM cluster (new resource model)

- ListCloudVmClusters
- GetCloudVmCluster
- ChangeCloudVmClusterCompartment
- UpdateCloudVmCluster
- DeleteCloudVmCluster

DB systems (old resource model):

- ListDbSystems
- GetDbSystem
- ChangeDbSystemCompartment
- UpdateDbSystem
- TerminateDbSystem

Virtual machines nodes (all Exadata Cloud Infrastructure instances):

- DbNodeAction
- ListDbNodes
- GetDbNode

Cloud Infrastructure Maintenance Updates

Oracle performs the updates to all of the Oracle-managed infrastructure components on Exadata Cloud Infrastructure.

You may manage contacts who are notified regarding infrastructure maintenance, set a maintenance window to determine the time your quarterly infrastructure maintenance will begin, and also view scheduled maintenance runs and the maintenance history of your Exadata Cloud Infrastructure in the Oracle Cloud Infrastructure Console. For details regarding the infrastructure maintenance process and configuring the maintenance controls refer to the following:

- About Oracle-managed Exadata Cloud Infrastructure Maintenance
  Oracle performs patches and updates to all of the Oracle-managed system components on Exadata Cloud Infrastructure.
- Overview of the Quarterly Infrastructure Maintenance Process
  By default, infrastructure maintenance updates the Exadata database server hosts in a rolling fashion, followed by updating the storage servers.
- Overview of Monthly Security Maintenance
  Security maintenance, performed alongside the quarterly maintenance, is executed in months when important security updates are needed and includes fixes for vulnerabilities across all CVSS scores.

- Understanding Monthly and Quarterly Maintenance in the Same Month
- Using the Console to Configure Oracle-Managed Infrastructure Updates
  Software updates are scheduled quarterly and monthly. You can use the the console to schedule and plan for them.
- Monitor Infrastructure Maintenance Using Lifecycle State Information
  The lifecycle state of your Exadata Infrastructure resource enables you to monitor when the maintenance of your infrastructure resource begins and ends.
- Receive Notifications about Your Infrastructure Maintenance Updates
  There are two ways to receive notifications. One is through email to infrastructure maintenance contacts and the other one is to subscribe to the maintenance events and get notified.
- Managing Infrastructure Maintenance Contacts
  Learn to manage your Exadata infrastructure maintenance contacts.
- Using the API to Manage Exadata Cloud Infrastructure Maintenance Controls
  Use these API operations to manage Exadata Cloud Infrastructure maintenance controls and resources.

About Oracle-managed Exadata Cloud Infrastructure Maintenance

Oracle performs patches and updates to all of the Oracle-managed system components on Exadata Cloud Infrastructure.

Oracle patches and updates include the physical database server hosts, Exadata Storage Servers, Network Fabric Switches, management switch, power distribution units (PDUs), integrated lights-out management (ILOM) interfaces, and Control Plane Servers. This is referred to as infrastructure maintenance.

The frequency of the updates depends on the region type, as follows:

- Commercial regions: Oracle performs full quarterly infrastructure updates and monthly security infrastructure updates.
- Government regions: Oracle performs monthly full infrastructure maintenance updates.

In all but rare exceptional circumstances, you receive advance communication about these updates to help you plan for them. If there are corresponding recommended updates for your VM cluster virtual machines (VMs), then Oracle provides notifications about them.

Wherever possible, scheduled updates are performed in a manner that preserves service availability throughout the update process. However, there can be some noticeable impact on performance and throughput while individual system components are unavailable during the update process.

For example, database server patching typically requires a reboot. In such cases, wherever possible, the database servers are restarted in a rolling manner, one at a time, to ensure that the service remains available throughout the process. However, each database server is unavailable for a short time while it restarts, and the overall service capacity diminishes accordingly. If your applications cannot tolerate the restarts, then take mitigating action as needed. For example, shut down an application while database server patching occurs.

> **Note:**
>
> Customers using Exadata Database on Dedicated Infrastructure in Oracle Cloud Infrastructure (OCI) US Government (OC2) and US Department of Defense (OC3) regions can use the OCI console to reschedule monthly and quarterly patching events.
>
> At this time specifying a maintenance schedule, all so known as “Setting Patch Management Schedule for Exadata Cloud Infrastructure”, is still not available in the OCI US Government (OC2) and US DOD (OC3) realms for Exadata patch management. For more information on Exadata Database on Dedicated Infrastructure on Patch Management Rescheduling can be found here.

Overview of the Quarterly Infrastructure Maintenance Process

By default, infrastructure maintenance updates the Exadata database server hosts in a rolling fashion, followed by updating the storage servers.

You can also choose non-rolling maintenance to update database and storage servers. The non-rolling maintenance method first updates your storage servers at the same time, then your database servers at the same time. Although non-rolling maintenance minimizes maintenance time, it incurs full system downtime while the storage servers and database servers are being updated.

Rolling infrastructure maintenance begins with the Exadata database server hosts. For the rolling maintenance method, database servers are updated one at a time. Each of the database server host's VMs is shut down, the host is updated, restarted, and then the VMs are started, while other database servers remain operational. This rolling maintenance impacts older applications not written to handle a rolling instance outage. This process continues until all servers are updated.

After database server maintenance is complete, storage server maintenance begins. For the rolling maintenance method, storage servers are updated one at a time and do not impact VM cluster VM's availability. However, the rolling storage server maintenance can result in reduced IO performance as storage servers are taken offline (reducing available IO capacity) and resynced when brought back online (small overhead on database servers). Properly sizing the database and storage infrastructure to accommodate increased work distributed to database and storage servers not under maintenance will minimize (or eliminate) any performance impact.

> **Note:**
>
> While databases are expected to be available during the rolling maintenance process, the automated maintenance verifies Oracle Clusterware is running but does not verify that all database services and pluggable databases (PDBs) are available after a server is brought back online. The availability of database services and PDBs after maintenance can depend on the application service definition. For example, a database service, configured with certain preferred and available nodes, may be relocated during the maintenance and wouldn't automatically be relocated back to its original node after the maintenance completes. Oracle recommends reviewing the documentation on *Achieving Continuous Availability for Your Applications on Exadata Cloud Systems* to reduce the potential for impact to your applications. By following the documentation's guidelines, the impact of infrastructure maintenance will be only minor service degradation as database servers are sequentially updated.

Oracle recommends that you follow the *Maximum Availability Architecture (MAA)* best practices and use Data Guard to ensure the highest availability for your critical applications. For databases with Data Guard enabled, Oracle recommends that you separate the maintenance windows for the infrastructure instances running the primary and standby databases. You may also perform a switchover prior to the maintenance operations for the infrastructure instance hosting the primary database. This allows you to avoid any impact on your primary database during infrastructure maintenance.

Prechecks are performed on the Exadata Cloud Infrastructure components prior to the start of the maintenance window. The goal of the prechecks is to identify issues that may prevent the infrastructure maintenance from succeeding. The Exadata infrastructure and all components remain online during the prechecks. An initial precheck is run approximately 5 days prior to the maintenance start and another precheck is run approximately 24 hours prior to maintenance start. If the prechecks identify an issue that requires rescheduling the maintenance notification is sent to the maintenance contacts.

> **Note:**
>
> Do not perform major maintenance operations on your databases or applications during the patching window, as these operations could be impacted by the infrastructure maintenance operations

- Time Estimates for Quarterly Maintenance Windows
  The time taken to update infrastructure components varies depending on the number of database servers and storage servers in the Exadata infrastructure, the maintenance method and whether custom action has been enabled.

Related Topics

- Achieving Continuous Availability For Your Applications
- Maximum Availability Architecture (MAA) Best Practices

Time Estimates for Quarterly Maintenance Windows

The time taken to update infrastructure components varies depending on the number of database servers and storage servers in the Exadata infrastructure, the maintenance method and whether custom action has been enabled.

The approximate times provided are estimates. Time for custom action, if configured, is not included in the estimates. Database server maintenance time may vary depending on the time required to shutdown each VM before the update and then start each VM and associated resources after the update of each node before proceeding to the next node. The storage server maintenance time will vary depending on the time required for the ASM rebalance, which is not included in the estimates below. If issues are encountered during maintenance this may also delay completion beyond the approximate time listed. In such a situation, if Oracle cloud operations determine resolution would extend beyond the expected window, they will send a notification and may reschedule the maintenance.

> **Note:**
>
> The timeframes mentioned below can change if Oracle cloud operations determine that additional maintenance work is needed. If additional time is necessary, Oracle will send a customer notification in advance to inform customers that additional time will be required for the next quarterly maintenance window.

Table 5-1 Approximate Times for Exadata Infrastructure Maintenance

Exadata Shape Configuration	Rolling Patching Method (Approximate Time)	Non-Rolling Patching Method (Approximate Time)
Quarter rack	5-6 hours	4-7 hours
Half rack	10 hours	4-7 hours
Full rack	20 hours	4-7 hours
Flexible shapes (X8M and higher)	1.5 hours per compute node + 1 hour per storage node	4-7 hours

Overview of Monthly Security Maintenance

Security maintenance, performed alongside the quarterly maintenance, is executed in months when important security updates are needed and includes fixes for vulnerabilities across all CVSS scores.

> **Note:**
>
> For more information about the CVE release matrix, see Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1).
> To view the CVE release matrix specific to an Exadata Infrastructure version, click the Exadata version, for example, Exadata 23. Version-specific CVE release matrices are listed in the Notes column of the table.

Security maintenance, when needed, is scheduled to be applied during a 21-day window that begins between the 18th-21st of each month and will run till the 9th-12th of the next month. Customers will receive notification of the proposed schedule at least 7 days before the start of the monthly maintenance window and can reschedule monthly maintenance to another date in the window if desired. The monthly security maintenance process updates database servers to fix critical security vulnerabilities and critical product issues. Monthly maintenance also updates storage servers to an Exadata Storage Software image that resolves known security vulnerabilities and product issues.

Updates to database servers are applied online via Ksplice technology, and have no impact to workloads running on the compute (database) servers, as database server security updates are applied online to the host server while your VM and all processes within the VM, including databases, remain up and running. Servers and VMs are not restarted. Updates to storage servers are applied in a rolling fashion. As with quarterly maintenance, the impact of rebooting storage servers should be minimal to applications.

While updating your services infrastructure, some operations including memory, and storage scaling, operating system and Grid Infrastructure patching (including prechecks), and elastic expansion of compute and storage servers may be blocked.

> **Note:**
>
> Only VM startup and shutdown operations are supported during monthly infrastructure maintenance.

Please plan to defer these operations until after the updates are complete. Application of security updates takes about 15 minutes per DB server host, plus 60 minutes per storage server depending on the I/O activity. If you attempt an affected operation, the console will notify you of the ongoing security updates. No software is updated in the guest VMs.

Related Topics

- https://support.oracle.com/rs?type=doc&id=888828.1

Understanding Monthly and Quarterly Maintenance in the Same Month

Special considerations are made when both quarterly and monthly security maintenance are scheduled to run in the same month. Quarterly maintenance will reapply any security fixes already applied by security maintenance, and neither quarterly nor monthly maintenance will apply a storage server update if the existing storage server version is the same or newer than the version contained in the update.

- The contents of the updates applied during quarterly maintenance are determined at the start of the maintenance quarter and use the latest Exadata release from the month prior to the start of the maintenance quarter. If any additional security fixes are available at that time, those updates are included in the quarterly maintenance. That image is then used throughout the quarter. For example, the January release is used for quarterly maintenance in Feb, March, and April.
- When quarterly maintenance is applied it is possible there are security updates previously installed on the database servers are not included in the quarterly maintenance release to be applied. In that case, the automation will apply the same security fixes to new release installed by the quarterly maintenance so there will not be any regression in security fixes. If the current image on the storage server is the same or newer than that to be applied by the quarterly or monthly security maintenance, that maintenance will be skipped for the storage servers.

If quarterly maintenance is scheduled within 24 hours of the time the monthly is scheduled, the scheduled monthly maintenance will be skipped and the monthly update will instead be applied immediately following the quarterly maintenance.

- When scheduled at the same time, the monthly update is executed immediately following the completion of the quarterly maintenance.

- If monthly maintenance is scheduled to begin 0-24 hours ahead of the quarterly maintenance, then the monthly maintenance will not execute as scheduled, but instead, wait and be executed immediately following the quarterly maintenance. If the quarterly maintenance is subsequently rescheduled, then the monthly security maintenance will begin immediately. Oracle, therefore, recommends scheduling quarterly and monthly maintenance at the same time. As a result, if you reschedule the quarterly at the last moment, the monthly maintenance will run at the scheduled time instead of immediately upon editing the schedule. You can also reschedule the monthly security maintenance when rescheduling the quarterly maintenance as long as you keep the monthly within the current maintenance window. Monthly maintenance can be rescheduled to another time in the maintenance window, but cannot be skipped.

Monthly Security Maintenance before Quarterly Maintenance

- To apply security maintenance before quarterly maintenance, reschedule the monthly security maintenance to occur more than 24 hours prior to the quarterly maintenance. The security maintenance will online apply security patches to the database servers with no impact to applications, and apply an update to the storage servers with minimal to no impact (may be slight performance degradation) on applications. The quarterly maintenance will follow as scheduled, and will perform rolling maintenance on the database servers, which will impact applications not written to handle a rolling reboot. As part of the quarterly maintenance, it will apply the same security updates to the database server that are already installed on the system (no security regression).
- If you are concerned about getting the latest security updates applied, schedule the monthly security maintenance to run after the new monthly maintenance window opens (usually on the 21st of the month).
- The impact of the monthly security maintenance rebooting the storage servers should be minimal, so impact to the applications during this month will only be due to the restart of the database servers during the quarterly maintenance. However, if you must coordinate a maintenance window with your end users for the security maintenance, this will require two maintenance windows.

Quarterly Maintenance before Monthly Security Maintenance

- To run the quarterly maintenance before the monthly security maintenance, reschedule the security maintenance to run no earlier than 24 hours before the quarterly maintenance is scheduled to start. The security maintenance will be deferred until the quarterly maintenance is completed. The quarterly maintenance will perform rolling maintenance on the database servers, which will impact applications not written to handle a rolling reboot. The quarterly maintenance may or may not skip the storage server patching. That depends on if it is newer or older than the release currently installed. In most cases, the version installed should be newer than the version associated with the quarterly maintenance. Exceptions to this rule may occur if it is the first month of a maintenance quarter, or you skipped the security maintenance in one or more prior months. The security maintenance will run either immediately after the quarterly maintenance is completed, or when scheduled, whichever is later. It will apply online updates to the database servers (no application impact) and will likely update the storage servers in a rolling manner. In some corner cases. the quarterly maintenance may contain the same storage server release as the security maintenance and the security maintenance storage server updates will be skipped.
- The impact to end users of running the quarterly maintenance before the security maintenance should be roughly the same as running the security maintenance first. The quarterly maintenance will be a disruptive event, but the security maintenance rebooting the storage servers should cause minimal disruption, and the security maintenance is applied to the database servers online. However, if you must coordinate a maintenance

window with your end users for the security maintenance, this will require two maintenance windows. You can schedule those two maintenance windows to run back-to-back, to appear as single maintenance window to end users. To do this, reschedule the security maintenance to start at the same time (or up to 24 hours prior) as the quarterly maintenance. The security maintenance will be deferred until the quarterly maintenance is completed. Assuming you have been regularly applying monthly security maintenance, the storage servers will be skipped by the quarterly maintenance and will be updated by the security maintenance immediately upon the completion of the quarterly maintenance.

Minimizing Maintenance Windows

- To minimize the number of maintenance windows (you have to negotiate those with end users), schedule the quarterly maintenance and monthly maintenance at the same time. The security maintenance will be blocked. The quarterly maintenance will update the database servers in a rolling manner and will most likely skip the storage server. The security maintenance will follow up immediately and update the database servers online and the storage servers in a rolling manner. The result is a single database and storage server restart in a single maintenance window.
- There are two exceptions to this. 1. If the quarterly and monthly maintenance contain the same storage server release, the quarterly maintenance will apply the storage server update, and the security maintenance will be skipped. From your perspective, this is still a single rolling reboot in a single maintenance window. 2. The currently installed release on the storage servers is older than that contained in the quarterly maintenance, which in turn is older than that in the security maintenance. That would cause the quarterly maintenance to update the storage, and then the security maintenance to do it as well. This can only happen if you skipped a prior month's security maintenance, because it requires the current image to be at least 2 months out of date. In such a scenario, you may want to schedule the security maintenance first and then the quarterly maintenance. This would result in one storage server reboot, but two distinct maintenance windows — the first for the security maintenance, and then later the quarterly maintenance.
- To minimize the impact to your end users, always apply the monthly security updates, and in months where both are scheduled, schedule them at the same time.

> **Note:**
>
> If the Exadata Infrastructure is provisioned before Oracle schedules the security maintenance, then it will be eligible for security maintenance.
> Any time before the scheduled monthly Exadata Infrastructure maintenance, you can reschedule it.

Using the Console to Configure Oracle-Managed Infrastructure Updates

Software updates are scheduled quarterly and monthly. You can use the the console to schedule and plan for them.

Full Exadata Cloud Infrastructure software updates are scheduled on a quarterly basis for commercial regions, and monthly for government regions. In addition, important security updates are scheduled monthly. While you cannot opt-out of these infrastructure updates, Oracle alerts you in advance through the Cloud Notification Portal and allows scheduling flexibility to help you plan for them.

For quarterly infrastructure maintenance, you can set a maintenance window to determine when the maintenance will begin. You can also edit the maintenance method, enable custom

action, view the scheduled maintenance runs and the maintenance history, and manage maintenance contacts in the in the Exadata Infrastructure Details page of the Oracle Cloud Infrastructure Console.

- To set the automatic quarterly maintenance schedule for Exadata Cloud Infrastructure
- To view or edit the properties of the next scheduled quarterly maintenance for Exadata Cloud Infrastructure
  Review and change the properties of the Exadata Cloud Infrastructure scheduled quarterly maintenance.
- To view the maintenance history of an Exadata Cloud Infrastructure resource
  This task describes how to view the maintenance history for a cloud Exadata infrastructure or DB system. resource.
- To set the node patching order for a scheduled infrastructure maintenance run
  This task describes how to set the node patching order for a scheduled infrastructure maintenance run for a cloud Exadata infrastructure or Exadata DB system resource.

To set the automatic quarterly maintenance schedule for Exadata Cloud Infrastructure

This task describes how to set the maintenance schedule for a cloud Exadata infrastructure resource.

> **Note:**
>
> Specifying a maintenance schedule is not available in Government regions.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Navigate to the Cloud Exadata infrastructure you want to access:

   In the **Oracle Exadata Database Service on Dedicated Infrastructure** section, click **Exadata Infrastructure**. In the list of infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.

   On the resource details page, under **Maintenance** or **Infrastructure Maintenance**, click the **edit** link in the **Quarterly Maintenance Schedule** field.
3. In the **Automatic Maintenance Schedule** dialog, select **Specify a schedule**.
4. Under **Maintenance months**, specify at least one month for each quarter during which Exadata infrastructure maintenance will take place. You can select more than one month per quarter. If you specify a long lead time for advanced notification (for example, 4 weeks), then you may want to specify two or three months per quarter during which maintenance runs can occur. This will ensure that your maintenance updates are applied in a timely manner after accounting for your required lead time. Lead time is discussed in the following steps.

   > **Note:**
   >
   > Maintenance quarters begin in February, May, August, and November, with the first maintenance quarter of the year beginning in February.
5. Under **Week of the month**, specify which week of the month maintenance will take place. Weeks start on the 1st, 8th, 15th, and 22nd days of the month, and have a duration of

seven days. Weeks start and end based on calendar dates, not days of the week. Maintenance cannot be scheduled for the fifth week of months that contain more than 28 days.

6. *Optional.* Under **Day of the week**, specify the day of the week on which the maintenance will occur. If you do not specify a day of the week, then Oracle will run the maintenance update on a weekend day to minimize disruption.
7. *Optional.* Under **Start hour**, specify the hour during which the maintenance run will begin. If you do not specify a start hour, then Oracle will choose the least disruptive time to run the maintenance update.
8. Under **Lead time**, specify the number of weeks ahead of the maintenance event you would like to receive a notification message. Your lead time ensures that a newly released maintenance update is scheduled to account for your required period of advanced notification.
9. **Update Maintenance Schedule**.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

To view or edit the properties of the next scheduled quarterly maintenance for Exadata Cloud Infrastructure

Review and change the properties of the Exadata Cloud Infrastructure scheduled quarterly maintenance.

NOT_SUPPORTED

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**."
2. Navigate to the Cloud Exadata infrastructure or DB system you want to access:
   Cloud Exadata infrastructure (new resource model): Under **Exadata at Oracle Cloud**, click **Exadata Infrastructure**. In the list of infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.

   DB systems: Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
3. On the resource details page, under **Maintenance**, click the **View** link in the **Next Quarterly Maintenance** field.
   The Exadata Infrastructure Maintenance page is displayed.
4. On the **Exadata Infrastructure Maintenance** page, scheduled maintenance details are listed.
   Target DB Server Version and **Target Storage Server Version**: These fields display the Exadata software version to be applied by the scheduled maintenance. The version applied will be the most recent certified update for Exadata infrastructures in the cloud. If the next quarterly update is not yet certified when the maintenance is scheduled, then the versions may show "LATEST" until the new quarterly update becomes available. Once the update becomes available the new version will be displayed.

To find information on the Database Server Exadata software version or the Storage Server Exadata software version, see My Oracle Support note *Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1)*.

For each scheduled Exadata Infrastructure resource maintenance event, the Maintenance page lists the following details:

- The status of the event
- The OCID of the event
- The scheduled start time and date of the event
- Click **Patch Now** to start the maintenance event immediately. When prompted, click **Run Maintenance** to confirm that you want to start the event now.

If a maintenance event is already in progress on one or more of the VM Clusters hosted by an Exadata Infrastructure resource when a maintenance event on that resource is to start, the Exadata Infrastructure resource maintenance event is queued and begins immediately after all VM Cluster maintenance events complete.

5. To change the next scheduled maintenance settings, click **Edit Maintenance Run**.
6. On the **Edit Maintenance** page, do the following:
   - Select a maintenance method, **Rolling** or **Non-rolling**.

     > **Note:**
     >
     > If you select the **Non-rolling** option, components will be updated simultaneously, resulting in full system downtime.

   - **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.
     - **Custom action timeout (in minutes):** Maximum timeout available to perform custom action before starting maintenance on the DB Servers.
       Default: 30 minutes

       Minimum: 15 minutes

       Maximum: 120 minutes
   - To reschedule the next maintenance run, enter a date and time in the **Scheduled Start time** field.
     The following restrictions apply:
     - You can reschedule the infrastructure maintenance to a date no more than 180 days from the prior infrastructure maintenance. If a new maintenance release is announced prior to your rescheduled maintenance run, the newer release will be applied on your specified date. You can reschedule your maintenance to take place earlier than it is currently scheduled. You cannot reschedule the maintenance if the current time is within 2 hours of the scheduled maintenance start time.

- Oracle reserves certain dates each quarter for internal maintenance operations, and you cannot schedule your maintenance on these dates.

- Click **Save Changes**.

7. To view estimated maintenance time details for various components, click the **View** link is displayed in the Total Estimated Maintenance Time field.
   The **View** link is displayed in the **Total Estimated Maintenance Time** field only if the Maintenance Method is **Rolling**.

   The **Estimated Maintenance Time Details** page is displayed with details that include:

   - Total Estimated Maintenance Time
   - Database Servers Estimated Maintenance Time
   - Storage Servers Estimated Maintenance Time
   - Order in which components are updated. In rolling maintenance, components are updated in the sequence displayed

   a. To view the number of VMs that will be restarted as part of Database Server maintenance, click the **Show details** link.
      The **VM Location** dialog is displayed.

   b. In the **VM Cluster Name** field, you can find out what VM cluster a particular VM belongs to.

   c. Click **Close**.

8. Click **Close** to close the **Estimated Maintenance Time Details** page.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.
- Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1)

To view the maintenance history of an Exadata Cloud Infrastructure resource

This task describes how to view the maintenance history for a cloud Exadata infrastructure or DB system. resource.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Navigate to the Cloud Exadata infrastructure or DB system you want to access:

   Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata Infrastructure**. In the list of infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.
3. On the resource details page, under **Bare Metal,VM, and Exadata**, click the **Maintenance History**.
4. The Maintenance jobs, State and type of patching are displayed.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

- My Oracle Support note Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1)

To set the node patching order for a scheduled infrastructure maintenance run

This task describes how to set the node patching order for a scheduled infrastructure maintenance run for a cloud Exadata infrastructure or Exadata DB system resource.

> **Note:**
>
> By default, all scheduled maintenance runs are initially set to use rolling patching. To use non-rolling patching, you must change this setting for each maintenance run after it is scheduled.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata on Oracle Public Cloud.**
2. Navigate to the cloud Exadata infrastructure or DB system you want to access:
   - *Cloud Exadata infrastructure (new resource model):* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata Infrastructure**. In the list of infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.
   - *DB systems:* Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
3. On the resource details page, under **Maintenance**, click the **view** link in the **Next Quarterly Maintenance** field.
4. On the **Maintenance** page, click the **edit** link in the **Maintenance Method** field for a scheduled cloud Exadata infrastructure maintenance run.
5. In **Update Exadata Infrastructure Node Patching Order**, change the maintenance method to either **Rolling** or **Non-rolling** as needed.

Monitor Infrastructure Maintenance Using Lifecycle State Information

The lifecycle state of your Exadata Infrastructure resource enables you to monitor when the maintenance of your infrastructure resource begins and ends.

In the Oracle Cloud Infrastructure Console, you can see lifecycle state details messages on the **Exadata Infrastructure Details** page when a tooltip is displayed beside the **Status** field. You can also access these messages using the `ListCloudExadataInfrastructures` API, and using tools based on the API, including *SDKs* and the *OCI CLI*.

During infrastructure maintenance operations, you can expect the following:

- If you specify a maintenance window, then patching begins at your specified start time. The infrastructure resource's lifecycle state changes from **Available** to **Maintenance in Progress**.

> **Note:**
>
> The prechecks are now done prior to the start of the maintenance.

- When Exadata database server maintenance starts, the infrastructure resource's lifecycle state is **Maintenance in Progress**, and the associated lifecycle state message is, **The underlying infrastructure of this system (dbnodes) is being updated.**
- When storage server maintenance starts, the infrastructure resource's lifecycle state is **Maintenance in Progress**, and the associated lifecycle state message is, **The underlying infrastructure of this system (cell storage) is being updated and this will not impact Database availability.**
- After storage server maintenance is complete, the networking switches are updated one at a time, in a rolling fashion.
- When maintenance is complete, the infrastructure resource's lifecycle state is **Available**, and the Console and API-based tools do not provide a lifecycle state message.

Related Topics

- ListCloudExadataInfrastructures
- Software Development Kits and Command Line Interface
- Command Line Interface (CLI)

Receive Notifications about Your Infrastructure Maintenance Updates

There are two ways to receive notifications. One is through email to infrastructure maintenance contacts and the other one is to subscribe to the maintenance events and get notified.

Oracle schedules maintenance run of your infrastructure based on your scheduling preferences and sends email notifications to all your infrastructure maintenance contacts. You can login to the console and view details of the schedule maintenance run. Appropriate maintenance related events will be generated as Oracle prepares for your scheduled maintenance run, for example, schedule reminder, patching started, patching end, and so on. For more information about all maintenance related events, see *Oracle Cloud Exadata Infrastructure Events*. In case, if there are any failures, then Oracle reschedules your maintenance run, generates related notification, and notifies your infrastructure maintenance contacts.

For more information about Oracle Cloud Infrastructure Events, see *Overview of Events*. To receive additional notifications other than the ones sent to infrastructure maintenance contacts, you can subscribe to infrastructure maintenance events and get notified using the Oracle Notification service, see *Notifications Overview*.

Related Topics

- Oracle Exadata Database Service on Dedicated Infrastructure Events
  Exadata Cloud Infrastructure resources emit events, which are structured messages that indicate changes in resources.
- Overview of Events
- Notifications Overview
- Managing Infrastructure Maintenance Contacts
  Learn to manage your Exadata infrastructure maintenance contacts.

Managing Infrastructure Maintenance Contacts

Learn to manage your Exadata infrastructure maintenance contacts.

- To manage maintenance contacts in an Exadata Cloud Infrastructure
  Manage contacts for Exadata infrastructure maintenance notifications using the Console.

To manage maintenance contacts in an Exadata Cloud Infrastructure

Manage contacts for Exadata infrastructure maintenance notifications using the Console.

To prevent an Exadata infrastructure administrator from being overwhelmed by system update notifications, you can specify up to 10 email addresses of people to whom maintenance notifications are sent.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. In the **Oracle Exadata Database Service on Dedicated Infrastructure** section, click **Exadata Infrastructure** to display a list of Exadata infrastructures in the default compartment. You can select a different compartment from the **Compartment** drop-down located in the **List Scope** section.
3. In the list of Exadata infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.
4. In the **Maintenance** section, click **Manage** in the **Customer Contacts** field to display the Manage Contacts dialog.
5. Click the **Add Contacts** button to display a field in which to enter a valid email address. You can have up to 10 maintenance contacts for each Exadata infrastructure.
6. To edit an email address, in the Manage Contacts dialog, select the box preceding the email address you want to edit and click the **Edit** button.
7. To remove an email address from the list, in the Manage Contacts dialog, select the box preceding the email address you want to remove and click the **Remove** button.

Using the API to Manage Exadata Cloud Infrastructure Maintenance Controls

Use these API operations to manage Exadata Cloud Infrastructure maintenance controls and resources.

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.

Use these API operations to manage Exadata Cloud Infrastructure maintenance controls.

Cloud Exadata infrastructure resource (new resource model):

- ListCloudExadataInfrastructures
- GetCloudExadataInfrastructure
- UpdateCloudExadataInfrastructure
- GetMaintenanceRun

- ListMaintenanceRuns
- UpdateMaintenanceRun

Manage VM Clusters

Learn how to manage your VM clusters on Exadata Cloud Infrastructure.

- Introduction to Scale Up or Scale Down Operations
  With the Multiple VMs per Exadata system (MultiVM) feature release, you can scale up or scale down your VM cluster resources.
- Overview of VM Cluster Node Subsetting
  VM Cluster Node Subsetting enables you to allocate a subset of database servers to new VM clusters to enable maximum flexibility in the allocation of compute (CPU, memory, local storage) resources.
- About Application VIP
  Oracle Exadata Database Service on Dedicated Infrastructure fully supports creating additional Virtual IP Addresses on an Exadata VM Cluster.
- Using the Console to Manage VM Clusters on Exadata Cloud Infrastructure
  Learn how to use the console to create, edit, and manage your VM Clusters on Oracle Exadata Database Service on Dedicated Infrastructure.
- Overview of Automatic Diagnostic Collection
  By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
- Incident Logs and Trace Files
  This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.
- Health Metrics
  Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.
- Using the API to Manage Exadata Cloud Infrastructure Instance
- Troubleshooting Virtual Machines Using Console Connections
  You can troubleshoot malfunctioning virtual machines using console connections. For example, a previously working Guest VM stops responding.

Related Topics

- Application Checklist for Continuous Service for MAA Solutions

Introduction to Scale Up or Scale Down Operations

With the Multiple VMs per Exadata system (MultiVM) feature release, you can scale up or scale down your VM cluster resources.

- Scale VM Resources in Multi VM Enabled Infrastructure
  Increase or decrease the OCPUs, memory, storage or local disk size (`/u02`) storage available to a VM cluster
- Resizing Memory and Large Pages
  You can scale the database server memory up and down in a VM Cluster. Scaling memory requires a rolling restart of the database servers to take effect.
- Calculating the ASM Storage

- Estimating How Much Local Storage You Can Provision to Your VMs
- Scaling Local Storage

Scale VM Resources in Multi VM Enabled Infrastructure

Increase or decrease the OCPUs, memory, storage or local disk size (`/u02`) storage available to a VM cluster

> **Note:**
>
> Oracle doesn't stop billing when a VM or VM Cluster is stopped. To stop billing for a VM Cluster, lower the OCPU count to zero.

Scaling up or down of these resources requires thorough auditing of existing usage and capacity management by the customer DB administrator. Review the existing usage to avoid failures during or after a scale down operation. While scaling up, consider how much of these resources are left for the next VM cluster you are planning to create. Oracle Exadata Database Service on Dedicated Infrastructure tooling calculates the current usage of memory, local disk, and ASM storage in the VM cluster, adds headroom to it, and arrives at a "minimum" value below which you cannot scale down, and expects that you specify the value below this minimum value.

Note:

- When creating or scaling a VM Cluster, setting the number of OCPUs to zero will shut down the VM Cluster and eliminate any billing for that VM Cluster, but the hypervisor will still reserve the minimum 2 OCPUs for each VM. These reserved OCPUs cannot be allocated to any other VMs, even though the VM to which they are allocated is shut down. The Control Plane does not account for reserved OCPUs when showing maximum available OCPU, so you should account for these reserved OCPU when performing any subsequent scaling operations to ensure the operation can acquire enough OCPUs to successfully complete the operation.
- For memory and `/u02` scale up or scale down operations, if the difference between the current value and the new value is less than 2%, then no change will be made to that VM. This is because memory change involves rebooting the VM, and `/u02` change involves bringing down the Oracle Grid Infrastructure stack and un-mounting `/u02`. Production customers will not resize for such a small increase or decrease, and hence such requests are a no-op.
- You can scale down the following resources in any combinations:
  - OCPU
  - Memory
  - Local storage
  - Exadata storage

  Each scaling operation can take several minutes to complete. The time for each operation will vary based on activity in the system, but as a general rule, most operations should complete within 15 minutes for a quarter rack, 20 minutes for a half rack, and 30 minutes for a full or larger rack. Performing multiple OCPU scaling operations over a short period of time can lengthen the time for completion. Although online, OCPU scaling is not implemented on all VMs in parallel so as to detect and protect from any anomalies before they affect the entire system. Memory and Local Storage scaling require a VM reboot, and are performed one VM at a time in a rolling manner.

  If you run multiple scale-down operations, then each operation is performed serially. For example, if you scale memory and local storage from the Console, then the system will first scale memory, and when that operation completes, it will scale storage. The time to complete all operations will be the sum of the time to complete individual operations.
- Storage servers added to the infrastructure but yet to run the 'Add Capacity' step will not have any disk groups created on them.

Resizing Memory and Large Pages

You can scale the database server memory up and down in a VM Cluster. Scaling memory requires a rolling restart of the database servers to take effect.

Changing the memory in a VM Cluster will affect the large pages (HugePages) settings for the VMs in that cluster. When a VM is initially created, each VM's operating system is configured with 50% of the memory allocated to the VM for large pages, and databases are configured to use that memory for their SGA. Oracle recommends that you not modify the large pages configuration unless you understand the implication of any changes you make. Improper

configurations can prevent all databases from starting, and even prevent the VM from starting up.

Although not recommended, you are allowed to modify the large pages configuration. Any changes you make may be modified by automation should you subsequently resize the memory available to the VM. After a memory resize operation, the cloud automation will attempt to maintain the same amount of large pages memory as a percentage of the total memory, with a cap of 60%. If you configure large pages to be greater than 60% of total memory, then the cloud automation will resize it to 60% of total memory. This automatic resize is to ensure sufficient conventional memory for the virtual machine to start. The automation will perform a precheck to determine the actual large pages memory in use by the running database instances, and ensure after the resize that there is enough large pages memory available to support those same databases. If there will not be sufficient memory available after the resize, then the precheck will fail and the resize will not continue.

Calculating the ASM Storage

Use the following formula to calculate the minimum required ASM storage:

- For each disk group, for example, `DATA`, `RECO`, note the total size and free size by running the `asmcmd lsdg` command on any Guest VM of the VM cluster.
- Calculate the used size as (Total size - Free size) / 3 for each disk group. The /3 is used because the disk groups are triple mirrored.
- DATA:RECO ratio is:

  80:20 if **Local Backups** option was NOT selected in the user interface.

  40:60 if **Local Backups** option was selected in the user interface.
- Ensure that the new total size as given in the user interface passes the following conditions:
  Used size for DATA * 1.15 <= (New Total size * DATA %)

  Used size for RECO * 1.15 <= (New Total size * RECO %)

Example 5-2 Calculating the ASM Storage

1. Run the `asmcmd lsdg` command in the Guest VM:
   - Without SPARSE:

     ```
     /u01/app/19.0.0.0/grid/bin/asmcmd lsdg
     ASMCMD>
     State    Type  Rebal  Sector  Logical_Sector  Block       AU  Total_MB
     Free_MB    Req_mir_free_MB   Usable_file_MB   Offline_disks
     Voting_files   Name
     MOUNTED  HIGH  N         512     512        4096  4194304  12591936
     10426224   1399104           3009040           0
     Y      DATAC5/
     MOUNTED  HIGH  N         512     512        4096  4194304   3135456
     3036336    348384            895984            0
     N      RECOC5/
     ASMCMD>
     ```

   - With SPARSE:

     ```
     /u01/app/19.0.0.0/grid/bin/asmcmd lsdg
     ASMCMD>
     ```

```
State    Type Rebal Sector Logical_Sector Block AU       Total_MB
Free_MB   Req_mir_free_MB   Usable_file_MB   Offline_disks
Voting_files   Name
MOUNTED HIGH N         512     512        4096 4194304   12591936
10426224  1399104            3009040             0
Y     DATAC5/
MOUNTED HIGH N         512     512        4096 4194304   3135456
3036336   348384             895984              0
N     RECOC5/
MOUNTED HIGH N         512     512        4096 4194304   31354560
31354500  3483840            8959840             0
N     SPRC5/
ASMCMD>
```

> **Note:**
>
> The listed values of all attributes for SPARSE diskgroup (SPRC5) present the virtual size. In Exadata DB Systems and Exadata Cloud Infrastructure, we use the ratio of 1:10 for `physicalSize:virtualSize`. Hence, for all purposes of our calculation we must use 1/10th of the values displayed above in case of SPARSE for those attributes.

2. Used size for a disk group = (Total_MB - Free_MB) /3
   - Without SPARSE:
     Used size for DATAC5 = (12591936 - 10426224) / 3 = 704.98 GB

     Used size for RECO5 = (3135456 - 3036336) / 3 = 32.26 GB
   - With SPARSE:
     Used size for DATAC5 = (12591936 - 10426224) / 3 ~= 704.98 GB

     Used size for RECO5 = (3135456 - 3036336) /3 ~= 32.26 GB

     Used size for SPC5 = (1/10 * (31354560 - 31354500)) / 3 ~= 0 GB
3. Storage distribution among diskgroups
   - Without SPARSE:
     DATA:RECO ratio is 80:20 in this example.
   - With SPARSE:
     DATA RECO: SPARSE ratio is 60:20:20 in this example.
4. New requested size should pass the following conditions:
   - Without SPARSE: (For example, 5 TB in user interface.)
     5 TB = 5120 GB ; 5120 *.8 = 4096 GB; 5120 *.2 = 1024 GB

     For DATA: (704.98 * 1.15) <= 4096 GB

     For RECO: (32.36 * 1.15) <= 1024 GB
   - With SPARSE: (For example, 8 TB in the user interface.)
     8 TB = 8192 GB; 8192 *.6 = 4915 GB; 8192 *.2 = 1638 GB; 8192 *.2 = 1638 GB

     For DATA: (704.98 * 1.15) <= 4915 GB

     For RECO: (32.36 * 1.15) <= 1638 GB

     For SPR: (0 * 1.15) <= 1638 GB

Above resize will go through. If above conditions are not met by the new size, then resize will fail the precheck.

Estimating How Much Local Storage You Can Provision to Your VMs

> **Note:**
>
> The following does not apply to X6, X7, X8, and Base Systems as they do not support multiple VMs. The Base System has 200 GB available for `/u02`.

VM Images include the files necessary to boot and run the VM and its operating system, as well as space for Oracle Homes stored in `/u02`. To estimate how much additional local storage space beyond the minimum can be allocated to any file system associated with a VM, subtract the size of the VM images for all VMs on a server from the total available space. If you have not modified the default VM Image size by expanding any file systems, use the VM Image size (default and minimum) below. If you have or plan to modify your VM Image size, you must use the OCI console and "Scale VM Cluster" action to check the allocated and available for an existing VM Cluster as expanding some non-/u02 file systems will consume more incremental storage than was added to the file system. This information is also available in the "Configure VM Cluster" action while creating a new VM Cluster.

X8M-2 Systems

- Total space available for VM images (X8M): 2243 GB
- VM Image size (default and minimum) including `/u02`: 244 GB
- Default (minimum) `/u02`: 60 GB

X9M-2 Systems

- Total Available for VM Images: 2243 GB
- VM Image size (default and minimum) including `/u02`: 244 GB
- Default (minimum) `/u02`: 60 GB

Example: If you have an X9M Elastic System with 2 VMs per physical server, and have not made any changes to any of the file systems, you will have 2243 GB available for all VMs, and each will consume 244 GB (488 total), leaving 1755 GB to expand any VM Images. The default VM image will include 60GB of `/u02` per VM to store Oracle homes. The 1755 GB of available space can be used to expand `/u02`, or can be used to expand other file systems in the VM Image. Every GB used to expand `/u02` will consume a GB of available space. Every GB used to expand other file systems in the VM image may consume more than a GB of space. Refer to the information in the console when expanding non-/u02 file systems to see the actual available space impact of expanding these file systems.

Scaling Local Storage

Scale Local Space Operation Guidelines

You can scale local storage by modifying the size of many of the individual file systems in a VM. By default, the file systems are created at their minimum size. You can increase the size of the file systems as required. However, note that you can only shrink `/u02`. Other file systems can only be increased in size. The maximum supported size of any file system is 900 GB.

The storage consumed by all file systems is greater than the sum of the file system sizes. Refer to the calculations displayed in the OCI console to see the effects on free local storage when resizing a file system.

Using the OCI Console or API, you can increase or decrease the size of the following local file systems:

- `/u02`

Using the OCI Console or API, you can increase the size of following local file systems:

- `/`
- `/u01`
- `/tmp`
- `/var`
- `/var/log`
- `/var/log/audit`
- `/home`

However, you cannot resize the following local file systems:

- `/crashfiles`
- `/boot`
- `/acfs01`
- `/u01/app/19.0.0.0/grid`

> **Note:**
>
> - With the exception of `/u02`, you can only expand the file systems and cannot reduce their size once they have been expanded.
> - A rolling restart of each VM is required for the resizing to take effect.
> - Each file system can only be expanded to a maximum of 900 GB
> - Ability to increase the size of additional local file systems is only supported on X8M and later systems.

For more information about resizing these file systems, see *Estimating How Much Local Storage You Can Provision to Your VMs.*

Resource Limit Based On Current Utilization

- Any scale-down operation must leave 15% buffer on top of highest local space utilization across all nodes in the cluster.
- The lowest local space per node allowed is higher of the above two limits.
- Run the `df -kh` command on each node to find out the node with the highest local storage.
- You can also use the utility like `cssh` to issue the same command from all hosts in a cluster by typing it just once.

- Lowest value of local storage each node can be scaled down to would be = 1.15x (highest value of local space used among all nodes).

ACFS File Systems

If requested by support, you can also resize the `/acfs01` file system. This file system is used by the system to stage software. It uses Exadata storage and is not subject to the limits described above for `/u02`. It is a shared file system visible from all nodes in the cluster, and can be online resized from the command line of any VM.

- **Default size:** The default size of `/acfs01` is 100 GB.
- **Scaling /acfs01:** You can scale `acfs01` as user `grid` from any VM via the `/sbin/acfsutil` command. No reboot is required. The resize operation will not affect the availability of the database service running in the VM cluster. The following command issued by the `grid` user will increase the size of `/acfs01` by 100 GB: `/sbin/acfsutil size +100 GB /acfs01`.
- You can create additional ACFS file systems if required. These will also consume storage from the Exadata Storage diskgroups and may be shared across all VMs in the cluster. Refer to the ACFS documentation for more information.

Overview of VM Cluster Node Subsetting

VM Cluster Node Subsetting enables you to allocate a subset of database servers to new VM clusters to enable maximum flexibility in the allocation of compute (CPU, memory, local storage) resources.

With VM Cluster Node Subsetting, you can:

- Create a smaller VM cluster to host databases that have low resource and scalability requirements or to host a smaller number of databases that require isolation from the rest of the workload.
- Expand or shrink an existing VM cluster by adding and removing nodes to ensure optimal utilization of available resources.

Consider reviewing the points below that will assist you in subsetting VM cluster nodes.

- VM Cluster Node Subsetting capability is available for new VM clusters in Oracle Exadata Database Service on Dedicated Infrastructure service.
- All VMs across a VM cluster will have the same resource allocation per VM irrespective of whether the VM was created during cluster provisioning or added later by extending an existing VM cluster.
- Any VM cluster must have a minimum of 2 DB servers.
- You can host a maximum of 4 VMs on X8M and above generations of DB Servers.
- Exadata Infrastructures with X8M and above generation of DB Servers can support a maximum of 8 VM clusters across all DB Servers.
- Maximum number of clusters across the infrastructure depends on the resources available per DB server and is subject to the per DB Server maximum VM limit.

With the release of Multi-VM, the add and remove virtual machine api for cloud vm clusters will not be supported via terraform.

You can perform these operations through UI, SDK, OCI CLI, OCI Ansible or similar tools. Terraform states should be managed similar to other operations which happen outside of terraform but need to be managed in terraform.

For more information, see *Detecting and Managing Drift with Terraform*.

- Add a VM to a VM Cluster
  Add a Virtual Machine to a VM Cluster
- Terminate a VM from a VM Cluster
  To remove a virtual machine from a provisioned cluster, use this procedure.

Related Topics

- Scaling Resources within an Exadata Infrastructure Instance
  If an Exadata Cloud Infrastructure instance requires more resources, you can scale up the number of DB servers, or storage servers.
- To create a cloud VM cluster resource
  Create a VM cluster in an Exadata Cloud Infrastructure instance.
- Add a VM to a VM Cluster
  Add a Virtual Machine to a VM Cluster
- Using the Console to View a List of DB Servers on an Exadata Infrastructure
  To view a list of database server hosts on an Oracle Exadata Database Service on Dedicated Infrastructure system, use this procedure.
- Terminate a VM from a VM Cluster
  To remove a virtual machine from a provisioned cluster, use this procedure.
- Detecting and Managing Drift with Terraform

Add a VM to a VM Cluster

Add a Virtual Machine to a VM Cluster

> **Note:**
>
> Once the VM cluster is upgraded to Exadata Database Service Guest VM OS 23.1, you will be able to add a new VM or a new database server to this VM cluster if Exadata Cloud Infrastructure is running an Exadata System Software version 22.1.16 and later.
> Upgrade to Exadata System Software 23.1 for Exadata Cloud Infrastructure will be available with February 2023 update cycle.

> **Note:**
>
> - This operation is only available with Multi-VM enabled Infrastructure.
> - To add a VM to a VM Cluster requires that all TCP ports to be open for the client subnet CIDR for ingress and egress.

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** and **Compartment** that contains the VM cluster for which you want to scale the CPU resources.
3. Click **VM Clusters**.

4. Click the name of the VM cluster to which you want to add a virtual machine.
5. Under Resources, select **Virtual Machines**, and click the **Add Virtual Machines** button.
6. In the Add Virual Machines window, select the DB server where you want the new VM to reside.

> **Note:**
>
> The VM that is added will have the same resources as the other VMs in the cluster.

7. Click **Add Virtual Machine**.

> **Note:**
>
> Add a VM to a VM Cluster is NOT supported using Terraform.

Related Topics

- Adding a VM to a VM Cluster Fails

Terminate a VM from a VM Cluster

To remove a virtual machine from a provisioned cluster, use this procedure.

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** and **Compartment** that contains the VM cluster for which you want to scale the CPU resources.
3. Click **VM Clusters**.
4. Click the name of the VM cluster for which you want to remove a virtual machine.
5. On the Exadata VM Cluster Details page, in the Virtual Machines section, select the Virtual Machine that will be removed, click the more commands symbol (three dots) and click **Terminate**

> **Note:**
>
> Remove a VM from a VM Cluster is NOT supported using Terraform at this time.

About Application VIP

Oracle Exadata Database Service on Dedicated Infrastructure fully supports creating additional Virtual IP Addresses on an Exadata VM Cluster.

These application VIPs are required to protect additional applications such as Oracle GoldenGate installed on an Oracle Exadata Database Service on Dedicated Infrastructure system or other services such as XA-Agents, and to provide High Availability to these additional applications. For more information, see Oracle Grid Infrastructure Standalone

Agents for Oracle Clusterware and Making Applications Highly Available Using Oracle Clusterware.

Within the Oracle Cloud Infrastructure, adding Virtual IP Addresses on the cluster stack alone is not sufficient as these additional (secondary) IP addresses also have to be added to the VCN layer as "**Floating IP**" addresses so that the VCN layer knows where these IP addresses are running, and in case of failover by the Clusterware to change the VNIC the floating IP address is attached to. For more information, see Creating an Application VIP Managed by Oracle Clusterware and Overview of IP Addresses.

Adding an Application VIP to an Exadata VM Cluster consists of the following steps:

1. Add the Virtual IP address to the Clusterware layer within the Exadata DomU, by following the standard Oracle Clusterware documentation, or the guide provided by the application, for example, by using

   ```
   appvipcfg create -network=1 -ip=10.10.10.10
            -vipname=applicationvip
   ```

   .

2. Attach the Application Floating IP address object on the Exadata VM Cluster to add the knowledge of the floating IP to the VCN layer. Ensure that you choose the same subnet as you created the backend application VIP, which normally is the client subnet.

   The private IP address needs to be the same as the one specified in the `appvipcfg` command above. The **Virtual IP Address Hostname** is the name under which the IP address is reachable via DNS and does not have to be the same as the `vipname`.

   If you have already started the VIP in the backend, ensure that the **Virtual Machine Name** reflects the host on which the VIP was started in the backend.

3. Test the relocation of the VIP. The VIP should stay available (test this via. ping), and the user interface should display after a short while that the floating IP also has moved to another host.

   If you did choose the wrong host while creating the VCN attachment, then simply relocate the VIP within the cluster. The VCN layer will detect the change and the user interface should get updated after a short while.

> **Note:**
>
> A single Virtual VM Cluster has a limitation of 8 additional Application VIPs. The limitation exists because a single VNIC can only have 31 additional secondary IP addresses. For more information, see Overview of IP Addresses. If all VIPs are started on the same node, then the application VIPs cannot be reached.

If more application VIPs are required, then raise an SR to have this limit increased. However, there are a few additional steps required then, to ensure that under no scenario more than 31 secondary IP addresses are attached to a single Exadata VM Cluster node. One way to accomplish this would be to ensure application VIPs are bound by the Clusterware to certain nodes so that this scenario is prevented.

A setup with 32 additional application VIPs would look as follows:

Floating IP	Node 1	Node 2	Node 3	Node 4
Private Hostname	1	1	1	1
VIP Hostname	4	4	4	4
SCAN	3	3	3	3
Appvip 1-8	8	8	-	-
Appvip 9-16	-	8	8	-
Appvip 17-24	-	-	8	8
Appvip 25-32	8	-	-	8
Max Possible VIP if all Floating IPs failover	24	24	24	24

Related Topics

- To Attach a Virtual IP Address
  Attach a Virtual IP address from a VM cluster using this procedure.
- To Detach a Virtual IP Address
  Attach a Virtual IP address from a VM cluster using this procedure.
- Resource-Types for Exadata Cloud Service Instances
- application-vips
  Review the list of permissions and API operations for `application-vips` resource-type.
- Permissions Required for Each API Operation
  The following tables list the API operations for Exadata Cloud Infrastructure instances in a logical order, grouped by resource type.
- Application VIP Event Types
  These are the event types that Application VIPs in Oracle Cloud Infrastructure emit.

Using the Console to Manage VM Clusters on Exadata Cloud Infrastructure

Learn how to use the console to create, edit, and manage your VM Clusters on Oracle Exadata Database Service on Dedicated Infrastructure.

- To create a cloud VM cluster resource
  Create a VM cluster in an Exadata Cloud Infrastructure instance.
- To add database server or storage server capacity to a cloud VM cluster
  This topic describes how to use the Oracle Cloud Infrastructure (OCI) Console to add the new capacity to your cloud VM cluster.
- Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.
- Using the Console to Update the License Type on a VM Cluster
  To modify licensing, be prepared to provide values for the fields required for modifying the licensing information.
- To add SSH keys to a VM cluster
  The VM cluster exists, and you wish to add a another user which requires another SSH key.

- Using the Console to Add SSH Keys After Creating a VM Cluster
- Using the Console to Stop, Start, or Reboot a VM Cluster Virtual Machine
  Use the console to stop, start, or reboot a virtual machine.
- Using the Console to Check the Status of a VM Cluster Virtual Machine
  Review the health status of a VM cluster virtual machine.
- Using the Console to Move a VM Cluster to Another Compartment
  To change the compartment that contains your VM cluster on Exadata Cloud Infrastructure, use this procedure.
- To change the VM cluster display name
- Using the Console to Terminate a VM Cluster
  Before you can terminate a VM cluster, you must first terminate the databases that it contains.
- To view details about private DNS configuration
- To Attach a Virtual IP Address
  Attach a Virtual IP address from a VM cluster using this procedure.
- To Detach a Virtual IP Address
  Attach a Virtual IP address from a VM cluster using this procedure.

To create a cloud VM cluster resource

Create a VM cluster in an Exadata Cloud Infrastructure instance.

> **Note:**
>
> To create a cloud VM cluster in an Exadata Cloud Infrastructure instance, you must have first created a Cloud Exadata infrastructure resource.

> **Note:**
>
> Multi-VM enabled Infrastructure will support creating multiple VM Clusters. Infrastructures created before the feature Create and Manage Multiple Virtual Machines per Exadata System (MultiVM) and VM Cluster Node Subsetting was released only support creating a single cloud VM cluster.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.

   > **Note:**
   >
   > Multiple VM clusters may be created only in a Multi-VM enabled Infrastructure.

3. Click **Create Exadata VM Cluster**.
   The **Create Exadata VM Cluster** page is displayed. Provide the required information to configure the VM cluster.

4. **Compartment:** Select a compartment for the VM cluster resource.
5. **Display name:** Enter a user-friendly display name for the VM cluster. The name doesn't need to be unique. An Oracle Cloud Identifier (OCID) will uniquely identify the DB system. Avoid entering confidential information.
6. **Select Exadata infrastructure:** Select the infrastructure resource that will contain the VM cluster. You must choose an infrastructure resource that has enough resources to create a new VM cluster. Click **Change Compartment** and pick a different compartment from the one you are working in to view infrastructure resources in other compartments.

   > **Note:**
   >
   > Multiple VM clusters may be created only in a Multi-VM enabled Infrastructure

7. **Choose the Oracle Grid Infrastructure version:** From the list, choose the Oracle Grid Infrastructure release (19c and 23ai) that you want to install on the VM cluster.
   The Oracle Grid Infrastructure release determines the Oracle Database releases that can be supported on the VM cluster. You cannot run an Oracle Database release that is later than the Oracle Grid Infrastructure software release.

   > **Note:**
   >
   > Minimum requirements for provisioning a VM Cluster with Grid Infrastructure 23ai:
   >
   > - Exadata Guest VM running Exadata System Software 23.1.8
   > - Exadata Infrastructure running Exadata System Software 23.1.x

8. **Choose an Exadata image version:**
   - **Exadata infrastructure with Oracle Linux 7 and Exadata image version 22.1.10.0.0.230422:**
     - The **Change image** button is not enabled.
     - The Oracle Grid Infrastructure version defaults to 19.0.0.0.0.
     - The Exadata guest version will be the same as that of the host OS.
   - **Exadata infrastructure with Oracle Linux 8 and Exadata image version 23.1.3.0.0.230613:**
     - The Exadata guest version defaults to the latest (23.1.3.0).
     - The Oracle Grid Infrastructure version defaults to 19.0.0.0.0
     - The **Change image** button is enabled.
     - Click **Change image**.
       The resulting Change image panel displays the list of available major versions of Exadata image (23.1.3.0 and 22.1.3.0).

       The most recent release for each major version is indicated by "(latest)".
     - Slide **Display all available versions**.
       Six past versions including the latest versions of Exadata images 23.1.3.0 and 22.1.3.0 are displayed.
     - Choose a version.

- Click **Save Changes**.

9. **Configure the VM cluster:** Specify the DB servers to used for new VM cluster (by default all DB Servers are selected). Click **Change DB Servers** to select from the available DB servers. In the **Resource allocation per VM** pane:
   - Specify the number of OCPU cores you want to allocate to each of the VM cluster's virtual machine compute nodes. Minimum is 2 OCPU per VM . The read-only **Requested OCPU count for the Exadata VM cluster** field displays the total number of OCPU cores you are allocating.
   - Specify the **Memory per VM** to allocate to each VM. The minimum per VM is 30 GB.
   - Specify the **Local Storage per VM** to allocate local storage to each VM. The minimum per VM is 60 GB.
     Each time when you create a new VM cluster, the space remaining out of the total available space is utilized for the new VM cluster.

     In addition to `/u02`, you can specify the size of additional local file systems.

     For more information and instructions to specify the size for each individual VM, see Introduction to Scale Up or Scale Down Operations.
     - Click **Show additional local file systems configuration options**.
     - Specify the size of `/`, `/u01`, `/tmp`, `/var`, `/var/log`, `/var/log/audit`, and `/home` file systems as needed.

       > **Note:**
       >
       > * You can only expand these file systems and cannot reduce the size once expanded.
       > * Due to backup partitions and mirroring, the `/` and `/var` file systems will consume twice the space they were allocated, which is indicated in the read-only **Total allocated storage for / (GB) due to mirroring** and **Total allocated storage for /tmp (GB) due to mirroring** fields.

     - After creating the VM Cluster, check the **Exadata Resources** section on the **Exadata Infrastructure Details** page to check the file size allocated to the local storage (`/u02`) and local storage (additional file systems).

10. **Configure Exadata storage:** Specify the following:
    - **Specify the usable Exadata storage TB**. Specify the storage in multiples of 1 TB. Minimum: 2 TB
    - **Allocate storage for Exadata sparse snapshots:** Select this configuration option if you intend to use snapshot functionality within your VM cluster. If you select this option, the SPARSE disk group is created, which enables you to use VM cluster snapshot functionality for PDB sparse cloning. If you do not select this option, the SPARSE disk group is not created and snapshot functionality will not be available on any database deployments that are created in the environment.

      > **Note:**
      >
      > The storage configuration option for sparse snapshots cannot be changed after VM cluster creation.

- **Allocate storage for local backups:** Select this option if you intend to perform database backups to the local Exadata storage within your Exadata Cloud Infrastructure instance. If you select this option, more space is allocated to the RECO disk group, which is used to store backups on Exadata storage. If you do not select this option, more space is allocated to the DATA disk group, which enables you to store more information in your databases.

> **Note:**
>
> The storage configuration option for local backups cannot be changed after VM cluster creation.

11. **Add SSH key:** Add the public key portion of each key pair you want to use for SSH access to the DB system:
    - **Generate SSH key pair** (Default option) Select this radio button to generate an SSH keypair. Then in the dialog below click **Save private key** to download the key, and optionally click **Save public key** to download the key.
    - **Upload SSH key files:** Select this radio button to browse or drag and drop .pub files.
    - **Paste SSH keys:** Select this radio button to paste in individual public keys. To paste multiple keys, click **+ Another SSH Key**, and supply a single key for each entry.
12. **Configure the network settings:** Specify the following:

> **Note:**
>
> IP addresses (100.64.0.0/10) are used for Exadata Cloud Infrastructure X8M interconnect.

    - **Virtual cloud network:** The VCN in which you want to create the VM cluster. Click **Change Compartment** to select a VCN in a different compartment.
    - **Client subnet:** The subnet to which the VM cluster should attach. Click **Change Compartment** to select a subnet in a different compartment.
      Do not use a subnet that overlaps with 192.168.16.16/28, which is used by the Oracle Clusterware private interconnect on the database instance. Specifying an overlapping subnet causes the private interconnect to malfunction.
    - **Backup subnet:** The subnet to use for the backup network, which is typically used to transport backup information to and from the **Backup Destination**, and for Data Guard replication. Click **Change Compartment** to select a subnet in a different compartment, if applicable.
      Do not use a subnet that overlaps with 192.168.128.0/20. This restriction applies to both the client subnet and backup subnet.

      If you plan to back up databases to Object Storage or Autonomous Recovery service, see the network prerequisites in Managing Exadata Database Backups.

> **Note:**
>
> In case Autonomous Recovery Service is used, a new dedicated subnet is highly recommended. Review the network requirements and configurations required to backup your Oracle Cloud databases to Recovery Service. See, Configuring Network Resources for Recovery Service.

- **Network Security Groups:** Optionally, you can specify one or more network security groups (NSGs) for both the client and backup networks. NSGs function as virtual firewalls, allowing you to apply a set of ingress and egress *security rules* to your Exadata Cloud Infrastructure VM cluster. A maximum of five NSGs can be specified. For more information, see *Network Security Groups* and *Network Setup for Exadata Cloud Infrastructure Instances*.

  Note that if you choose a subnet with a *security list*, the security rules for the VM cluster will be a union of the rules in the security list and the NSGs.

  To use network security groups:

  – Check the **Use network security groups to control traffic** check box. This box appears under both the selector for the client subnet and the backup subnet. You can apply NSGs to either the client or the backup network, or to both networks. Note that you must have a virtual cloud network selected to be able to assign NSGs to a network.

  – Specify the NSG to use with the network. You might need to use more than one NSG. If you're not sure, contact your network administrator.

  – To use additional NSGs with the network, click **+;Another Network Security Group**.

- *To use private DNS Service*

  > **Note:**
  >
  > A Private DNS must be configured before it can be selected. See *Configure Private DNS*

  – Check the **Use private DNS Service** check box.

  – Select a private view. Click **Change Compartment** to select a private view in a different compartment.

  – Select a private zone. Click **Change Compartment** to select a private zone in a different compartment.

- **Hostname prefix:** Your choice of hostname for the Exadata DB system. The host name must begin with an alphabetic character and can contain only alphanumeric characters and hyphens (-). The maximum number of characters allowed for an Exadata DB system is 12.

  > **Caution:**
  >
  > The hostname must be unique within the subnet. If it is not unique, the VM cluster will fail to provision.

- **Host domain name:** The domain name for the VM cluster. If the selected subnet uses the Oracle-provided Internet and VCN Resolver for DNS name resolution, this field displays the domain name for the subnet and it can't be changed. Otherwise, you can provide your choice of the domain name. Hyphens (-) are not permitted.
  If you plan to store database backups in Object Storage or Autonomous Recovery service, Oracle recommends that you use a VCN Resolver for DNS name resolution for the client subnet because it automatically resolves the Swift endpoints used for backups.
- **Host and domain URL:** This read-only field combines the host and domain names to display the fully qualified domain name (FQDN) for the database. The maximum length is 63 characters.

13. **Choose a license type:** The type of license you want to use for the VM cluster. Your choice affects metering for billing.
    - **License Included** means the cost of the cloud service includes a license for the Database service.
    - **Bring Your Own License (BYOL)** means you are an Oracle Database customer with an Unlimited License Agreement or Non-Unlimited License Agreement and want to use your license with Oracle Cloud Infrastructure. This removes the need for separate on-premises licenses and cloud licenses.
14. **Diagnostics Collection:** By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.

> **Note:**
>
> You are opting in with the understanding that the above list of events (or metrics, log files) can change in the future. You can opt out of this feature at any time.

- **Enable Diagnostic Events**: Allow Oracle to collect and publish critical, warning, error, and information events to me.
- **Enable Health Monitoring**: Allow Oracle to collect health metrics/events such as Oracle Database up/down, disk space usage, and so on, and share them with Oracle Cloud operations. You will also receive notification of some events.
- **Enable Incident Logs and Trace Collection**: Allow Oracle to collect incident logs and traces to enable fault diagnosis and issue resolution.

> **Note:**
>
> You are opting in with the understanding that the above list of events (or metrics, log files) can change in the future. You can opt-out of this feature at any time.

All three checkboxes are selected by default. You can leave the default settings as is or clear the checkboxes as needed. You can view the Diagnostic Collection settings on the **VM Cluster Details** page under **General Information >> Diagnostics Collection**.

- **Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (all three options).
- **Disabled:** When you choose not to collect diagnostics, health metrics, incident logs, and trace files (all three options).

- **Partially Enabled**: When you choose to collect diagnostics, health metrics, incident logs, and trace files (one or two options).

15. Click **Show Advanced Options** to specify advanced options for the VM cluster:

    - **Time zone:** This option is located in the **Management** tab. The default time zone for the DB system is UTC, but you can specify a different time zone. The time zone options are those supported in both the Java.util.TimeZone class and the Oracle Linux operating system. For more information, see *DB System Time Zone* .

      > **Note:**
      >
      > If you want to set a time zone other than UTC or the browser-detected time zone, and if you do not see the time zone you want, try selecting the **Select another time zone**, option, then selecting "Miscellaneous" in the **Region or country** list and searching the additional **Time zone** selections.

    - **SCAN Listener Port**: This option is located in the **Network** tab. You can assign a SCAN listener port (TCP/IP) in the range between 1024 and 8999. The default is 1521

      > **Note:**
      >
      > Manually changing the SCAN listener port of a VM cluster after provisioning using the backend software is not supported. This change can cause Data Guard provisioning to fail.

      .

    - **Tags**: If you have permissions to create a resource, then you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.

16. Click **Create Exadata VM Cluster**.

WHAT NEXT?

After your VM cluster is successfully created and in the Available state, you can view the VM Cluster Details page by clicking the name of the VM cluster in the list of clusters. From the VM Cluster Details page, you can *create your first database* in the cluster by clicking **Create Database**.

Related Topics

- Network Security Groups
- Network Setup for Exadata Cloud Infrastructure Instances
  This topic describes the recommended configuration for the VCN and several related requirements for the Exadata Cloud Infrastructure instance.
- Security Lists
- Configure Private DNS
  Prerequistes needed to use Private DNS
- DB System Time Zone

- Resource Tags
- To create a database in an existing Exadata Cloud Infrastructure instance
  This topic covers creating your first or subsequent databases.
- Getting Started with Events
- Overview of Database Service Events
- Overview of Automatic Diagnostic Collection
  By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.

To add database server or storage server capacity to a cloud VM cluster

This topic describes how to use the Oracle Cloud Infrastructure (OCI) Console to add the new capacity to your cloud VM cluster.

> **Note:**
>
> This procedure does not apply to Multi-VM enabled Infrastructure

If you have used the task *To add compute and storage resources to a flexible cloud Exadata infrastructure resource* by adding additional database (compute) or storage servers to the service instance, you must add the additional capacity to the cloud VM cluster to utilize the additional resources.

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.
3. In the list of cloud VM clusters, click the name of the cluster to which you want to add capacity.
4. On the VM Cluster Details page, click **Scale VM Cluster**.
5. If you have additional capacity available as a result of scaling the cloud Exadata infrastructure resource, a banner at the top of the **Scale VM Cluster** panel provides a message telling you the type and amount of additional capacity available to the VM cluster. Check the **Add Capacity** box.
6. Select either the **Add Database Server** or the **Add Storage** radio button, depending on which type of capacity you want to add to the cloud VM cluster.
7. Click **Update**. The cloud VM cluster goes into the Updating state. When the capacity has been successfully added, the cluster returns to the Available state.

> **Note:**
>
> If you have added additional database servers to the cluster, you can allocate the new CPU cores once the cluster is in the Available state by clicking the **Scale VM Cluster** button again. See *To scale CPU cores in an Exadata Cloud Service cloud VM cluster or DB system* for more information on adding CPU cores to your cloud VM cluster.

Related Topics

- To add compute and storage resources to a flexible cloud Exadata infrastructure resource
  This task describes how to use the Oracle Cloud Infrastructure Console to scale a flexible cloud Exadata infrastructure resource.
- To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster or DB system

Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection

You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.

> **Note:**
>
> - You are opting in with the understanding that the list of events, metrics, and log files collected can change in the future. You can opt-out of this feature at any time.
> - Oracle may add more metrics in the future, but if you have already chosen to collect metrics, you need not update your opt-in value. It will remain enabled/disabled based on your current preference.
> - If you have previously opted in for incident log and trace file collection and decide to opt out when Oracle Cloud operations run a log collection job, then the job will run its course and will not cancel. Future log collections won't happen until you opt-in again to the incident logs and trace file collection option.

1. Open the navigation menu. Under **Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Click **VM Clusters**.
4. Click the name of the VM cluster you want to enable or disable diagnostic data collection.
5. On the VM Cluster Details page, under **General Information**, enable, partially enable, or disable **Diagnostics Collection** beside **Diagnostics Collection**.
6. In the **Edit Diagnostics Collection Settings** dialog, enable or disable any of the Diagnostics Collections. By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
   - **Enable Diagnostics Events** Allow Oracle to collect and publish critical, warning, error, and information events to me. For more information, see *Overview of Database Service Events*
   - **Enable Health Monitoring** Allow Oracle to collect health metrics/events such as Oracle Database up/down, disk space usage, and so on, and share them with Oracle Cloud operations. You will also receive notification of some events.
   - **Enable Incident logs and trace collection**. Allow Oracle to collect incident logs and traces to enable fault diagnosis and issue resolution.
     Note: You had previously opted in for incident log and trace file collection and decide to opt-out when Oracle Cloud operations run a log collection job, the job will run its

course and will not cancel. Future log collections will not run until you opt-in again to the incident logs and trace file collection option.

7. Select or clear the checkboxes and then click **Save Changes**.

Related Topics

- Overview of Database Service Events

Using the Console to Update the License Type on a VM Cluster

To modify licensing, be prepared to provide values for the fields required for modifying the licensing information.

1. Open the navigation menu. Under **Oracle Database**, click **Exadata Cloud Infrastructure**.
2. Choose the **Region** and **Compartment** that contains the VM cluster for which you want to update the license type.
3. Click **VM Clusters**.
4. Click the name of the VM cluster for which you want to update the license type.

   The VM Cluster Details page displays information about the selected VM cluster.
5. Click **Update License Type**.
6. In the dialog box, choose one of the following license types and then click **Save Changes**.
   - **Bring Your Own License (BYOL):** Select this option if your organization already owns Oracle Database software licenses that you want to use on the VM cluster.
   - **License Included:** Select this option to subscribe to Oracle Database software licenses as part of Exadata Cloud Infrastructure.

   Updating the license type does not change the functionality or interrupt the operation of the VM cluster. Customers are permitted to change the license type for a VM Cluster at most once per month.

To add SSH keys to a VM cluster

The VM cluster exists, and you wish to add a another user which requires another SSH key.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters**.
4. In the list of VM clusters, find the cluster you want to manage and click its highlighted name.
5. Click **Add SSH Keys**.
6. Select one of the following options:
   - **Generate SSH key pair**: Use this option to create a new SSH key pair. Click both **Save Private Key** and **Save Public Key** when using this option. The private key is downloaded to your local machine, and should be stored in a safe location. You cannot download another copy of the private key generated during this operation after completing the operation.
   - **Upload SSH key files**: Select this option to browse or drag and drop .pub files.

- **Paste SSH keys**: Select this option to paste in individual public keys. To paste multiple keys, click **+ Another SSH Key**, and supply a single key for each entry.

7. Click **Save Changes**.

Using the Console to Add SSH Keys After Creating a VM Cluster

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Click **VM Clusters**.
3. Click the name of the VM cluster that you want to add SSH key(s).
4. In the VM Cluster Details page, click **Add SSH Keys**.
5. In the ADD SSH Keys dialog, choose any one of the methods:
   - **Generate SSH key pair:** Select this option if you want the Control Plane to generate public/private key pairs for you.
     Click **Save Private Key** and **Save Public Key** to download and save SSH Key pair.
   - **Upload SSH key files:** Select this option to upload the file that contains SSH Key pair.
   - **Paste SSH keys:** Select this option to paste the SSH key string.
     To provide multiple keys, click **Another SSH Key**. For pasted keys, ensure that each key is on a single, continuous line. The length of the combined keys cannot exceed 10,000 characters.
6. Click **Save Changes**.

Related Topics

- Managing Key Pairs on Linux Instances

Using the Console to Stop, Start, or Reboot a VM Cluster Virtual Machine

Use the console to stop, start, or reboot a virtual machine.

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** and **Compartment** that is associated with the VM cluster that contains the virtual machine that you want to stop, start, or reboot.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that contains the virtual machine that you want to stop, start, or reboot.

   The VM Cluster Details page displays information about the selected VM cluster.
5. In the **Resources** list, click **Virtual Machines**.

   The list of virtual machines is displayed.
6. In the list of nodes, click the **Actions** icon (three dots) for a node, and then click one of the following actions:
   a. **Start:** Restarts a stopped node. After the node is restarted, the **Stop** action is enabled.
   b. **Stop:** Shuts down the node. After the node is stopped, the **Start** action is enabled.
   c. **Reboot:** Shuts down the node, and then restarts it.

Using the Console to Check the Status of a VM Cluster Virtual Machine

Review the health status of a VM cluster virtual machine.

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** and **Compartment** that is associated with the VM cluster that contains the virtual machine that you are interested in.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that contains the virtual machine that you are interested in.

   The VM Cluster Details page displays information about the selected VM cluster.
5. In the **Resources** list, click **Virtual Machines**.

   The list of virtual machines displays. For each virtual machine in the VM cluster, the name, state, and client IP address are displayed.
6. In the node list, find the virtual machine that you are interested in and check its state.

   The color of the icon and the associated text it indicates its status.

   - **Available:** Green icon. The node is operational.
   - **Starting:** Yellow icon. The node is starting because of a start or reboot action in the Console or API.
   - **Stopping:** Yellow icon. The node is stopping because of a stop or reboot action in the Console or API.
   - **Stopped:** Yellow icon. The node is stopped.
   - **Failed:** Red icon. An error condition prevents the continued operation of the virtual machine.

Using the Console to Move a VM Cluster to Another Compartment

To change the compartment that contains your VM cluster on Exadata Cloud Infrastructure, use this procedure.

When you move a VM cluster, the compartment change is also applied to the virtual machines and databases that are associated with the VM cluster. However, the compartment change does not affect any other associated resources, such as the Exadata infrastructure, which remains in its current compartment.

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** and **Compartment** that contains the VM cluster that you want to move.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that you want to move.

   The VM Cluster Details page displays information about the selected VM cluster.
5. Click **Move Resource**.
6. In the resulting dialog, choose the new compartment for the VM cluster, and click **Move Resource**.

To change the VM cluster display name

> **Note:**
>
> This topic only applies to Exadata Cloud Infrastructure instances using the new Exadata Cloud Infrastructuree instance resource model.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
4. In the list of Exadata VM Clusters resources, click the name of the VM Cluster you're interested in
5. On rthe **Infrastructure Details** page, click **More Actions** and **Update Display Name** .
6. In the **Update Display Name** dialog, Enter the **New display name**, and the **current display name** as instructed.
7. Click **Update Display Name**.

Using the Console to Terminate a VM Cluster

Before you can terminate a VM cluster, you must first terminate the databases that it contains.

Terminating a VM cluster removes it from the Cloud Control Plane. In the process, the virtual machines and their contents are destroyed.

> **Note:**
>
> You cannot terminate a VM cluster from an infrastructure with less than 5 storage servers

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** and **Compartment** that contains the VM cluster that you want to terminate.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that you want to terminate.

   The VM Cluster Details page displays information about the selected VM cluster.
5. Click **More Actions**, and then click **Terminate**.
6. In the resulting dialog:
   - Review the message about the backup retention policy
   - Enter the name of the VM cluster
   - Click **Terminate VM Cluster** to confirm the action.

> **Note:**
>
> The database stays in a terminated state with backups listed until all backups are expired.

To view details about private DNS configuration

1. Open the navigation menu. Under **Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Choose the **Compartment** that contains your Exadata infrastructure.
4. Click **VM Clusters**.
5. Click the name of the VM cluster that is configured with a private DNS you want to view.
6. Under the Network section, Private DNS and Private Zone are displayed, if a private DNS is configured.
7. Click the **Private View** name to edit the configuration.

Related Topics

- Using the Console to manage private DNS

To Attach a Virtual IP Address

Attach a Virtual IP address from a VM cluster using this procedure.

1. Open the navigation menu. Under **Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Choose the **Compartment** that contains your Exadata infrastructure.
4. Click **VM Clusters**.
5. Under the **Resources**, click **Virtual IP Address**.
6. Click **Attach Virtual IP Address**.
7. In the Attach Virtual IP Address dialog:
   a. Select a subnet from the **Subnet** drop-down list.
   b. Enter a hostname for the Virtual IP Address in the Virtual IP Address Hostname field.
   c. (Optional) Enter an IP address in the **Virtual IP address** field.
   d. (Optional) You can enter a VIrtual Machine name to be the default attachment in the **Virtual Machine** field.
   e. Click **Attach**.

To Detach a Virtual IP Address

Attach a Virtual IP address from a VM cluster using this procedure.

1. Open the navigation menu. Under **Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.

2. Choose the **Region** that contains your Exadata infrastructure.
3. Choose the **Compartment** that contains your Exadata infrastructure.
4. Click **VM Clusters**.
5. Under the **Resources**, click **Virtual IP Address**.
6. Click the Actions icon (three dots) to the right of the Virtual IP Address you wish to detach.
7. Click **Detach**.
8. In the Detach Virtual IP Address dialog, confirm by entering the VIP Address that you wish to detach and click **Detach**.

Overview of Automatic Diagnostic Collection

By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.

- **Enable Diagnostic Events**

  Allow Oracle to collect and publish critical, warning, error, and information events to you. For more information, see *Database Service Events*.

- **Enable Health Monitoring**

  Allow Oracle to collect health metrics/events such as Oracle Database up/down, disk space usage, and so on, and share them with Oracle Cloud operations. You will also receive notification of some events. For more information, see *Health Metrics*.

- **Enable Incident Logs and Trace Collection**

  Allow Oracle to collect incident logs and traces to enable fault diagnosis and issue resolution. For more information, see *Incident Logs and Trace Files*.

Diagnostics Collection is:

- **Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (all three options).
- **Disabled:** When you choose not to collect diagnostics, health metrics, incident logs, and trace files (all three options).
- **Partially Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (one or two options).

Disabling diagnostic events and health monitoring will only stop the collection and notification of data/events from the time you uncheck the checkboxes tied to the options. However, historical data will not be purged from Oracle Cloud Operations data repositories.

Related Topics

- Database Service Events
  The Database Service emits events, which are structured messages that indicate changes in resources.
- Incident Logs and Trace Files
  This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.
- Health Metrics
  Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.

- To create a cloud VM cluster resource
  Create a VM cluster in an Exadata Cloud Infrastructure instance.
- Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.

Incident Logs and Trace Files

This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.

> **Note:**
>
> - Oracle will create a service request (SR) against the infrastructure Customer Support Identifier (CSI) when an issue is detected and needs customer interaction to resolve.
> - The customer's Oralce CLoud Infrastructure tenancy admin email will be used as the CSI contact to create SR and attach logs to it. Ensure tenancy admin is added as a CSI contact in My Oracle Support (MOS).

Oracle Trace File Analyze (TFA) Component Driven Logs Collections

The directories are generally assigned to a component and that component can then be used to guide TFA to the files it needs to collect, for example, requesting the CRS component would tell TFA to look at directories mapped to the CRS component and find files that match the required collection time frame.

> **Note:**
>
> If have previously opted in for incident log and trace file collection and decide to opt out when Oracle Cloud operations run a log collection job, then the job will run its course and will not cancel. Future log collections won't happen until you opt-in again to the incident logs and trace file collection option.
>
> TFA is shipped with scripts that run when a particular component is requested, for example, for CRS component, `crscollect.pl` will run a number of `crsctl` commands and gather the input. By default, TFA does not redact collected logs.

Table 5-2 Oracle Trace File Analyze (TFA) Component Driven Logs Collections

Component	Script	Files/Directories
`OS`: Operating system logs	`oscollect.pl`	• `/var/log/messages` • OSWatcher archive • **Exadata Only:** ExaWatcher archive `/opt/oracle.ExaWatcher/archive/`

Table 5-2 (Cont.) Oracle Trace File Analyze (TFA) Component Driven Logs Collections

Component	Script	Files/Directories
CRS: Grid Infrastructure and cluster logs	crscollect.pl	• /etc/oracle • GIHOME/crf/db/HOSTNAME1 • GIHOME/crs/log • GIHOME/css/log • GIHOME/cv/log • GIHOME/evm/admin/log • GIHOME/evm/admin/logger • GIHOME/evm/log • GIHOME/log/-/client • GIHOME/log/HOSTNAME1 • GIHOME/log/HOSTNAME1/admin • GIHOME/log/HOSTNAME1/client • GIHOME/log/HOSTNAME1/crflogd • GIHOME/log/HOSTNAME1/crfmond • GIHOME/log/HOSTNAME1/crsd • GIHOME/log/HOSTNAME1/cssd • GIHOME/log/HOSTNAME1/ctssd • GIHOME/log/HOSTNAME1/diskmon • GIHOME/log/HOSTNAME1/evmd • GIHOME/log/HOSTNAME1/gipcd • GIHOME/log/HOSTNAME1/gnsd • GIHOME/log/HOSTNAME1/gpnpd • GIHOME/log/HOSTNAME1/mdnsd • GIHOME/log/HOSTNAME1/ohasd • GIHOME/log/HOSTNAME1/racg • GIHOME/log/HOSTNAME1/srvm • GIHOME/log/HOSTNAME1/xag

Table 5-2 (Cont.) Oracle Trace File Analyze (TFA) Component Driven Logs Collections

Component	Script	Files/Directories
		• GIHOME/log/diag/asmtool • GIHOME/log/diag/clients • GIHOME/log/procwatcher/PRW_SYS_HOSTNAME1 • GIHOME/network/log • GIHOME/opmn/logs • GIHOME/racg/log • GIHOME/scheduler/log • GIHOME/srvm/log • GRIDBASE/crsdata/@global/cvu • GRIDBASE/crsdata/HOSTNAME1/core • GRIDBASE/crsdata/HOSTNAME1/crsconfig • GRIDBASE/crsdata/HOSTNAME1/crsdiag • GRIDBASE/crsdata/HOSTNAME1/cvu • GRIDBASE/crsdata/HOSTNAME1/evm • GRIDBASE/crsdata/HOSTNAME1/output • GRIDBASE/crsdata/HOSTNAME1/ovmmwallets • GRIDBASE/crsdata/HOSTNAME1/scripts • GRIDBASE/crsdata/HOSTNAME1/trace • GRIDBASE/diag/crs/-/crs/cdump • GRIDBASE/diag/crs/HOSTNAME1/crs/cdump • GRIDBASE/diag/crs/HOSTNAME1/crs/incident • GRIDBASE/diag/crs/HOSTNAME1/crs/trace

Table 5-2 (Cont.) Oracle Trace File Analyze (TFA) Component Driven Logs Collections

Component	Script	Files/Directories
`Database`: Oracle Database logs	No DB Specific Script - runs `opatch lsinventory` for the `ORACLE_HOME` the DB runs from TFA will run ipspack based on the time range for certain DB incidents.	• `ORACLE_BASE/diag/rdbms/<dbname>/<instance_name>/cdump` • `ORACLE_BASE/diag/rdbms/<dbname>/<instance_name>/trace` • `ORACLE_BASE/diag/rdbms/<dbname>/<instance_name>/incident`

Cloud Tool Logs

- **Creg files:** `/var/opt/oracle/creg/*.ini` files with masked sensitive info
- **Cstate file:** `/var/opt/oracle/cstate.xml`
- **Database related tooling logs:**

  If `dbName` specified, `/var/opt/oracle/log/<dbName>`, else collect logs for all databases `/var/opt/oracle/log/`

  If `dbName` specified, `/var/opt/oracle/dbaas_acfs/log/<dbName>`, else collect logs for all databases `/var/opt/oracle/log/<dbName>`
- **Database env files:** If `dbName` specified, `/home/oracle/<dbName>.env`, else collect logs for all databases `/home/oracle/*.env`
- **Pilot logs:** `/home/opc/.pilotBase/logs`
- **List of log directories:**
  - `/var/opt/oracle/log`
  - `/var/opt/oracle/dbaas_acfs/log`
  - `/var/opt/oracle/dbaas_acfs/dbsystem_details`
  - `/var/opt/oracle/dbaas_acfs/job_manager`
  - `/opt/oracle/dcs/log`

DCS Agent Logs

- `/opt/oracle/dcs/log/`

Tooling-Related Grid Infrastructure/Database Logs

- **Grid Infrastructure:** `GI_HOME/cfgtoollogs`
- **Database alertlog:** `/u02/app/oracle/diag/rdbms/*/*/alert*.log`

Related Topics

- Overview of Automatic Diagnostic Collection
  By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
- Health Metrics
  Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.
- To create a cloud VM cluster resource
  Create a VM cluster in an Exadata Cloud Infrastructure instance.
- Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.

Health Metrics

Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.

> **Note:**
>
> Oracle may add more metrics in the future, but if you have already chosen to collect metrics, you need not update your opt-in value. It will remain enabled/disabled based on your current preference.

Guest VM Health Metrics List - Database Metrics

Table 5-3 Guest VM Health Metrics List - Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Interval	Collection Frequency	Description
`CpuUtilization`	CPU Utilization	Percentage	Mean	One minute	Five minutes	The CPU utilization is expressed as a percentage, which is aggregated across all consumer groups. The utilization percentage is reported with respect to the number of CPUs the database is allowed to use, which is two times the number of OCPUs.
`StorageUtilization`	Storage Utilization	Percentage	Mean	One hour	One hout	The percentage of provisioned storage capacity currently in use. Represents the total allocated space for all tablespaces.
`BlockChanges`	DB Block Changes	Changes per second	Mean	One minute	Five minutes	The Average number of blocks changed per second.
`ExecuteCount`	Execute Count	Count	Sum	One minute	Five minutes	The number of user and recursive calls that executed SQL statements during the selected interval.

Table 5-3 (Cont.) Guest VM Health Metrics List - Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Interval	Collection Frequency	Description
CurrentLogons	Current Logons	Count	Sum	One minute	Five minutes	The number of successful logons during the selected interval.
TransactionCount	Transaction Count	Count	Sum	One minute	Five minutes	The combined number of user commits and user rollbacks during the selected interval.
UserCalls	User Calls	Count	Sum	One minute	Five minutes	The combined number of logons, parses, and execute calls during the selected interval.
ParseCount	Parse Count	Count	Sum	One minute	Five minutes	The number of hard and soft parses during the selected interval.
StorageUsed	Storage Space Used	GB	Max	One hour	One hour	Total amount of storage space used by the database at the collection time.
StorageAllocated	Storage Space Allocated	GB	Max	One hour	One hour	Total amount of storage space allocated to the database at the collection time.

Table 5-3 (Cont.) Guest VM Health Metrics List - Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Interval	Collection Frequency	Description
StorageUsedByTablespace	Storage Space Used By Tablespace	GB	Max	One hour	One hour	Total amount of storage space used by tablespace at the collection time. In the case of container databases, this metric provides root container tablespaces.
StorageAllocatedByTablespace	Allocated Storage Space By Tablespace	GB	Max	One hour	One hour	Total amount of storage space allocated to the tablespace at the collection time. In the case of container databases, this metric provides root container tablespaces.
StorageUtilizationByTablespace	Storage Space Utilization By Tablespace	Percentage	Mean	One hour	One hour	This indicates the percentage of storage space utilized by the tablespace at the collection time. In the case of container databases, this metric provides root container tablespaces.

Guest VM Health Metrics List - Non-Database Metrics

Table 5-4 Guest VM Health Metrics List - Non-Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Collection Frequency	Description
`ASMDiskgroup Utilization`	ASM Diskgroup Utilization	Percentage	Max	10 minutes	Percentage of usable space used in a Disk Group. Usable space is the space available for growth. DATA disk group stores our Oracle database files. RECO disk group contains database files for recovery such as archives and flashback logs.
`FilesystemUtilization`	Filesystem Utilization	Percentage	Max	One minute	Percent utilization of provisioned filesystem.
`CpuUtilization`	CPU Utilization	Percentage	Mean	One minute	Percent CPU utilization.
`MemoryUtilization`	Memory Utilization	Percentage	Mean	One minute	Percentage of memory available for starting new applications, without swapping. The available memory can be obtained via the following command: `cat /proc/meminfo`.
`SwapUtilization`	Swap Utilization	Percentage	Mean	One minute	Percent utilization of total swap space.
`LoadAverage`	Load Average	Number	Mean	One minute	System load average over 5 minutes.
`NodeStatus`	Node Status	Integer	Mean	One minute	Indicates whether the host is reachable.

Table 5-4 (Cont.) Guest VM Health Metrics List - Non-Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Collection Frequency	Description
`OcpusAllocated`	OCPU Allocated	Integer	Max	One minute	The number of OCPUs allocated.

Related Topics

- Overview of Automatic Diagnostic Collection
  By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
- Incident Logs and Trace Files
  This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.
- To create a cloud VM cluster resource
  Create a VM cluster in an Exadata Cloud Infrastructure instance.
- Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.

Using the API to Manage Exadata Cloud Infrastructure Instance

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.

Use these API operations to manage Exadata Cloud Infrastructure instance components.

Cloud Exadata infrastructure resource (new resource model):

- ListCloudExadataInfrastructures
- GetCloudExadataInfrastructure
- ChangeCloudExadataInfrastructureCompartment
- UpdateCloudExadataInfrastructure
- DeleteCloudExadataInfrastructure

Cloud VM cluster (new resource model)

- ListCloudVmClusters
- GetCloudVmCluster
- ChangeCloudVmClusterCompartment
- UpdateCloudVmCluster
- DeleteCloudVmCluster

DB systems (old resource model):

- ListDbSystems
- GetDbSystem
- ChangeDbSystemCompartment
- UpdateDbSystem
- TerminateDbSystem

Virtual machines nodes (all Exadata Cloud Infrastructure instances):

- DbNodeAction
- ListDbNodes
- GetDbNode

Troubleshooting Virtual Machines Using Console Connections

You can troubleshoot malfunctioning virtual machines using console connections. For example, a previously working Guest VM stops responding.

> **Note:**
>
> Exadata System Software 23.1.13 is the minimum required version. Also, make sure to review all prerequisites stated below, including setting a password for either the `opc` or the `root` user. Failure to make necessary changes to meet these requirements in advance will result in the inability to urgently connect to the serial console when the need arises when the VM is not otherwise accessible.

To connect to a running instance for administration and general use, use a Secure Shell (SSH). For more information, see Connecting to a Virtual Machine with SSH.

To make an SSH connection to the serial console, follow these configuration steps.

1. Ensure that you have the correct permissions.
2. Complete the prerequisites, including creating your SSH key pair (in case you don't have one yet).
3. Create the Virtual Machine Serial Console.
4. Connect to the serial console via SSH.

To check the DB server version installed, follow these steps:

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Click **Exadata Infrastructure** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
4. Click the name of the infrastructure that you are interested in.
5. In the resulting Infrastructure Details page, go to the Version section to find the DB Server version installed.

Required IAM Policies

An administrator must grant you secure access to the virtual machine console on the Oracle Exadata Database Service on Dedicated Infrastructure system through an IAM policy.

This access is required whether you're using the Console or the REST API with an SDK, CLI, or other tools. If you get a message that you don’t have permission or are unauthorized, verify with your administrator what type of access you have and which compartment to work in.

To create virtual machine console connections, an administrator needs to grant user access to read and manage virtual machine console connections through an IAM policy. The resource name for virtual machine console connections is `dbnode-console-connection`. The resource name for the virtual machine is `db-nodes`. The following policies grant users the ability to create virtual machine console connections:

```
Allow group <group_name> to manage dbnode-console-connection in tenancy
Allow group <group_name> to read db-nodes in tenancy
```

Prerequisites

You must install an SSH client and create SSH key pairs.

Install an SSH Client and a Command-line Shell (Microsoft Windows)

Microsoft Windows does not include an SSH client by default. If you are connecting from a Windows client, you need to install an SSH client. You can use PuTTY plink.exe with Windows PowerShell or software that includes a version of OpenSSH such as:

- Git for Windows
- Windows Subsystem for Linux

The instructions in this topic frequently use PuTTY and Windows PowerShell.

If you want to make the console connection from Windows with Windows PowerShell, PowerShell might already be installed on your Windows operating system. If not, follow the steps at the link. If you are connecting to the instance from a Windows client using PowerShell, plink.exe is required. plink.exe is the command link connection tool included with PuTTY. You can install PuTTY or install plink.exe separately. For installation information, see http://www.putty.org.

Create SSH Key Pairs

To create the secure console connection, you need an SSH key pair. The method to use for creating key pairs depends on your operating system. When connecting to the serial console, you must use an RSA key. The instructions in this section show how to create an RSA SSH key pair.

Create the SSH key Pair for Linux

If you're using a UNIX-style system, you probably already have the `ssh-keygen` utility installed. To determine whether the utility is installed, type `ssh-keygen` on the command line. If the utility isn't installed, you can download OpenSSH for UNIX from http://www.openssh.com/portable.html and install it.

1. Open a shell or terminal for entering the commands.
2. At the prompt, enter `ssh-keygen` and provide a name for the key when prompted. Optionally, include a passphrase.
   The keys will be created with the default values: RSA keys of 2048 bits.

   Alternatively, you can type a complete `ssh-keygen` command, for example:

   ```
   ssh-keygen -t rsa -N "" -b 2048 -C "<key_name>" -f <path/root_name>
   ```

Argument	Description
`-t rsa`	Use the RSA algorithm.
`-N "<passphrase>"`	A passphrase to protect the use of the key (like a password). If you don't want to set a passphrase, don't enter anything between the quotes. A passphrase is not required. You can specify one as a security measure to protect the private key from unauthorized use. If you specify a passphrase, when you connect to the instance you must provide the passphrase, which typically makes it harder to automate connecting to an instance.
`-b 2048`	Generate a 2048-bit key. You don't have to set this if 2048 is acceptable, as 2048 is the default. A minimum of 2048 bits is recommended for SSH-2 RSA.
`-C "<key_name>"`	A name to identify the key.
-f *<path/root_name>*	The location where the key pair will be saved and the root name for the files.

Create the SSH Key Pair for Windows Using PuTTY

If you are using a Windows client to connect to the instance console connection, use an SSH key pair generated by PuTTY.

> **Note:**
>
> Ensure that you are using the latest version of PuTTY, see http://www.putty.org.

1. Find `puttygen.exe` in the PuTTY folder on your computer, for example, `C:\Program Files (x86)\PuTTY`. Double-click `puttygen.exe` to open it.
2. Specify a key type of SSH-2 RSA and a key size of 2048 bits:
   - In the **Key** menu, confirm that the default value of **SSH-2 RSA key** is selected.
   - For the **Type of key to generate**, accept the default key type of **RSA**.
   - Set the **Number of bits in a generated key** to 2048 if not already set.
3. Click **Generate**.
4. To generate random data in the key, move your mouse around the blank area in the PuTTY window.
   When the key is generated, it appears under **Public key for pasting into OpenSSH authorized_keys file**.
5. A **Key comment** is generated for you, including the date and timestamp. You can keep the default comment or replace it with your own more descriptive comment.
6. Leave the **Key passphrase** field blank.
7. Click **Save private key**, and then click **Yes** in the prompt about saving the key without a passphrase.
   The key pair is saved in the PuTTY Private Key (PPK) format, which is a proprietary format that works only with the PuTTY tool set.

   You can name the key anything you want but use the `ppk` file extension. For example, `mykey.ppk`.
8. Select all of the generated keys that appear under **Public key** for pasting into OpenSSH `authorized_keys` file, copy it using **Ctrl + C**, paste it into a text file, and then save the file in the same location as the private key.

   > **Note:**
   >
   > Do not use the **Save public key** option because it does not save the key in the OpenSSH format.

   You can name the key anything you want, but for consistency, use the same name as the private key and a file extension of `pub`. For example: `mykey.pub`.
9. Write down the names and locations of your public and private key files. You need the public key when creating an instance console connection. You need the private key to connect to the instance console connection using PuTTY. For example: `$HOME\Documents\mykey.ppk`.

To create a connection using the SSH key pair generated using PuTTY

For more information about generating SSH key pairs, see Create the SSH Key Pair for Windows Using PuTTY

Do the following on the Create serial console access window:

1. Paste the SSH Key generated from OpenSSH format or choose **Upload SSH key file** and provide the path of the public key saved at step 8 in Create the SSH Key Pair for Windows Using PuTTY.
2. Once the connection is **Active**, click **Copy serial console connection for Windows**.
3. Paste the connection string copied from the previous step into a text file.

4. In the text file, replace `<PATH_FILE_PUTTY_PRIVATE.ppk>` to point to your PuTTY Private Key (PPK) file path on your computer. For example, if you have saved `.ppk` file at `$HOME\Documents\mykey.ppk`.
5. Paste the modified connection string into the PowerShell window, and then press **Enter** to connect to the console.

Sign in to a Virtual Machine From the Serial Console

If you want to sign in to a virtual machine using a virtual machine console connection, you can use Secure Shell (SSH) connection to sign in. If you want to sign in with a username and password, you need a user account with a password. Oracle Exadata Cloud does not set a default password for the `opc` or `root` users. Therefore, if you want to sign in as the `opc` or `root` user, you need to create a password for the `opc` or `root` user. Otherwise, add a different user with a password and sign in as that user. This should be completed in advance, before a potential situation that might require you to log in to the serial console.

Connect Through Firewalls

If the client you will use to access the serial console is behind a firewall, you must ensure that this client can reach the required endpoint to access the serial console of the virtual machine. The client system connecting to the serial console must be able to reach the serial console server (for example, `vm-console-ad1.exacs.us-ashburn-1.oci.oraclecloud.com`) over SSH using port 443, directly or through a proxy.

Create the Virtual Machine Serial Console Connection

Before you can make a local connection to the serial console, you need to create the virtual machine console connection.

Virtual machine console connections are limited to one client at a time. If the client fails, the connection remains active for approximately five minutes. During this time, no other client can connect. After five minutes, the connection is closed, and a new client can connect. During the five-minute timeout, any attempt to connect a new client fails with the following message:

```
channel 0: open failed: administratively prohibited: console access is
limited to one connection at a time
```

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Click the VM Cluster that you're interested in.
3. In the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
   Under **Resources**, **Console connection** is selected by default.
4. Click **Create serial console access**.
5. In the resulting Create serial console access window, you have three options for adding the SSH key
   - **Generate a key pair for me**: You can have Oracle Cloud Infrastructure generate an SSH key pair to use. If you are using PowerShell or PuTTY to connect to the instance from a Windows client, you cannot use the generated SSH key pair without first converting it to a `.ppk` file.

- **Upload public key file**: Browse to a public key file on your computer. If you followed the steps in Creating SSH Key Pairs in the Prerequisites section to create a key pair, use this option to navigate to the `.pub` file.
- **Paste public key**: Paste the content of your public key file into the text box.

6. Click **Create console connection**.
   When the console connection has been created and is available, the state changes to **Active**.

Make an SSH Connection to the Serial Console

After you create the console connection for the virtual machine, you can connect to the serial console using a Secure Shell (SSH) connection. When making an SSH connection to the serial console, you must use an RSA key. You can use the same SSH key for the serial console that was used when you launched the instance, or you can use a different SSH key.

When you are finished with the serial console and have terminated the SSH connection, you should delete the serial console connection. If you do not disconnect from the session, Oracle Cloud Infrastructure terminates the serial console session after 24 hours and you must reauthenticate to connect again.

Validate Server Host Keys

When you first connect to the serial console, you're prompted to validate the fingerprint of the server host key. The fingerprint of the server host key is the SHA256 hash of the server host's public SSH key. The server SSH handshake response is signed with the associated private key. Validating the server host key's fingerprint protects against potential attacks.

When you make a manual connection to the serial console, the fingerprint of the server host key is not automatically validated. To manually validate the fingerprint, compare the fingerprint value displayed in the Oracle Cloud Infrastructure Console to the value of the RSA key fingerprint that appears in the terminal when you connect.

To find the fingerprint of the server host key in the Console, on the Virtual Machine details page, under **Resources**, click **Console connection**. The table displays the fingerprint of the server host key. The fingerprint in the Console should match the value of the **RSA key fingerprint** shown in the terminal when you connect to the serial console.

The server host keys are periodically rotated for security purposes. Key rotation reduces the risk posed when keys are compromised by limiting the amount of data encrypted or signed by one key version. When your key is rotated and you try to connect to the serial console, a warning appears indicating a potential attack. The warning includes an `Host key verification failed` error and a line number in your `.ssh/known_hosts` file. Delete that line in your `.ssh/known_hosts` file and then reconnect to the serial console. You are then prompted to accept a new server host key fingerprint.

Connect from Mac OS X and Linux Operating Systems

Use an SSH client to connect to the serial console. Mac OS X and most Linux and UNIX-like operating systems include the SSH client OpenSSH by default.

To connect to the serial console using OpenSSH on Mac OS X or Linux:

1. Open the navigation menu. Under **Oracle Database**, click **Exadata Database Service on Cloud@Customer**.
2. Click the VM Cluster that you're interested in.

3. In the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
4. On the Virtual Machine details page in the Oracle Cloud Infrastructure Console, under **Resources**, click **Console connection**.
5. Click the Actions menu (three dots), and then click **Copy serial console connection for Linux/Mac**.
6. Paste the connection string into a terminal window on a Mac OS X or Linux system, and then press Enter to connect to the console.
   If you are not using the default SSH key or ssh-agent, modify the serial console connection string to include the identity file flag, `-i` to specify the private key portion for the SSH key to use, for example, `id_rsa`. Specify this flag for both the SSH connection and the SSH Proxy Command, as shown in the following line:

   ```
   ssh -i /<path>/<ssh_key> -o ProxyCommand='ssh -i /<path>/<ssh_key> -W
   %h:%p -p 443...
   ```

7. If prompted, validate and accept the fingerprint of the server host key.
   If you have previously accepted a fingerprint for the server host key but the key has been rotated, a warning appears indicating a potential attack. The warning includes an `Host key verification failed` error and a line number in your `.ssh/known_hosts` file. Delete the specified line in your `.ssh/known_hosts` file and then reconnect to the serial console. Validate and accept the new server host key fingerprint.
8. Press **Enter** again to activate the console.
   If the connection is active, a message appears in the console:

   ```
   ================================================
   IMPORTANT: You are now connected to the serial console for this VM. This should be
   used in emergency situations only.
   See product documentation for more details and alternative connectivity options for
   normal operations
   ================================================
   ```

9. Reboot your virtual machine.
   You do not need to enter a user name or password. If the Virtual Machine is functional and the connection is active, the serial output appears in your console. If the serial output does not appear in the console, the Guest VM operating system is not booting.

   For more troubleshooting options, see *Troubleshooting Virtual Machines from Guest VM Console Connections on Linux Operating Systems*.

   a. Go to the ExaDB-C@C VM Cluster Details page.
   b. Under **Resources**, click **Virtual Machines**.
   c. Select **Reboot** from the Actions menu (three dots) for the virtual machine that you want to reboot.

Connect from Windows Operating Systems

The steps to connect to the serial console from Microsoft Windows PowerShell are different from the steps for OpenSSH. The following steps do not work in the Windows terminal.

If you are connecting to the instance from a Windows client using PowerShell, `plink.exe` is required. `plink.exe` is the command link connection tool included with PuTTY. You can install PuTTY or install `plink.exe` separately. For more information, see *Installing an SSH Client and a Command-line Shell (Windows)*.

To connect to the serial console on Microsoft Windows:

1. On the Virtual Machine details page in the Oracle Cloud Infrastructure Console, under **Resources**, click **Console connection**.
2. Click the Actions menu (three dots).
   Depending on which SSH client you are using, do one of the following:
   - If you are using Windows PowerShell, click **Copy serial console connection for Windows**.
   - If you are using OpenSSH, click **Copy serial console connection for Linux/Mac**.

   > **Note:**
   >
   > The copied connection string for Windows contains the parameter `-i` specifying the location of the private key file. The default value for this parameter in the connection string references an environment variable that might not be configured on your Windows client, or it might not represent the location where the private key file is saved. Verify the value specified for the `-i` parameter and make any required changes before proceeding to the next step.

3. Paste the connection string copied from the previous step into a text file so that you can add the file path to the private key file.
4. In the text file, replace `$env:homedrive$env:homepath\oci\console.ppk` with the file path to the `.ppk` file on your computer. This file path appears twice in the string. Replace it in both locations.
5. Paste the modified connection string into the PowerShell window or your OpenSSH client, and then press **Enter** to connect to the console.
6. If prompted, validate and accept the fingerprint of the server host key.
   If you have previously accepted a fingerprint for the server host key, but the key has been rotated, a warning appears indicating a potential attack. The warning includes a Host key verification failed error and a line number in your `.ssh/known_hosts` file. Delete the specified line in your `.ssh/known_hosts` file and then reconnect to the serial console. Validate and accept the new server host key fingerprint.
7. Press **Enter** again to activate the console.
8. Reboot your virtual machine.
   You do not need to enter a user name or password. If the Virtual Machine is functional and the connection is active, the serial output appears in your console. If the serial output does not appear in the console, the Guest VM operating system is not booting.

   For more troubleshooting options, see Troubleshooting Virtual Machines from Guest VM Console Connections.

   a. Go to the ExaDB-C@C VM Cluster Details page.
   b. Under **Resources**, click **Virtual Machines**.
   c. Select **Reboot** from the Actions menu (three dots) for the virtual machine that you want to reboot.

To create a connection using the SSH key pair generated using the OCI Console:

Do the following on the Create serial console access window:

1. Click **Generate a key pair for me**.
2. Click **Save Private Key**.
3. Click **Create console connection**.

> **Note:**
>
> Ensure that you are using the latest version of PuTTY, see http://www.putty.org.

4. Find `puttygen.exe` in the PuTTY folder on your computer, for example, `C:\Program Files (x86)\PuTTY. Double-click puttygen.exe` to open it.
5. On the PuTTY Key Generator, click the **Conversions** menu and then click **Import**.
6. On the Windows Explorer, select OCI Console generated SSH key (step 1) and then click **Open**.
   PuTTY imports the key and displays information about the key on the PuTTY Key Generator window.
7. Click **Save private key**.
8. Click **Yes** when prompted about saving the key without a passphrase.
   The key pair is saved in the PuTTY Private Key (PPK) format, which is the proprietary format that works only with the PuTTY tool set.

   You can name the key anything you want but use the `.ppk` file extension. For example, `$HOME\Desktop\key-vm-console.ppk`.
9. Use a text editor to change the command to point to your PuTTY Private Key (PPK) path. Replace *<PATH_FILE_PUTTY_PRIVATE.ppk>* to point to your PuTTY Private Key (PPK) file path on your computer. For example, if you have saved `.ppk` file at `$HOME\Desktop\key-vm-console.ppk`.
10. Paste the modified connection string into the PowerShell window, and then press **Enter** to connect to the console.

To convert a generated .key private key file:

1. Open PuTTYgen.
2. Click **Load**, and select the private key generated when you created the instance.
   The extension for the key file is `.key`.
3. Click **Save private key**.
4. Specify a name for the key.
   The extension for the new private key is `.ppk`.
5. Click **Save**.

Troubleshooting Virtual Machines from Guest VM Console Connections on Linux Operating Systems

After you are connected with an instance console connection, you can perform various tasks, such as:

- Edit system configuration files.
- Add or reset the SSH keys for the `opc` user.
- Reset the password for the `opc` user.

These tasks require you to boot into a Bash shell in maintenance mode.

To boot into maintenance mode

> **Note:**
>
> **Default user and password:**
>
> - **Account**: Grub boot loader
> - **Username**: root
> - **Default Password**: sos1Exadata
> - **Account Type**: Operating system user
>
> For more information, see *Default User Accounts for Oracle Exadata*.

1. Reboot the VM from the VM Cluster.
2. For virtual machines running Oracle Linux 7.x or Oracle Linux 8.x, when the reboot process starts, switch back to the terminal window, and you see Console messages start to appear in the window. As soon as the **GRUB boot menu** appears, use the **up/down arrow key** to stop the automatic boot process, enabling you to use the boot menu.
3. In the boot menu, highlight the top item in the menu, and press **e** to edit the boot entry.
4. In edit mode, use the **down arrow key** to scroll down through the entries until you reach the line that starts with **linux16**.
5. At the end of that line, add the following:

   ```
   init=/bin/bash
   ```

6. Reboot the instance from the terminal window by entering the keyboard shortcut **CTRL+X**. When the instance has rebooted, you see the Bash shell command-line prompt, and you can proceed with the following procedures.

To edit the system configuration files

1. From the Bash shell, run the following command to load the SElinux policies to preserve the context of the files you are modifying:

   ```
   /usr/sbin/load_policy -i
   ```

2. Run the following command to remount the root partition with read/write permissions:

   ```
   /bin/mount -o remount, rw /
   ```

3. Edit the configuration files as needed to try to recover the instance.
4. After you have finished editing the configuration files, to start the instance from the existing shell, run the following command:

   ```
   exec /usr/lib/systemd/systemd
   ```

Alternatively, to reboot the instance, run the following command:

```
/usr/sbin/reboot -f
```

To add or reset the SSH key for the opc user

1. From the Bash shell, run the following command to load the SElinux policies to preserve the context of the files you are modifying:

   ```
   /usr/sbin/load_policy -i
   ```

2. Run the following command to remount the root partition with read/write permissions:

   ```
   /bin/mount -o remount, rw /
   ```

3. From the Bash shell, run the following command to change to the SSH key directory for the `opc` user:`cd ~opc/.ssh`
4. Include your public key entry to the `authorized_keys` file.

   > **Note:**
   >
   > You can edit the file and remove your previous key if you want to. However, make sure to keep the cloud automation keys to prevent cloud automation from breaking.
   >
   > ```
   > echo '<contents of public key file>' >> authorized_keys
   > ```

5. Restart the instance by running the following command:

   ```
   /usr/sbin/reboot -f
   ```

To reset the password for the opc user

1. From the Bash shell, run the following command to load the SElinux policies to preserve the context of the files you are modifying.
   This step is necessary to sign in to your instance using SSH and the Console.

   ```
   /usr/sbin/load_policy -i
   ```

2. Run the following command to remount the root partition with read/write permissions:

   ```
   /bin/mount -o remount, rw /
   ```

3. Run the following command to reset the password for the `opc` user:

   ```
   sudo passwd opc
   ```

4. Restart the instance by running the following command:

   ```
   sudo reboot -f
   ```

> **Note:**
>
> Setting a `root` password would be an acceptable alternative to setting an `opc` password.

Exiting the Virtual Machine Serial Console Connection

To exit the serial console connection

When using SSH, the `~` character at the beginning of a new line is used as an escape character.

1. To exit the serial console, enter:

   ```
   ~.
   ```

2. To suspend the SSH session, enter:

   ```
   ~^z
   ```

   The `^` character represents the **CTRL** key.

3. To see all the SSH escape commands, enter:

   ```
   ~?
   ```

To delete the serial console connection for a Virtual Machine

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Click the VM Cluster that you're interested in.
3. In the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
   Under **Resources**, **Console connection** is selected by default.
4. Click the Actions menu, and then click **Delete**. Confirm when prompted.

Manage Software Images

- Using Software Images in Oracle Cloud Infrastructure
- Using a Software Image with an Exadata Cloud Infrastructure Instance
  Create, save, and reuse a Software Image.
- Using the Console for Software Images
- Using the API to manage database software images
  Use these API operations to manage database software images:

Using Software Images in Oracle Cloud Infrastructure

- Creation and Storage of Software Images
  Software images are resources within your tenancy that you create before provisioning or updating a DB system, Exadata Cloud Infrastructure instance, Database Home, database, or Grid Infrastructure.
- Using the OPatch lsinventory Command to Verify the Patches Applied to an Oracle Home
  OPatch utility enables you to apply the interim patches to Oracle Database Home or Oracle Grid Infrastructure Home. You can find the `opatch` utility in the `$ORACLE_HOME/Opatch` directory.

Creation and Storage of Software Images

Software images are resources within your tenancy that you create before provisioning or updating a DB system, Exadata Cloud Infrastructure instance, Database Home, database, or Grid Infrastructure.

There are two types of software image resources:

- **Grid Infrastructure software image:** Grid Infrastructure software images are resources containing Oracle Grid Infrastructure software used to update Oracle Grid Infrastructure. Grid Infrastructure software images are either Oracle-published software releases or custom software images created by the customer that include the desired Grid Infrastructure release updates (GIRU) and additional one-off (interim) patches.
- **Database software image:** Database software images are resources containing Oracle Database software used to provision and update Oracle Databases and Oracle Database Homes. Database software images are either Oracle-published software releases or custom software images created by the customer that include the desired Database release updates (DBRU) and additional one-off (interim) patches.

There is no limit on the number of software images you can create in your tenancy, and you can create your images with any Oracle Database software or Oracle Grid Infrastructure version and update supported in Oracle Cloud Infrastructure.

Software images are automatically stored in Oracle-managed Object Storage and can be viewed and managed in the Oracle Cloud Infrastructure Console. Software images are regional-level resources but they can be accessed from any region within your tenancy.

Note: The software images incur Object Storage usage costs.

Using the OPatch lsinventory Command to Verify the Patches Applied to an Oracle Home

OPatch utility enables you to apply the interim patches to Oracle Database Home or Oracle Grid Infrastructure Home. You can find the `opatch` utility in the `$ORACLE_HOME/Opatch` directory.

Using the `lsinventory` command provided by OPatch, you can create a file that lists the interim patches applied to an Oracle Database Home or Oracle Grid Infrastructure Home. This file can then be uploaded to the OCI Console during the creation of a custom software Image to add the exact set of patches used by the source Oracle Database Home or Oracle Grid Infrastructure Home to the list of patches included in the software image. You can find the `opatch` utility in the `$ORACLE_HOME/Opatch` directory. The following example shows how to use the `lsinventory` command to create the `lsinventory` file:

1. Run the `opatch lsinventory` command to get the list of interim patches applied.

```
$ORACLE_HOME/OPatch/opatch lsinventory
Oracle Interim Patch Installer version 12.2.0.1.21
Copyright (c) 2021, Oracle Corporation. All rights reserved.

Oracle Home : /u02/app/oracle/product/19.0.0.0/dbhome_2
Central Inventory : /u01/app/oraInventory
from : /u02/app/oracle/product/19.0.0.0/dbhome_2/oraInst.loc
OPatch version : 12.2.0.1.21
OUI version : 12.2.0.7.0
Log file location : /u02/app/oracle/product/19.0.0.0/dbhome_2/cfgtoollogs/
opatch/opatch2021-01-21_09-22-45AM_1.log

Lsinventory Output file location : /u02/app/oracle/product/19.0.0.0/
dbhome_2/cfgtoollogs/opatch/lsinv/lsinventory2021-01-21_09-22-45AM.txt


Oracle Interim Patch Installer version 12.2.0.1.41
Copyright (c) 2024, Oracle Corporation.  All rights reserved.


Oracle Home       : /u01/app/oracle/product/19.0.0.0/gridhome_1
Central Inventory : /u01/app/oraInventory
   from           : /u01/app/oracle/product/19.0.0.0/gridhome_1/oraInst.loc
OPatch version    : 12.2.0.1.41
OUI version       : 12.2.0.7.0
Log file location : /u01/app/oracle/product/19.0.0.0/gridhome_1/
cfgtoollogs/opatch/opatch2024-04-19_19-24-22PM_1.log

Lsinventory Output file location : /u01/app/oracle/product/19.0.0.0/
gridhome_1/cfgtoollogs/opatch/lsinv/lsinventory2024-04-19_19-24-22PM.txt
```

2. Use the `lsinventory` output file to extract the additional interim patches applied to a specific Oracle Database Home or Oracle Grid Infrastructure Home.

Using a Software Image with an Exadata Cloud Infrastructure Instance

Create, save, and reuse a Software Image.

Creating a Software Image enables you to:

- Create custom Database and Grid Infrastructure images based on Software Images, RU, and one-off (interim) patches.
- Save a custom image automatically to Object Storage as a resource.
- Provision an Oracle Database home or Oracle Database with the desired RU and one-off (interim) patches.
- Update the Database Home and Grid Infrastructure using the Software Image.
- Clone Software Image to another service in the Data Guard creation process.

> **Note:**
>
> The Software Images are created and managed by the customer and they are available for use until explicitly deleted.

Using the Console for Software Images

- To view the list of software images
- To create a database software image
- To create a Grid Infrastructure software image
- To create a database software image from a Database Home
- To view the image details of a software image
- To move a software image to a different compartment
- To update database software using custom database software image
  Use the following instructions to update database software using a custom database software image.
- To update Grid Infrastructure software using custom Grid Infrastructure software image
  Use the following instructions to update Grid Infrastructure software using a custom Grid Infrastructure software image.
- To delete a software image
  Use the following instructions to delete a software image.

To view the list of software images

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Resources**, click **Database Software Images**.

   The resulting Software images page displays the list of custom software images, which includes details such as Image type (Database, Grid infrastructure), and version.

To create a database software image

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Resources**, click **Software Images**.
3. Click **Create Software Image**.
4. In the resulting Create software image page, click **Database Software image**.
5. In the **Display name** field, provide a display name for your image. Avoid entering confidential information.
6. Choose your **Compartment**.
7. Choose a **Database release**.

8. Choose the **Database version** for your image. You can create a database software image using any supported Oracle Database release update (RU).
9. Optionally, you can enter a comma-separated list of one-off (interim) patch numbers.
10. Optionally, you can upload an Oracle Home inventory file from an existing Oracle Database. See *Using the OPatch lsinventory Command to Verify the Patches Applied to an Oracle Home* for instructions on creating an inventory file using OPatch.
11. Click **Show Advanced Options** to add tags to your database software image. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
12. Click **Create software image**.

To create a Grid Infrastructure software image

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Resources**, click **Software Images**.
3. Click **Create Software Image**.
4. In the resulting Create software image page, click **Grid Infrastructure Software image**.
5. In the **Display name** field, provide a display name for your image. Avoid entering confidential information.
6. Choose your **Compartment**.
7. Choose a **Grid Infrastructure release**.
8. Choose the **Grid Infrastructure** version for your image. You can create a Grid Infrastructure software image using any supported Oracle Grid Infrastructure release update (RU).
9. Optionally, you can enter a comma-separated list of one-off (interim) patch numbers.
10. Optionally, you can upload a Grid Infrastructure Home inventory file from an existing Oracle Grid Infrastructure. See *Using the OPatch lsinventory Command to Verify the Patches Applied to an Oracle Home* for instructions on creating an inventory file using OPatch.
11. Click **Show Advanced Options** to add tags to your database software image. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
12. Click **Create software image**.

To create a database software image from a Database Home

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure** .
2. Choose your **Compartment**.
3. Navigate to the Database Home: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

4. Click **Database Homes** under **Resources**.
5. Find the Database Home you want to use to create the database software image in the list of Database Homes. Click the name of the Database Home to display details about it.
6. Click **Create Image from Database Home**.
7. In the **Create Database Software Image** panel, enter a **Display name** and select a compartment for the software image.
8. Click **Create**.

To view the image details of a software image

Use this procedure to view the image details such as Image type, Oracle release, and version used for creating the software image and one-off (interim) patches included (if any) in a software image.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure** .
2. Under **Resources**, click **Software images**.
3. In the list of software images, find the image you want to view and click on the display name of the image.
4. On the Software Image details page for your selected image, details about the image are displayed:
   - The **General Information** section includes details such as Image type.
   - The **Patch Information** section includes details such as release, version, and available interim patches.
   - The **One-Off Patches** field displays the number of one-off patches included (if any) in a software image. The count includes all patches specified when creating the image (including patches listed in `lsinventory`).
     - To view the included patches (if any are included),
       * click the **Show** link to view the list in the One-Off Patches overlay window.
       * click the **Copy** link and paste the list of included patches into a text editor. The copied list of patch numbers is comma-separated and can be used to create additional database software images.

To move a software image to a different compartment

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure** .
2. Under **Resources**, click **Software images**.
3. In the list of software images, find the image you want to move and click the Actions icon (three dots) at the end of the row.
4. Click **Move resource**.
5. On the resulting Move resource to a different compartment dialog, choose a target compartment.
6. Click **Move resource**.

To update database software using custom database software image

Use the following instructions to update database software using a custom database software image.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.
3. Click the name of the VM cluster that you want to update the database software image.
4. On the resulting VM cluster details page, click the **View updates** link in the **Version** section.
5. On the resulting Updates page, click the **Custom Database Software Images** tab under the **Database Home** section.
6. Choose a **Compartment**.
7. Choose a **Region**.
   Region filter defaults to the currently connected region and lists all the software images created in that region. When you choose a different region, the software image list is refreshed to display the software images created in the selected region.
8. Click the Actions button (three dots) for the update you're interested in, and select **Run Precheck**.
9. On the resulting Confirm dialog, click **OK** to continue.
10. After running the precheck successfully, select **Apply** from the Actions button (three dots).
11. On the resulting Confirm dialog, click **OK** to continue.

To update Grid Infrastructure software using custom Grid Infrastructure software image

Use the following instructions to update Grid Infrastructure software using a custom Grid Infrastructure software image.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.
3. Click the name of the VM cluster that you want to update the Grid Infrastructure software image.
4. On the resulting VM cluster details page, click the **View updates** link in the **Version** section.
5. On the resulting Updates page, click the **Custom Grid Infrastructure Software Images** tab under the **Grid Infrastructure Updates** section.
6. Choose a **Compartment**.
7. Choose a **Region**.
   Region filter defaults to the currently connected region and lists all the software images created in that region. When you choose a different region, the software image list is refreshed to display the software images created in the selected region.

8. Click the Actions button (three dots) for the update you're interested in, and select **Run Precheck**.
9. On the resulting Confirm dialog, click **OK** to continue.
10. After running the precheck successfully, select **Apply Grid Infrastructure Patch** from the Actions button (three dots).
11. On the resulting Confirm dialog, click **OK** to continue.

To delete a software image

Use the following instructions to delete a software image.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Resources**, click **Software Images**.
3. In the list of software images, find the image you want to delete and click the Actions icon (three dots) at the end of the row.
4. Click **Delete**.
5. In the resulting Delete software image dialog, enter the name of the software image to confirm your action.
6. Click **Delete software image**.

Using the API to manage database software images

Use these API operations to manage database software images:

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.

- CreateDatabaseSoftwareImage
- ListDatabaseSoftwareImages
- GetDatabaseSoftwareImage
- DeleteDatabaseSoftwareImage
- ChangeDatabaseSoftwareImageCompartment

Create Oracle Database Homes on an Exadata Cloud Infrastructure System

Learn to create Oracle Database Homes on Exadata Cloud Infrastructure.

- About Creating Oracle Database Homes on an Exadata Cloud Infrastructure System
  You can add Oracle Database homes (referred to as **Database Homes** in Oracle Cloud Infrastructure) to an existing VM cluster by using the Oracle Cloud Infrastructure Console, the API, or the CLI.
- To create a new Database Home in an existing Exadata Cloud Infrastructure instance
  To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.
- To create a database software image from a Database Home

- Using the API to Create Oracle Database Home on Exadata Cloud Infrastructure
  To create an Oracle Database home, review the list of API calls.

About Creating Oracle Database Homes on an Exadata Cloud Infrastructure System

You can add Oracle Database homes (referred to as **Database Homes** in Oracle Cloud Infrastructure) to an existing VM cluster by using the Oracle Cloud Infrastructure Console, the API, or the CLI.

A Database Home is a directory location on the Exadata database virtual machines that contains Oracle Database software binary files.

> **Note:**
>
> Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

You can also add and remove Database homes, and perform other management tasks on a Database home by using the `dbaascli` utility.

Related Topics

- Using the dbaascli Utility on Exadata Cloud Infrastructure
  Learn to use the `dbaascli` utility on Exadata Cloud Infrastructure.

To create a new Database Home in an existing Exadata Cloud Infrastructure instance

To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.

> **Note:**
>
> Minimum requirements for provisioning a Database 23ai home:
>
> - Grid Infrastructure 23ai
> - Exadata Guest VM running Exadata System Software 23.1.8

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system you want to create the new Database Home on:
   - *Cloud VM clusters (*new resource model*):* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters,

find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

- *DB systems:* Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

4. Under **Resources**, click **Database Homes**.

   A list of Database Homes is displayed.

5. Click **Create Database Home**.

6. In the **Create Database Home** dialog, enter the following:

   - **Database Home display name:** The display name for the Database Home. Avoid entering confidential information.
     Database image: Determines what Oracle Database version is used for the database. You can have databases with different minor versions the same database home. The major versions must remain the same. By default, the latest Oracle-published database software image is selected.

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

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

     – **Custom Database Software Images:** These images are created by your organization and contain customized configurations of software updates and patches. Use the **Select a compartment**, **Select a region**, and **Select a Database version** selectors to limit the list of custom database software images to a specific compartment, region, or Oracle Database software major release version.

       Region filter defaults to the currently connected region and lists all the software images created in that region. When you choose a different region, the software image list is refreshed to display the software images created in the selected region.

> **Note:**
>
> For the Oracle Database major version releases available in Oracle Cloud Infrastructure, images are provided for the current version plus the three most recent older versions (N through N - 3). For example, if an instance is using Oracle Database 19c, and the latest version of 19c offered is 19.8.0.0.0, images available for provisioning are for versions 19.8.0.0.0, 19.7.0.0, 19.6.0.0 and 19.5.0.0.

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

After choosing a software image, click **Select** to return to the Create Database dialog.

- **Unified Auditing:** Select this check box to enable Unified Auditing framework.

  > **Note:**
  >
  > You cannot disable Unified Auditing after provisioning the Database Home.

  – **Oracle Database version 12.1 or lower:** You cannot use the Unified Auditing framework. Instead, use the Traditional Audit - legacy Oracle Database audit framework.

  – **Oracle Database version 12.2 or higher:** You can enable Unified Auditing from the OCI Console. For Oracle Database versions 12.2 or higher but lower than version 23, the **Unified Auditing** checkbox is not selected by default. However, it is selected by default for Oracle Database version 23.

  Unified Auditing field in the **General Information** section on the **Database Home Details** page displays if Unified Auditing is **Enabled** or **Disabled**.

- Click **Show Advanced Options** to specify advanced options for the Database Home.

  – **Tags:** If you have permissions to create a resource, then you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see Resource Tags. If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.

7. Click **Create**.

When the Database Home creation is complete, the status changes from Provisioning to Available.

To create a database software image from a Database Home

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure** .
2. Choose your **Compartment**.
3. Navigate to the Database Home: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. Click **Database Homes** under **Resources**.
5. Find the Database Home you want to use to create the database software image in the list of Database Homes. Click the name of the Database Home to display details about it.

6. Click **Create Image from Database Home**.
7. In the **Create Database Software Image** panel, enter a **Display name** and select a compartment for the software image.
8. Click **Create**.

Using the API to Create Oracle Database Home on Exadata Cloud Infrastructure

To create an Oracle Database home, review the list of API calls.

For information about using the API and signing requests, see "REST APIs" and "Security Credentials". For information about SDKs, see "Software Development Kits and Command Line Interface".

To create Database Homes in Exadata Cloud Infrastructure, use the API operation `CreateDbHome`.

For the complete list of APIs, see "Database Service API".

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- CreateDbHome
- Database Service API

Managing Oracle Database Homes on an Exadata Cloud Infrastructure Instance

You can delete or view information about Oracle Database Homes (referred to as "Database Homes" in Oracle Cloud Infrastructure) by using the Oracle Cloud Infrastructure Console, the API, or the CLI.

For information on how to perform these tasks manually, see About Using the dbaascli Utility on Exadata Cloud Infrastructure .

- Manage Database Home Using the Console
  Use the OCI console to manage the various operations needed on a Database Home.
- Using the API to Manage Oracle Database Home on Exadata Cloud Infrastructure
  Review the list of API calls to manage Oracle Database home.

Manage Database Home Using the Console

Use the OCI console to manage the various operations needed on a Database Home.

- To view information about a Database Home
- To delete a database home
  You cannot delete a Database Home that contains databases. You must first terminate the databases to empty the Database Home. See To terminate a database to learn how to terminate a database.

- To manage tags for your Database Home
- Using the Console to Move a Database to Another Database Home
  Learn to move a database to another Database Home.

To view information about a Database Home

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system containing the Database Home:.
   - *Cloud VM clusters (new resource model)*: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster
   - *DB systems*: Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. On the DB System Details page, under Resources, click **Database Homes**.
5. In the list of Database Homes, find the Database Home you are interested in, and then click its name to display details about it.

(Optional) Enter the result of the procedure here.

To delete a database home

You cannot delete a Database Home that contains databases. You must first terminate the databases to empty the Database Home. See To terminate a database to learn how to terminate a database.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system containing the Database Home you want to delete:
   - *Cloud VM clusters (new resource model):* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
   - *DB systems:* Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. On the DB System Details page, under Resources, click **Database Homes**.
5. In the list of Database Homes, find the Database Home you want to delete, and then click its name to display details about it.
6. On the Database Home Details page, click **Delete**.

   If the Database Home contains databases, you will not be able to proceed. You must cancel the deletion, empty the Database Home as applicable, and then retry the deletion.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

To manage tags for your Database Home

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system containing the Database Home:
   - *Cloud VM clusters (*new resource model*):* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
   - *DB systems:* Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. Under **Resources**, click **Database Homes**.
5. In the list of Database Homes, find the Database Home you want to administer.
6. Click the the Actions icon (three dots) on the row listing the Database Home, and then click **Add Tags**.

Using the Console to Move a Database to Another Database Home

Learn to move a database to another Database Home.

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment** that contains the VM cluster that hosts the database that you want to move.
3. Click **Exadata VM Clusters** in the left hand navigation.
4. Click the name of the VM cluster that contains the database that you want to move.
5. In the Resources list of the VM Cluster Details page, click **Databases**.
6. Click the name of the database that you want to move.

   The Database Details page displays information about the selected database.
7. Click **More Actions** and **Move To Another Home**.
8. In the resulting dialog, select the target Database Home.

> **Note:**
>
> Oracle recommends using Database Homes, which are running the latest (N) to 3 versions from the latest (N-3) RU versions when updating the software version of the database by moving them to a target DB Home. Only DB Homes provisioned with database versions, which meet this best practice criterion are available as target homes to move your database.

9. Click **Move Database**.

The database will be stopped in the current home and then restarted in the destination home. While the database is being moved, the Database Home status displays as **Moving Database**. When the operation completes, Database Home is updated with the current home. If the operation is unsuccessful, the status of the database displays as **Failed**, and the Database Home field provides information about the reason for the failure.

Using the API to Manage Oracle Database Home on Exadata Cloud Infrastructure

Review the list of API calls to manage Oracle Database home.

For information about using the API and signing requests, see "REST APIs" and "Security Credentials". For information about SDKs, see "Software Development Kits and Command Line Interface".

Use these API operations to manage Database Homes:

- `ListDbHomes`
- `GetDbHome`
- `DeleteDbHome`

For the complete list of APIs, see "Database Service API".

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- ListDbHomes
- GetDbHome
- DeleteDbHome
- Database Service API

Manage Databases on Exadata Cloud Infrastructure

- Prerequisites and Limitations for Creating and Managing Oracle Databases on Oracle Exadata Database Service on Dedicated Infrastructure
  Review the prerequisites for creating and managing Oracle Databases on Oracle Exadata Database Service on Dedicated Infrastructure.

- Oracle Database Releases Supported by Oracle Exadata Database Service on Dedicated Infrastructure
  Learn about the versions of Oracle Database that Oracle Exadata Database Service on Dedicated Infrastructure supports.
- Provisioning and Managing Exadata Databases
  This topic describes creating and managing Oracle Databases on an Exadata Cloud Infrastructure instance instance.
- Using the API to manage Databases
- Create and Manage Exadata Pluggable Databases
  You can create and manage pluggable databases (PDBs) in Exadata Cloud Infrastructure using the Console and APIs.
- Restoring an Exadata Pluggable Database
  You can perfrom in-place and out of place restore of an Exadata pluggable database.
- Changing the Database Passwords
  To change the SYS password, or to change the TDE wallet password, use this procedure.

Prerequisites and Limitations for Creating and Managing Oracle Databases on Oracle Exadata Database Service on Dedicated Infrastructure

Review the prerequisites for creating and managing Oracle Databases on Oracle Exadata Database Service on Dedicated Infrastructure.

Before you can create and use an Oracle Database on Exadata Cloud Infrastructure, you must:

- Provision Exadata Cloud Infrastructure infrastructure
- Configure a VM cluster
- Create any required backup destinations

You can create one or more databases on each Oracle Exadata Database Service on Dedicated Infrastructure system. Other than the storage and processing limits of your Oracle Exadata system, there is no maximum for the number of databases that you can create. By default, databases on Exadata Cloud Infrastructure use Oracle Database Enterprise Edition - Extreme Performance. This edition provides all the features of Oracle Database Enterprise Edition, plus all of the database enterprise management packs, and all of the Enterprise Edition options, such as Oracle Database In-Memory, and Oracle Real Application Clusters (Oracle RAC). If you use your own Oracle Database licenses, then your ability to use various features is limited by your license holdings. TDE Encryption is required for all cloud databases. All new tablespaces will automatically be enabled for encryption.

Oracle Database Releases Supported by Oracle Exadata Database Service on Dedicated Infrastructure

Learn about the versions of Oracle Database that Oracle Exadata Database Service on Dedicated Infrastructure supports.

Exadata Cloud Infrastructure supports the following Oracle Database software releases:

- Oracle Database 19c (19.0)
- Oracle Database 12c Release 2 (12.2) (requires a valid Upgrade Support contract).
- Oracle Database 12c Release 1 (12.1) (requires a valid Upgrade Support contract).

- Oracle Database 11g Release 2 (11.2) (requires a valid Upgrade Support contract).

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

Related Topics

- https://support.oracle.com/epmos/faces/DocContentDisplay?id=742060.1

Provisioning and Managing Exadata Databases

This topic describes creating and managing Oracle Databases on an Exadata Cloud Infrastructure instance instance.

In this documentation, "database" refers to a container database (CDB). When you provision a database in an Exadata cloud VM cluster, the database includes an initial pluggable database (PDB). For more information on these resource types, see Multitenant Architecture in the Oracle Database documentation. See Exadata Pluggable Database Operations for more information on pluggable databases in Exadata Cloud Infrastructure.

You can create Database Homes, databases, and pluggable databases at any time by using the Console or the Database APIs.

When you add a database to a VM cluster on an Exadata instance, the database versions you can select from depend on the current patch level of that resource. You may have to patch your VM cluster to add later database versions.

After you provision a database, you can move it to another Database Home. Consolidating databases under the same home can facilitate management of these resources. All databases in a given Database Home share the Oracle Database binaries and therefore, have the same database version. The Oracle-recommended way to patch a database to a version that is different from the current version is to move the database to a home running the target version. For information about patching, see Patching an Exadata Cloud Service Instance.

> **Note:**
>
> When provisioning databases, make sure your VM cluster has enough OCPUs enabled to support the total number of database instances on the system. Oracle recommends the following general rule: for each database, enable 1 OCPU per node. See To scale CPU cores in an Exadata Cloud Service cloud VM cluster or DB system for information on scaling your OCPU count up or down.

When you create an Exadata database, you can choose to encrypt the database using your own encryption keys that you manage. You can rotate encryption keys, periodically, to maintain security compliance and, in cases of personnel changes, to disable access to a database.

> **Note:**
>
> - The encryption key you use must be AES-256.
> - To ensure that your Exadata database uses the most current versions of the Vault encryption key, rotate the key from the Database Details page on the Oracle Cloud Infrastructure Console. Do not use the Vault service's Console pages to rotate your Database keys.

If you want to use your own encryption keys to encrypt a database that you create, then you must create a dynamic group and assign specific policies to the group for customer-managed encryption keys. See Managing Dynamic Groups and Let security admins manage vaults, keys, and secrets. Additionally, see To integrate customer-managed key management into Exadata Cloud Service if you need to update customer-managed encryption libraries for the Vault service.

You can also add and remove databases, and perform other management tasks on a database by using command line utilities. For information and instructions on how to use these utilities, see Creating and Managing Exadata Databases Manually.

- Database Memory Initialization Parameters
- Customer-Managed Keys in Exadata Cloud Infrastructure
  Customer-managed keys for Exadata Cloud Infrastructure is a feature of Oracle Cloud Infrastructure (OCI) Vault service that enables you to encrypt your data using encryption keys that you control.
- Using the Console to Manage Databases on Oracle Exadata Database Service on Dedicated Infrastructure
  To create or terminate a database, complete procedures using the Oracle Exadata console.
- Known Issues in Exadata Cloud Infrastructure
  `rac stopdb` failed

Database Memory Initialization Parameters

- When creating a container database, the initialization parameter, `SGA_TARGET` is set by the automation. This will automatically size the SGA memory pools. The setting will vary depending on the size of the database VM total memory. If the VM has less than or equal to 60 GB of system memory, `SGA_TARGET` is set to 3800 MB. If the VM has 60 GB or more system memory, `SGA_TARGET` is set to 7600 MB.
- The database initialization parameter `USE_LARGE_PAGES` is set to ONLY upon database creation, which will require the use of large pages for SGA memory. If the VM is configured with insufficient large pages, the instance will fail to start.

Customer-Managed Keys in Exadata Cloud Infrastructure

Customer-managed keys for Exadata Cloud Infrastructure is a feature of Oracle Cloud Infrastructure (OCI) Vault service that enables you to encrypt your data using encryption keys that you control.

The OCI Vault service provides you with centralized key management capabilities that are highly available and durable. This key-management solution also offers secure key storage using isolated partitions (and a lower-cost shared partition option) in FIPS 140-2 Level 3-certified hardware security modules, and integration with select Oracle Cloud Infrastructure services. Use customer-managed keys when you need security governance, regulatory compliance, and homogenous encryption of data, while centrally managing, storing, and monitoring the life cycle of the keys you use to protect your data.

You can:

- Enable customer-managed keys when you create databases in Exadata Cloud Infrastructure
- Switch from Oracle-managed keys to customer-managed keys

- Rotate your keys to maintain security compliance

Requirements

To enable management of customer-managed encryption keys, you must create a policy in the tenancy that allows a particular dynamic group to do so, similar to the following: `allow dynamic-group dynamic_group_name to manage keys in tenancy`.

Another policy is needed if the Vault being used by the customer is replicated (https://docs.oracle.com/en-us/iaas/Content/KeyManagement/Tasks/replicatingvaults.htm). For vaults that are replicated, this policy is needed: `allow dynamic-group dynamic_group_name to read vaults in tenancy`

Limitations

To enable Data Guard on Exadata Cloud Infrastructure databases that use customer-managed keys, the primary and standby databases must be in the same realm.

Task 1. Create a Vault and a Master Encryption Key

Create a vault in the Vault service by following the instructions in To create a new vault in Oracle Cloud Infrastructure Documentation. When following these instructions, Oracle recommends that you create the vault in a compartment created specifically to contain the vaults containing customer-managed keys, as described in Before You Begin: Compartment Hierarchy Best Practice.

After creating the vault, create at least one master encryption key in the vault by following the instructions in To create a new master encryption key in Oracle Cloud Infrastructure Documentation. When following these instructions, make these choices:

- **Create in Compartment**: Oracle recommends that you create the master encryption key in the same compartment as its vault; that is, the compartment created specifically to contain the vaults containing customer-managed keys.
- **Protection Mode**: Choose an appropriate value from the drop-down list:
  - **HSM** to create a master encryption key that is stored and processed on a hardware security module (HSM).
  - **Software** to create a master encryption key that is stored in a software file system in the Vault service. Software-protected keys are protected at rest using an HSM-based root key. You may export software keys to other key management devices or to a different OCI cloud region. Unlike HSM keys, software-protected keys are free of cost.
- **Key Shape Algorithm**: AES
- **Key Shape Length**: 256 bits

Oracle strongly recommends that you create a separate master encryption key for each of your container databases (CDBs). Doing so makes management of key rotation over time much simpler.

Task 2. Create a Service Gateway, a Route Rule, and an Egress Security Rule

Create a service gateway in the VCN (Virtual Cloud Network) where your Oracle Exadata Database Service on Dedicated Infrastructure resources reside by following the instructions in Task 1: Create the service gateway in Oracle Cloud Infrastructure Documentation.

After creating the service gateway, add a route rule and an egress security rule ***to each subnet*** (in the VCN) where Oracle Exadata Database Service on Dedicated Infrastructure resources reside so that these resources can use the gateway to access the Vault service:

1. Go to the **Subnet Details** page for the subnet.

2. In the **Subnet Information** tab, click the name of the subnet's **Route Table** to display its **Route Table Details** page.
3. In the table of existing **Route Rules**, check whether there is already a rule with the following characteristics:
   - **Destination**: All IAD Services In Oracle Services Network
   - **Target Type**: Service Gateway
   - **Target**: The name of the service gateway you just created in the VCN

   If such a rule does not exist, click **Add Route Rules** and add a route rule with these characteristics.
4. Return to the **Subnet Details** page for the subnet.
5. In the subnet's **Security Lists** table, click the name of the subnet's security list to display its **Security List Details** page.
6. In the side menu, under **Resources**, click **Egress Rules**.
7. In the table of existing **Egress Rules**, check whether there is already a rule with the following characteristics:
   - **Stateless**: No
   - **Destination**: All IAD Services In Oracle Services Network
   - **IP Protocol**: TCP
   - **Source Port Range**: All
   - **Destination Port Range**: 443

   If such a rule does not exist, click **Add Egress Rules** and add an egress rule with these characteristics.

Task 3. Create a Dynamic Group and a Policy Statement

To grant your Oracle Exadata Database Service on Dedicated Infrastructure resources permission to access customer-managed keys, you create an IAM dynamic group that identifies these resources and then create an IAM policy that grants this dynamic group access to the master encryption keys you created in the Vault service.

When defining the dynamic group, you identify your Oracle Exadata Database Service on Dedicated Infrastructure resources by specifying the OCID of the compartment containing your Exadata Infrastructure resource.

1. Copy the OCID of the compartment containing your Exadata Infrastructure resource. You can find this OCID on the **Compartment Details** page of the compartment.
2. Create a dynamic group by following the instructions in To create a dynamic group in Oracle Cloud Infrastructure Documentation. When following these instructions, enter a matching rule of this format:

   ```
   ALL {resource.compartment.id ='<compartment-ocid>'}
   ```

   where *<compartment-ocid>* is the OCID of the compartment containing your Exadata Infrastructure resource.

After creating the dynamic group, navigate to (or create) an IAM policy in a compartment higher up in your compartment hierarchy than the compartment containing your vaults and keys. Then, add a policy statement of this format:

```
allow dynamic-group <dynamic-group-name>
to manage keys
in compartment <vaults-and-keys-compartment>
where all {
target.key.id='<key_ocid>',
request.permission!='KEY_DELETE',
request.permission!='KEY_MOVE',
request.permission!='KEY_IMPORT',
request.permission!='KEY_BACKUP'
}
```

If you are using a replicated virtual private vault for the Oracle Data Guard deployment, add an additional policy statement in this format:

```
allow dynamic-group <dynamic-group>
to read vaults
in tenancy | compartment <vaults-and-keys-compartment>
```

where *<dynamic-group>* is the name of the dynamic group you created and *<vaults-and-keys-compartment>* is the name of the compartment in which you created your vaults and master encryption keys.

- To integrate customer-managed key management into Exadata Cloud Infrastructure
  If you choose to encrypt databases in an Exadata Cloud Infrastructure instance using encryption keys that you manage, then you may update the following two packages (using Red Hat Package Manager) to enable DBAASTOOLS to interact with the APIs that customer-managed key management uses.

Related Topics

- To create a database in an existing Exadata Cloud Infrastructure instance
  This topic covers creating your first or subsequent databases.
- To administer Vault encryption keys
  Use this procedure to rotate the Vault encryption key or or change the encryption management configuration.
- Known Issues for Exadata Cloud Infrastructure and Data Guard
  Possible TDE key replication issue, and MRP and DG LCM operation failures.
- To integrate customer-managed key management into Exadata Cloud Infrastructure
  If you choose to encrypt databases in an Exadata Cloud Infrastructure instance using encryption keys that you manage, then you may update the following two packages (using Red Hat Package Manager) to enable DBAASTOOLS to interact with the APIs that customer-managed key management uses.

To integrate customer-managed key management into Exadata Cloud Infrastructure

If you choose to encrypt databases in an Exadata Cloud Infrastructure instance using encryption keys that you manage, then you may update the following two packages (using Red

Hat Package Manager) to enable DBAASTOOLS to interact with the APIs that customer-managed key management uses.

KMS TDE CLI

To update the KMS TDE CLI package, you must complete the following task on all nodes in the Exadata Cloud Infrastructure instance:

1. Deinstall current KMS TDE CLI package, as follows:

```
rpm -ev kmstdecli
```

2. Install the updated KMS TDE CLI package, as follows:

```
rpm -ivh kms_tde_cli
```

LIBKMS

LIBKMS is a library package necessary to synchronize a database with customer-managed key management through PKCS11. When a new version of LIBKMS is installed, any databases converted to customer-managed key management continue to use the previous LIBKMS version, until the database is stopped and restarted.

To update the LIBKMS package, you must complete the following task on all nodes in the Exadata Cloud Infrastructure instance:

1. Confirm that the LIBKMS package is already installed, as follows:

```
rpm -qa --last | grep libkmstdepkcs11
```

2. Install a new version of LIBKMS, as follows:

```
rpm -ivh libkms
```

3. Use SQL*Plus to stop and restart all databases converted to customer-managed key management, as follows:

```
shutdown immediate;
startup;
```

4. Ensure that all converted databases are using the new LIBKMS version, as follows:

```
for pid in $(ps aux | grep "<dbname>" | awk '{print $2;}'); do echo $pid;
sudo lsof -p $pid | grep kms | grep "pkcs11_[0-9A-Za-z.]*" | sort -u; done
| grep pkcs11
```

5. Deinstall LIBKMS packages that are no longer being used by any database, as follows:

```
rpm -ev libkms
```

Using the Console to Manage Databases on Oracle Exadata Database Service on Dedicated Infrastructure

To create or terminate a database, complete procedures using the Oracle Exadata console.

- To create a database in an existing Exadata Cloud Infrastructure instance
  This topic covers creating your first or subsequent databases.
- To manage SYS user and TDE Wallet passwords
  Learn to manage administrator (SYS user) and TDE wallet passwords.
- To view details of a Protected Database
  To view the details of a Protected Database, use this procedure.
- To create a database from a backup
- To create a database from the latest backup
- To move a database to another Database Home
  This task explains how to patch a single Oracle Database in your Exadata Cloud Infrastructure instance by moving it to another Database Home.
- To terminate a database
- To administer Vault encryption keys
  Use this procedure to rotate the Vault encryption key or or change the encryption management configuration.

To create a database in an existing Exadata Cloud Infrastructure instance

This topic covers creating your first or subsequent databases.

> **Note:**
>
> If IORM is enabled on the Exadata Cloud Infrastructure instance, then the default directive will apply to the new database and system performance might be impacted. Oracle recommends that you review the IORM settings and make applicable adjustments to the configuration after the new database is provisioned.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system you want to create the database in: **Cloud VM clusters (The New Exadata Cloud Infrastructure Resource Model)**: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems: Under **Oracle Base Database**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. Click **Create Database**.
5. In the **Create Database** dialog, enter the following:

   > **Note:**
   >
   > You cannot modify the `db_name`, `db_unique_name`, and SID prefix after creating the database.

- **Database name:** The name for the database. The database name must meet the requirements:
  - Maximum of 8 characters
  - Contain only alphanumeric characters
  - Begin with an alphabetic character
  - Cannot be part of the first 8 characters of a `DB_UNIQUE_NAME` on the VM cluster
  - DO NOT use the following reserved names: `grid`, `ASM`
- **Database unique name suffix:**
  Optionally, specify a value for the `DB_UNIQUE_NAME` database parameter. The value is case insensitive.

  The unique name must meet the requirements:
  - Maximum of 30 characters
  - Contain only alphanumeric or underscore (_) characters
  - Begin with an alphabetic character
  - Unique across the VM cluster. Recommended to be unique across the tenancy.

  If not specified, the system automatically generates a unique name value, as follows:

  <db_name>_<3_chars_unique_string>_<region-name>
- **Database version:** The version of the database. You can mix database versions on the Exadata DB system.
- **Database Home:** The Oracle Database Home for the database. Choose the applicable option:
  - **Select an existing Database Home:** The Database Home display name field allows you to choose the Database Home from the existing homes for the database version you specified. If no Database Home with that version exists, you must create a new one.
  - **Create a new Database Home**: Use this option to provision a new Database Home for your Data Guard peer database.
    Click **Change Database Image** to use a desired Oracle-published image or a custom *database software image* that you have created in advance, then select an **Image Type**:
    * **Oracle Provided Database Software Images:**
      then you can use the **Display all available version** switch to choose from all available PSUs and RUs. The most recent release for each major version is indicated with a **latest** label.

      > **Note:**
      >
      > For the Oracle Database major version releases available in Oracle Cloud Infrastructure, images are provided for the current version plus the three most recent older versions (N through N - 3). For example, if an instance is using Oracle Database 19c, and the latest version of 19c offered is 19.8.0.0.0, images available for provisioning are for versions 19.8.0.0.0, 19.7.0.0, 19.6.0.0 and 19.5.0.0.

        - **Custom Database Software Images:** These images are *created by your organization* and contain customized configurations of software updates and patches. Use the **Select a compartment**, **Select a region**, and **Select a Database version** selectors to limit the list of custom database software images to a specific compartment, region, or Oracle Database software major release version.
          Region filter defaults to the currently connected region and lists all the software images created in that region. When you choose a different region, the software image list is refreshed to display the software images created in the selected region.

- **PDB name:** *(Optional)* For Oracle Database 12*c* (12.1.0.2) and later, you can specify the name of the pluggable database. The PDB name must begin with an alphabetic character, and can contain a maximum of eight alphanumeric characters. The only special character permitted is the underscore (_).
  To avoid potential service name collisions when using Oracle Net Services to connect to the PDB, ensure that the PDB name is unique across the entire VM cluster. If you do not provide the name of the first PDB, then a system-generated name is used.
- **Create administrator credentials:** *(Read only)* A database administrator `SYS` user will be created with the password you supply.
  - **Username:** SYS
  - **Password:** Supply the password for this user. The password must meet the following criteria:
    A strong password for SYS, SYSTEM, TDE wallet, and PDB Admin. The password must be 9 to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters. The special characters must be _, #, or -. The password must not contain the username (SYS, SYSTEM, and so on) or the word "**oracle**" either in forward or reversed order and regardless of casing.
  - **Confirm password:** Re-enter the SYS password you specified.
  - Using a **TDE wallet password** is optional. If you are using customer-managed encryption keys stored in a vault in your tenancy, the TDE wallet password is not applicable to your DB system. Use **Show Advanced Options** at the end of the Create Database dialog to configure customer-managed keys.
    If you are using customer-managed keys, or if you want to specify a different TDE wallet password, uncheck the **Use the administrator password for the TDE wallet box**. If you are using customer-managed keys, leave the TDE password fields blank. To set the TDE wallet password manually, enter a password in the **Enter TDE wallet password** field, and then confirm by entering it into the **Confirm TDE wallet password** field.
- **Configure database backups:** Specify the settings for backing up the database to Autonomous Recovery Service or Object Storage:
  - **Enable automatic backup**: Check the check box to enable automatic incremental backups for this database. If you are creating a database in a security zone compartment, you must enable automatic backups.
  - **Backup Destination**: Your choices are **Autonomous Recovery Service** or **Object Storage**.
  - **Backup Scheduling**:
    - **Object Storage (L0)**:
      - **Full backup scheduling day**: Choose a day of the week for the initial and future L0 backups to start.

        - **Full backup scheduling time (UTC)**: Specify the time window when the full backups start when the automatic backup capability is selected.
        - **Take the first backup immediately**: A full backup is an operating system backup of all datafiles and the control file that constitute an Oracle Database. A full backup should also include the parameter file(s) associated with the database. You can take a full database backup when the database is shut down or while the database is open. You should not normally take a full backup after an instance failure or other unusual circumstances.

          If you choose to defer the first full backup your database may not be recoverable in the event of a database failure.
      - **Object Storage (L1)**:
        - **Incremental backup scheduling time (UTC)**: Specify the time window when the incremental backups start when the automatic backup capability is selected.
      - **Autonomous Recovery Service (L0)**:
        - **Scheduled day for initial backup**: Choose a day of the week for the initial backup.
        - **Scheduled time for initial backup (UTC)**: Select the time window for the initial backup.
        - **Take the first backup immediately**: A full backup is an operating system backup of all datafiles and the control file that constitute an Oracle Database. A full backup should also include the parameter file(s) associated with the database. You can take a full database backup when the database is shut down or while the database is open. You should not normally take a full backup after an instance failure or other unusual circumstances.

          If you choose to defer the first full backup your database may not be recoverable in the event of a database failure.
      - **Autonomous Recovery Service (L1)**:
        - **Scheduled time for daily backup (UTC)**: Specify the time window when the incremental backups start when the automatic backup capability is selected.
    - **Deletion options after database termination**: Options that you can use to retain protected database backups after the database is terminated. These options can also help restore the database from backups in case of accidental or malicious damage to the database.
      - **Retain backups for the period specified in your protection policy or backup retention period**: Select this option if you want to retain database backups for the entire period defined in the Object Storage Backup retention period or Autonomous Recovery Service protection policy after the database is terminated.
      - **Retain backups for 72 hours, then delete**: Select this option to retain backups for a period of 72 hours after you terminate the database.
    - **Backup Retention Period/Protection Policy**: If you choose to enable automatic backups, you can choose a policy with one of the following preset retention periods, or a Custom policy.

Object Storage Backup retention period: 7, 15, 30, 45, 60. Default: 30 days. The system automatically deletes your incremental backups at the end of your chosen retention period.

Autonomous Recovery Service protection policy:

* **Bronze:** 14 days
* **Silver:** 35 days
* **Gold:** 65 days
* **Platinum:** 95 days
* Custom defined by you
* **Default:** Silver - 35 days

– **Enable Real-Time Data Protection**: Real-time protection is the continuous transfer of redo changes from a protected database to **Autonomous Recovery Service**. This reduces data loss and provides a recovery point objective (RPO) near 0. This is an extra cost option.

6. Click **Show Advanced Options** to specify advanced options for the database:

- **Management:**

  Oracle SID prefix: The Oracle Database instance number is automatically added to the SID prefix to create the `INSTANCE_NAME` database parameter. The `INSTANCE_NAME` parameter is also known as the `SID`. The `SID` is unique across the cloud VM Cluster. If not specified, `SID` prefix defaults to the `db_name`.

  > **Note:**
  >
  > Entering an `SID` prefix is only available for Oracle 12.1 databases and above.

  The `SID` prefix must meet the requirements:

  – Maximum of 12 characters
  – Contain only alphanumeric characters. You can, however, use underscore (_), which is the only special character that is not restricted by this naming convention.
  – Begin with an alphabetic character
  – Unique in the VM cluster
  – DO NOT use the following reserved names: `grid`, `ASM`

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

  If you are creating a database in an Exadata Cloud Service VM Cluster, then you can choose to use encryption based on encryption keys that you manage. By default, the database is configured using Oracle-managed encryption keys. To configure the database with encryption based on encryption keys you manage:

  a. Select **Use customer-managed keys**. You must have a valid encryption key in Oracle Cloud Infrastructure Vault service. See Let security admins manage vaults, keys, and secrets.

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

b. Choose a **Vault**.

c. Select a **Master encryption key**.

d. To specify a key version other than the latest version of the selected key, check **Choose the key version** and enter the OCID of the key you want to use in the **Key version OCID** field.

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

- **Tags**: If you have permissions to create a resource, then you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see *Resource Tags* . If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.

7. Click **Create Database**.

After database creation is complete, the status changes from **Provisioning** to **Available**, and on the database details page for the new database, the **Encryption** section displays the encryption key name and the encryption key OCID.

> **WARNING:**
>
> Do not delete the encryption key from the vault. This causes any database protected by the key to become unavailable.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.
- security zone compartment
- Resource Tags
- Let security admins manage vaults, keys, and secrets

To manage SYS user and TDE Wallet passwords

Learn to manage administrator (SYS user) and TDE wallet passwords.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment** that contains the VM cluster that hosts the database that you want to change passwords.

3. Click the name of the VM cluster that contains the database that you want to change passwords.
4. In the **Resources** list of the VM Cluster Details page, click **Databases**.
5. Click the name of the database that you want to change passwords.
   The Database Details page displays information about the selected database.
6. On the Database Details page, click More actions, and then click **Manage passwords**.
7. In the resulting **Manage passwords** dialog, click **Update Administrator Password** or **Update TDE Wallet Password**.
   Depending on the option you select, the system displays the fields to edit.
   - **Update Administrator Password**: Enter the new password in both the New administrator password and Confirm administrator password fields.

     > **Note:**
     >
     > The **Update Administrator Password** option will change the sys user password only. Passwords for other administrator accounts such as system, pdbadmin, and TDE wallet will not be changed.

   - **Update TDE Wallet Password**: Enter the current wallet password in the **Enter existing TDE wallet password** field, and then enter the new password in both the **New TDE wallet password** and **Confirm TDE wallet password** fields.
8. Click **Apply** to update your chosen password.

To view details of a Protected Database

To view the details of a Protected Database, use this procedure.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the database:
   Cloud VM clusters (The New Exadata Cloud Infrastructure Resource Model): Under **Exadata at Oracle Cloud**, click **Exadata VM Clusters**.

   In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems: Under Oracle Base Database, click **DB Systems**.

   In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

   On the cloud **VM cluster** or **DB system** details page, in the Databases table, click the name of the database to display the **Database Details** page.The Backup section displays the state of the automatic backups. If the Autonomous Recovery Service is the destination, a link will be available which includes additional details. You can also check if Real-time Data Protection is enabled or disabled. Click the **Autonomous Recovery Service** link to be taken to the page containing the Protected Database details.For more information about Protected Databases, see *Viewing Protected Database Details*.

Related Topics

- Viewing Protected Database Details

To create a database from a backup

Before you begin, note the following:

- When you create a database from a backup, the availability domain is the same as the availability domain that hosts the backup or a different one within the same region.
- The Oracle Database software version you specify must be the same or later version as that of the backed-up database.
- If you are creating a database from an automatic backup, then you can choose any level 0 weekly backup, or a level 1 incremental backup created after the most recent level 0 backup. For more information on automatic backups, see Using the Console
- If the backup being used to create a database is in a security zone compartment, the database cannot be created in a compartment that is not in a security zone. See the Security Zone Policies topic for a full list of policies that affect Database service resources.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to a backup.
   - *Standalone backups:* Click **Standalone Backups** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
   - *Automatic backups:* Navigate to the Database Details page of the database associated with the backup:
     - *Cloud VM clusters (*new resource model*):* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
     - *DB systems:* Under **Exadata at Oracle Cloud**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

     Click the name of the database associated with the backup that you will use to create the new database. Locate the backup in the list of backups on the Database Details page.
4. Click the Actions icon (three dots) for the backup you chose.
5. Click **Create Database**. On the **Create Database from Backup** page, configure the database as follows.
6. In the **Provide basic information for the Exadata infrastructure** section:
   - **Select an availability domain:** It could be the same as the availability domain that hosts the backup or a different one within the same region
   - **Select Exadata infrastructure:** Select an Exadata infrastructure from the chosen compartment. Click the **Change Compartment** hyperlink to choose a different compartment.
7. In the **Configure your DB system** section:
   - *Backups created in cloud VM clusters:* Choose a cloud VM cluster to run the database from the **Select a VM cluster** drop-down list.

- *Backups created in DB systems:* Choose a shape from the **Select a shape** drop-down list, then choose a DB system to run the database from the **Select a DB system** drop-down list.

8. In the **Configure Database Home** section:
   - **Select an existing Database Home**: If you choose this option, make a selection from the **Select a Database Home** drop-down list.

     > **Note:**
     >
     > You can not create a database from backup in the same Database Home where the source database exists.

   - **Create a new Database home**: If you choose this option, enter a name for the new Database Home in the **Database Home display name** field. Click **Change Database Image** to select a database software image for the new Database Home. In the **Select a Database Software Image** panel, do the following:
     a. Select the compartment containing the database software image you want to use to create the new Database Home.
     b. Select the region containing the database software image you want to use to create the new Database Home. Region filter defaults to the currently connected region and lists all the software images created in that region. When you choose a different region, the software image list is refreshed to display the software images created in the selected region.
     c. Select the Oracle Database software version that the new Database Home will use, then choose an image from the list of available images for your selected software version.

        > **Note:**
        >
        > Database restore operations for Databases of 12.2.0.1 and earlier are not allowed at this time.

     d. Click **Select**.

9. In the **Configure database** section:

   > **Note:**
   >
   > You cannot modify the `db_name`, `db_unique_name` , and SID prefix after creating the database.

   - In the **Database name** field, name the database or accept the default name. The database name must meet the requirements:
     – Maximum of 8 characters
     – Contain only alphanumeric characters
     – Begin with an alphabetic character

  - Cannot be part of first 8 characters of a different database's `db_unique_name` on the VM cluster
  - DO NOT use the following reserved names: grid, ASM

- **Database unique name:** Specify a value for the `DB_UNIQUE_NAME` database parameter. The unique name must meet the requirements:
  - Maximum of 30 characters
  - Contain only alphanumeric or underscore (_) characters
  - Begin with an alphabetic character
  - Unique across the VM cluster. Recommended to be unique across the tenancy.

  If not specified, the system automatically generates a unique name value, as follows:

  ```
  <db_name>_<3_chars_unique_string>_<region-name>
  ```

- **Administrator username**: This read-only field displays the username for the administrator, "sys".
- In the **Password** and **Confirm password** fields, enter and re-enter a password. A strong password for SYS administrator must be 9 to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters. The special characters must be _, #, or -. The password must not contain the user name (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reverse order and regardless of casing.

10. In the **Enter the source database's TDE wallet or RMAN password** field, enter a password that matches either the Transparent Data Encryption (TDE) wallet password or RMAN password for the source database.
11. Click **Show Advanced Options** to specify advanced options for the database:
    - **Management**
      Oracle SID prefix: This option is in the **Management** tab. The Oracle Database instance number is automatically added to the SID prefix to create the `INSTANCE_NAME` database parameter. If not provided, then the SID prefix defaults to the first twelve characters of the `db_name`.

      > **Note:**
      >
      > Entering an SID prefix is only available for Oracle 12.1 databases and above.

      The SID prefix must meet the requirements:
      - Maximum of 12 characters
      - Contain only alphanumeric characters
      - Begin with an alphabetic character
      - Unique in the VM cluster
      - DO NOT use the following reserved names: grid, ASM
12. Click **Create Database**.

NOT_SUPPORTED

1. Click the Exadata cloud VM cluster or DB system name that contains the specific database to display the details page.
2. From the list of databases, click the database name associated with the backup you want to use to display a list of backups on the database details page. You can also access the list of backups for a database by clicking **Backups** in the **Resources** section.

NOT_SUPPORTED

1. Click **Standalone Backups** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. In the list of standalone backups, find the backup you want to use to create the database.

- To navigate to the list of standalone backups for your current compartment

To navigate to the list of standalone backups for your current compartment

1. Click **Standalone Backups** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. In the list of standalone backups, find the backup you want to use to create the database.

To create a database from the latest backup

Before you begin, note the following:

- When you create a database from a backup, the availability domain is the same as the availability domain that hosts the backup or a different one within the same region.
- The Oracle Database software version you specify must be the same or later version as that of the backed-up database.
- If the backup being used to create a database is in a security zone compartment, the database cannot be created in a compartment that is not in a security zone. See the Security Zone Policies topic for a full list of policies that affect Database service resources.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster that contains the source database you are using to create the new database:
   - *Cloud VM clusters* (new resource model) Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
   - *DB systems* Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. Under **Databases**, click the name of the database you are using as the source for the new database.
5. On the Database Details page, click **Create Database from Last Backup**.
6. In the **Provide basic information for the Exadata infrastructure** section:

- **Select an availability domain:** It could be the same as the availability domain that hosts the backup or a different one within the same region.
- **Select Exadata infrastructure:** Select an Exadata infrastructure from the chosen compartment. Click the **Change Compartment** hyperlink to choose a different compartment.

7. On the **Create Database from Backup** page, configure the database as follows.
8. In the **Configure your DB system** section: *Backups created in cloud VM clusters:* Choose a cloud VM cluster to run the database from the **Select a VM cluster** drop-down list.
   - *Backups created in cloud VM clusters:* Choose a cloud VM cluster to run the database from the **Select a VM cluster** drop-down list.
   - *Backups created in DB systems:* Choose a shape from the **Select a shape** drop-down list, then choose a DB system to run the database from the **Select a DB system** drop-down list.
9. In the **Configure Database Home** section:
   - **Select an existing Database Home**: If you choose this option, make a selection from the **Select a Database Home** drop-down list.
   - **Create a new Database home**: If you choose this option, enter a name for the new Database Home in the **Database Home display name** field. Click **Change Database Image** to select a database software image for the new Database Home. In the **Select a Database Software Image** panel, do the following:
     a. Select the compartment containing the database software image you want to use to create the new Database Home.
     b. Select the Oracle Database software version that the new Database Home will use, then choose an image from the list of available images for your selected software version.
     c. Click **Select**.
10. In the **Configure database** section:

> **Note:**
>
> You cannot modify the `db_name`, `db_unique_name`, and SID prefix after creating the database.

- **Database name:** The name for the database. The database name must meet the requirements:
  - Maximum of 8 characters
  - Contain only alphanumeric characters
  - Begin with an alphabetic character
  - Cannot be part of first 8 characters of a DB_UNIQUE_NAME on the VM cluster
  - DO NOT use the following reserved names: grid, ASM
- **Database unique name:** Optionally, specify a value for the `DB_UNIQUE_NAME` database parameter. The value is case insensitive.

  The unique name must meet the requirements:
  - Maximum of 30 characters

- Contain only alphanumeric or underscore (_) characters
- Begin with an alphabetic character
- Unique across the VM cluster. Recommended to be unique across the tenancy.

If not specified, the system automatically generates a unique name value, as follows:

```
<db_name>_<3_chars_unique_string>_<region-name>
```

- **Administrator username:** This read-only field displays the username for the administrator, "sys".
- In the **Password** and **Confirm password** fields, enter and re-enter a password. A strong password for SYS administrator must be 9 to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters. The special characters must be _, #, or -. The password must not contain the user name (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reverse order and regardless of casing.

11. In the **Enter the source database's TDE wallet or RMAN password** field, enter a password that matches either the Transparent Data Encryption (TDE) wallet password or RMAN password for the source database.
12. Click **Show Advanced Options** to specify advanced options for the database.
    - **Management**
      Oracle SID prefix: The Oracle Database instance number is automatically added to the SID prefix to create the INSTANCE_NAME database parameter. he INSTANCE_NAME parameter is also known as the SID. The SID is unique across the cloud VM cluster. If not specified, SID prefix defaults to the first 12 characters of the `db_name`.

      > **Note:**
      > Entering an SID prefix is only available for Oracle 12.1 databases and above.

      The SID prefix must meet the requirements:
      - Maximum of 12 characters
      - Contain only alphanumeric characters
      - Begin with an alphabetic character
      - Unique in the VM cluster
      - DO NOT use the following reserved names: grid, ASM
13. Click **Create Database**.

To move a database to another Database Home

This task explains how to patch a single Oracle Database in your Exadata Cloud Infrastructure instance by moving it to another Database Home.

You can move a database to any Database Home that meets at either of the following criteria:

- The target Database Home uses the same Oracle Database software version (including patch updates) as the source Database Home

- The target Database Home is based on either the latest version of the Oracle Database software release used by the database, or one of the three prior versions of the release

Moving a database to a new Database Home brings the database up to the patch level of the target Database Home. For information on patching Database Homes, see Database Home Patching and .

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the database you want to move.:
   Cloud VM clusters (The New Exadata Cloud Infrastructure Resource Model *):* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, click the name of the VM cluster that contains the database you wan to move.

   DB systems: Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find you want to access, and then click the name of the Exadata DB system that contains the database you want to move..
4. Click **More Actions**, then click **Move to Another Home**.
5. Select the target Database Home.
6. Click **Move Database**.
7. Confirm the move operation.

   The database is moved in a rolling fashion. The database instance will be stopped, node by node, in the current home and then restarted in the destination home. While the database is being moved, the Database Home status displays as **Moving Databse**. When the operation completes, Database Home is updated with the current home. Datapatch is executed automatically, as part of the database move, to complete post-patch SQL actions for all patches, including one-offs, on the new Database Home. If the database move operation is unsuccessful, then the status of the database displays as `Failed`, and the Database Home field provides information about the reason for the failure.

To terminate a database

You'll get the chance to back up the database prior to terminating it. This creates a standalone backup that can be used to create a database later. We recommend that you create this final backup for any production (non-test) database.

> **Note:**
>
> Terminating a database removes all automatic incremental backups of the database from Oracle Cloud Infrastructure Object Storage. However, all full backups that were created on demand, including your final backup, will persist as standalone backups.

You cannot terminate a database that is assuming the primary role in a Data Guard association. To terminate it, you can switch it over to the standby role.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.

3. Navigate to the database:

   Cloud VM clusters (The New Exadata Cloud Infrastructure Resource Model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems: Under **Oracle Base Database**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

   On the cloud VM cluster or DB system details page, in the Databases table, click the name of the database to display the Database Details page.

4. Click **More Actions**, and then click **Terminate**.
   For the database using Oracle Cloud Infrastructure Object Storage or Oracle Database Autonomous Recovery Service: In the confirmation dialog,
   - Review the message about the backup retention policy.
   - Configure automatic backups as needed.
   - Type the name of the database to confirm the termination

5. Click **Terminate Database**.
   The database's status indicates Terminating.

> **Note:**
>
> The database stays in a terminated state with backups listed until all backups are expired.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

To administer Vault encryption keys

Use this procedure to rotate the Vault encryption key or or change the encryption management configuration.

After you provision a database in an Exadata DB system or cloud VM cluster, you can rotate the Vault encryption key or change the encryption management configuration for that database.

> **Note:**
>
> - To ensure that your Exadata database uses the most current version of the Vault encryption key, rotate the key from the database details page on the Oracle Cloud Infrastructure Console. Do not use the Vault service.
> - You can rotate Vault encryption keys only on databases that are configured with customer-managed keys.
> - You can change encryption key management from Oracle-managed keys to customer-managed keys but you cannot change from customer-managed keys to Oracle-managed keys.
> - Oracle supports administering encryption keys on databases after Oracle Database 11*g* release 2 (11.2.0.4).

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your compartment from the **Compartment** drop-down.
3. Navigate to the cloud VM cluster that contains the database for which you want to change encryption management or to rotate a key.
   Cloud VM clusters: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, locate the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. In the **Databases** section, click the name of the database for which you want to change encryption management or to rotate a key to display its details page.
5. Click the **More Actions** drop-down.
6. Click **Administer Encryption Key**.
   To rotate an encryption key on a database using customer-managed keys:
   a. Click **Rotate Encryption Key** to display a confirmation dialog.
   b. Click **Rotate Key**.

   To change key management type from Oracle-managed keys to customer-managed keys:
   a. Click **Change Key Management Type**.
   b. Select **Use customer-managed keys**.
   You must have a valid encryption key in Oracle Cloud Infrastructure Vault service and provide the information in the subsequent steps. See Key and Secret Management Concepts.
   c. Choose a vault from the **Vault in *compartment*** drop-down. You can change the compartment by clicking the **Change Compartment** link.
   d. Select an encryption key from the **Master encryption key in *compartment*** drop-down. You can change the compartment containing the encryption key you want to use by clicking the **Change Compartment** link.
   e. If you want to use an encryption key that you import into your vault, then select **Change Compartment** and enter the OCID of the key you want to use in the **Key version OCID** field.
7. Click **Apply**.

> **Note:**
>
> Changing key management causes the database to become briefly unavailable.

> **Caution:**
>
> After changing key management to customer-managed keys, do not delete the encryption key from the vault as this can cause the database to become unavailable.

On the database details page for this database, the **Encryption** section displays the encryption key name and the encryption key OCID.

Known Issues in Exadata Cloud Infrastructure

`rac stopdb` failed

`rac stopdb` failed to stop db

When GI version is 19.17 then creating a database against 11.2.0.4 Oracle home with patchsets July ’22 RU or older will fail with error mentioned in bug#28326679

Example:

ERROR : rac stopdb, failed to stop db viacmd export ORACLE_HOME=/u02/app/oracle/product/11.2.0/dbhome_1 ;/u02/app/oracle/product/11.2.0/dbhome_1/bin/srvctl stop database -d db008077-o immediate, out : PRCD-1120 : The resource for database db008077 could notbe found. PRCR-1001 : Resource ora.db008077.db does not exist, err :1 }

Solution:

Option 1: (Create new oracle home with Custom Image):

- Create custom image for 11.2.0.4 with patchsets July ’22 RU or older along with bug#28326679 one off
- Create Oracle home using above customer image
- Create database against the home

Option 2 (Apply one-off to existing Oracle home) :

- Download the patch for bug#28326679
- Apply the patch using opatch

Applicability:

- For ExaCS and ExaCC-Gen2, Both options given above will work.
- For ExaCC – Gen1, Option 2 (Apply one-off to existing Oracle home) will work.

Using the API to manage Databases

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.

Use these API operations to manage databases.

- ListDatabases
- GetDatabase
- CreateDatabase
- UpdateDatabase - Use this operation to move a database to another Database Home
- DeleteDatabase

For the complete list of APIs for the Database service, see Database Service API.

Create and Manage Exadata Pluggable Databases

You can create and manage pluggable databases (PDBs) in Exadata Cloud Infrastructure using the Console and APIs.

In this documentation, "database" refers to a container database, also called a CDB. For more information on these resource types, see Multitenant Architecture in the Oracle Database documentation. See Provisioning and Managing Exadata Databases for information on container databases in Exadata Cloud Infrastructure.

Oracle 19c or later databases created in Exadata Cloud Infrastructure include an initial PDB that you can access from the Database Details page in the Console. You can create and manage additional PDBs in the database using the Console or APIs.

- **Backup**
  You can take a backup of the PDB optionally during create, clone, or relocate operations when the CDB is configured with the auto-backup feature. The PDB backup destination will always be the same as CDB, and the backups cannot be accessed directly or created on demand. Oracle recommends immediately backing up the PDB after you create or clone it. This is because the PDB will not be recoverable until the next daily auto-backup completes successfully, leading to a possible data loss.
- **Restore**
  - **Base Database Service / Oracle Exadata Database Service on Dedicated Infrastructure:**
    - **In place restore:** You can restore a PDB within the same CDB to last known good state or to a specified timestamp.
    - **Out of place restore:** You can restore a PDB by creating a database (CDB) from the backup, then selecting a PDB or a subset of them you want to restore on the new database.
  - **Oracle Exadata Database Service on Cloud@Customer:**
    - **In place restore:** You can restore a PDB within the same CDB to last known good state and specified timestamp.
    - **Out of place restore:** It's not available.

    You can perform an in-place restore when you want to move a PDB back to a specified state or time. Both the CDB and PDB must be up and running and only one PDB can be restored at a time.
    - If you have multiple PDBs in your CDB and want to restore multiple of them to the same CDB, then you could restore each individual PDB, one PDB at a time, from the CDB backup.

    - When the CDB is down, you could restore the complete CDB and all the PDBs in that CDB will also be restored.
    - You could either restore the database to the specified timestamp or to its last known good state.

- **Relocate**
  You can relocate a PDB from one CDB to another CDB within the same availability domain (AD):
  - Across compartments, VM clusters, DB system (for BaseDB only), or VCNs (not applicable to ExaDB-C@C). If two different VCNs are used, then both VCNs must be peered before relocating.
  - To the same or a higher database version.

  During relocate, the PDB will be removed from the source CDB and moved to the destination CDB that is up and running. In a Data Guard association, a PDB relocated to the primary will be synchronized with the standby as well.
- **Clone**
  A clone is an independent and complete copy of the given database as it existed at the time of the cloning operation. You can create clones of your PDB within the same CDB or a different CDB and refresh the cloned PDB.

  The following types of clones are supported:
  - **Local clone:** A copy of the PDB is created within the same CDB.
  - **Remote clone:** A copy of the PDB is created in a different CDB.

  You can perform a remote clone of a PDB from one CDB to another CDB within the same availability domain (AD):
  - Across compartments, VM clusters, DB system (for BaseDB only), or VCNs (not applicable to ExaDB-C@C). If two different VCNs are used, then both VCNs must be peered before cloning.
  - To the same or a higher database version.
  - **Refreshable clone:** A copy of the PDB is created in a different CDB, and you will be able to refresh the cloned PDB.
    You can perform a refreshable clone of a PDB from one CDB to another CDB within the same availability domain (AD):
    - Across compartments, VM clusters, DB system (for BaseDB only), or VCNs (not applicable to ExaDB-C@C). If two different VCNs are used, then both VCNs must be peered before cloning.
    - To the same or a higher database version.
- **Refreshable Clone**
  A refreshable clone enables you to keep your remote clone updated with the source PDB. You can only refresh while the PDB is in mount mode. The only open mode you can have is read-only and refresh cannot be done while it is in read-only mode.
  - A database link user credential is required for creating a refreshable clone.
  - Clone, relocate, and in-place restore operations are not supported in the refreshable clone. Relocate and in-place restore operations are not supported in the source, and the source can only be deleted after disconnecting or deleting the refreshable clone.
  - In a Data Guard association, a refreshable clone cannot be created on standby, but it can be created on the primary. However, the primary will not be synced to the standby.

> **Note:**
>
> A PDB in standby cannot be used as the source for a refreshable PDB.

- **Convert Refreshable PDB to Regular PDB**
  You can convert a refreshable PDB to a regular PDB by disconnecting the refreshable clone (destination PDB) from the source PDB at any time. If the refresh PDB is in a Data Guard association, when it is converted to a regular PDB the PDB will be synced to the standby as part of the conversion process.
- **Open Modes**
  On the Console, you can see the open modes of a PDB, such as read-write, read-only, and mounted. If the PDB status is the same across all nodes, the system displays the same status for all PDBs. If the PDB statuses are different across the nodes, the system displays a message indicating on which nodes the PDBs are opened in read-write mode. You cannot change the open mode of a PDB through the API or Console. However, you can start or stop a PDB. Starting the PDB will start it in read-write mode. Stopping the PDB will close it and it will remain in mount mode.
- Limitations for Pluggable Database Management
- Creating an Exadata Pluggable Database
- Managing an Exadata Pluggable Database
  This topic includes the procedures to connect to, start, stop, and delete a pluggable database (PDB).
- Cloning an Exadata Pluggable Database
  You can create local, remote, and refreshable clones.

Limitations for Pluggable Database Management

- New PDBs created with SQL are not immediately discovered by OCI's control plane and displayed in the Console. However, OCI does perform a sync operation on a regular basis to discover manually-created PDBs, and they should be visible in the Console and with API-based tools within 45 minutes of creation. Oracle recommends using the Console or API-based tools (including the OCI CLI , SDKs, and Terraform) to create PDBs.
- Pluggable database operations are supported only for databases using Oracle Database 19c and later.
- PDBs are backed up at the CDB level when using the OCI Console or APIs, and each backup includes all the PDBs in the database. However, the dbaascli utility's dbaascli database backup command allows you to create backups of specified PDBs. See Using the dbaascli Utility on Exadata Cloud Infrastructure for more information.
- Restore operations are performed at the CDB level when using the OCI Console or APIs. However, the dbaascli utility's dbaascli pdb recover command allows you to restore backups of specified PDBs. See Using the dbaascli Utility on Exadata Cloud Infrastructure for more information.

Creating an Exadata Pluggable Database

You can create a pluggable database (PDB) in Exadata Cloud Service from the OCI Console, or with the APIs and API-based tools (the OCI CLI, SDKs, and Terraform). PDBs must be created one at a time. During the PDB create operation, the parent database (CDB) is in the "Updating" state. Creating a new PDB has no impact on existing PDBs in the database.

- Using the console to create pluggable database
- Using the console to relocate a pluggable database
- Using the API to create pluggable database

Using the console to create pluggable database

> **Note:**
>
> Creating a pluggable database (PDB) is not supported for databases using Data Guard.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the database:

   Cloud VM clusters (new resource model) Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

   On the cloud VM cluster or DB system details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. On the Database Details page, click **Pluggable Databases** in the **Resources** section of the page.
5. Click **Create Pluggable Database**.
6. In the **Create Pluggable Database** dialog, enter the following:
   - **PDB Name**: Enter a name for the PDB. The name must begin with an alphabetic character and can contain a maximum of 30 alphanumeric characters. Note: For bare metal DB systems, you cannot have two PDBs in the same database that use the same PDB name. You can use the same name for PDBs in different databases within the same DB system.
   - **Unlock my PDB Admin account**: *Optional.* Select this option to specify a PDB Admin password and configure the PDB to be unlocked at creation.
   - **PDB Admin password**: If you clicked **Unlock my PDB Admin** account, create and enter a PDB admin password. The password must contain:
     - A minimum of 9 and a maximum of 30 characters
     - At least two uppercase characters
     - At least two lowercase characters
     - At least two special characters. The valid special characters are: underscore (_), a hash sign (#), and a dash (-). You can use two of the same characters or any combination of two of the same characters.
     - At least two numeric characters (0 - 9)

- **Confirm PDB Admin password**: Reenter the PDB admin password.
- **TDE wallet password**: *Applicable only to databases using Oracle-managed encryption keys*. Enter the TDE wallet password for the parent CDB.
- **Take a backup of the PDB immediately after creating it:** You must enable auto-backup on the CDB to back up a PDB immediately after creating it. This check box is checked by default if auto-backup was enabled on the CDB.

> **Note:**
>
> If the check box is unchecked, the system displays a warning stating that PDB cannot be recovered until the next daily backup has been successfully completed.

7. Click **Create Pluggable Database**.

WHAT NEXT?

After creating your PDB, you can get connection strings for the administrative service using the OCI Console.

Using the console to relocate a pluggable database

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the database:

   Cloud VM clusters (new resource model) Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

   On the cloud VM cluster or DB system details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. On the Database Details page, click **Pluggable Databases** in the **Resources** section of the page.
5. Click the name of the PDB that you want to relocate.
   From the Pluggable Database details page, click **More Actions**, and then select **Relocate**.

   (or)

   Click the Actions menu (three dots) and select **Relocate**.
6. In the resulting Relocate Pluggable Database window, enter the following:
   - **VM Cluster:** Use the menu to select the destination VM cluster.
   - **Destination database:** Use the menu to select an existing database where the PDB will be created. This database can be of the same version as the CDB the source PDB is in or of a higher version.

- **New PDB name for the clone:** The name must begin with an alphabetic character and can contain up to 30 characters. To keep the PDB name the same, just re-enter the source PDB name.
- **Database TDE wallet password:** Enter the TDE wallet password for the parent CDB of the source PDB.
- **Unlock my PDB Admin Account:**
  - To enter the administrator's password, check this check box.
    - **PDB Admin Password:** Enter PDB admin password. The password must contain:
      - a minimum of 9 and a maximum of 30 characters
      - at least two uppercase characters
      - at least two lowercase characters
      - at least two special characters. The valid special characters are underscore (_), a pound or hash sign (#), and dash (-). You can use two of the same characters or any combination of two of the same characters.
      - at least two numeric characters (0 - 9)
    - **Confirm PDB Admin Password:** Enter the same PDB Admin password in the confirmation field.
  - To skip entering the administrator's password, uncheck this check box. If you uncheck this check box, then the PDB is created but you cannot use it. To use the PDB, you must reset the administrator password.

> **Note:**
>
> When you create a new PDB, a local user in the PDB is created as the administrator and granted the `PDB_DBA` role locally to the administrator.

To reset the password:

a. Connect to the container where your PDB exists using the SQL*Plus `CONNECT` statement.

```
SQL> show con_name;
CON_NAME
------------------------
CDB$ROOT
```

For more information, see *Administering a CDB* and *Administering PDBs* in the *Oracle® Multitenant Administrator's Guide*.

b. Find the administrator name of your PDB:

```
SQL> select grantee from cdb_role_privs where con_id = (select
con_id from cdb_pdbs where pdb_name = '<PDB_NAME>') and
granted_role = 'PDB_DBA';
```

c. Switch into your PDB:

```
SQL> alter session set container=<PDB_NAME>;
Session altered.
SQL> show con_name;
CON_NAME
------------------------
<PDB_NAME>
```

d. Reset the PDB administrator password:

```
SQL> alter user <PDB_Admin> identified by <PASSWORD>;
User altered.
```

- **Source database SYS password:** Enter the database admin password.
- **Database link:** Enter the user name and password for the database link. Note that the user must be precreated in the source database. The DB link will be created in the destination using that username and password.
- **Take a backup of the PDB immediately after creating it:** You must enable auto-backup on the CDB to back up a PDB immediately after creating it. This check box is checked by default if auto-backup was enabled on the CDB.

> **Note:**
>
> If the checkbox is unchecked, the system displays a warning stating that PDB cannot be recovered until the next daily backup has been successfully completed.

- **Advanced Options:**
  - **Tags:** Optionally, you can apply tags. If you have permission to create a resource, you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

7. Click **Relocate pluggable database**.

> **Note:**
>
> Relocate will incur downtime during the process and that the time required is based on the size of the PDB.

Using the API to create pluggable database

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.

Use the CreatePluggableDatabase API to create pluggable databases on Exadata Cloud Infrastructure.

For the complete list of APIs for the Database service, see Database Service API.

Managing an Exadata Pluggable Database

This topic includes the procedures to connect to, start, stop, and delete a pluggable database (PDB).

It also includes instructions for getting PDB connection strings for the administrative service.

- To start a pluggable database
- To stop a pluggable database
- To delete a pluggable database
- To get connection strings for a pluggable database
- Using the API to manage pluggable databases

To start a pluggable database

> **Note:**
>
> The PDB must be available and stopped to use this procedure.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

   *Cloud VM clusters (*new resource model*)* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

   On the cloud VM cluster or DB system details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to start. Click the PDB name to display details about it.
6. Click **Start**.
7. In the **Start PDB** dialog, click **Start PDB** to confirm the start operation.

To stop a pluggable database

> **Note:**
>
> The PDB must be available and running (started) to use this procedure.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

   *Cloud VM clusters (*new resource model*)* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

   On the cloud VM cluster or DB system details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to stop. Click the PDB name to display details about it.
6. Click **Start**.
7. In the **Stop PDB** dialog, click **Stop PDB** to confirm the stop operation.

To delete a pluggable database

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

   *Cloud VM clusters (*new resource model*)* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

   On the cloud VM cluster or DB system details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to delete. Click the PDB name to display details about it.
6. Click **More Actions**, then choose **Delete**.
7. In the **Delete PDB** dialog box, enter the name of the PDB that you want to delete to confirm the action, then click **Delete PDB**.

To get connection strings for a pluggable database

> **Note:**
>
> This topic explains how to get connection strings for the administrative service of a PDB. Oracle recommends that you connect applications to an application service, using strings created for the application service.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

   *Cloud VM clusters (*new resource model*)* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

   On the cloud VM cluster or DB system details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the PDB, and then click its name to display details about it.
6. Click **PDB Connection**.
7. In the **Pluggable Database Connection** dialog, use the **Show** and **Copy** links to display and copy connection strings, as needed.
8. Click **Close** to exit the dialog.

Using the API to manage pluggable databases

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.

Use these APIs to manage pluggable databases.

- ListPluggableDatabases
- GetPluggableDatabase
- StartPluggableDatabase
- StopPluggableDatabase
- UpdatePluggableDatabase
- DeletePluggableDatabase

> **Note:**
>
> Use the GetPluggableDatabase API to get administration service connection strings and other details about a PDB.

For the complete list of APIs for the Database service, see Database Service API.

Cloning an Exadata Pluggable Database

You can create local, remote, and refreshable clones.

A clone is an independent and complete copy of the given database as it existed at the time of the cloning operation. You can create clones of your PDB within the same CDB or a different CDB and also refresh the cloned PDB.

The following types of clones are supported:

- **Local clone:** A clone of the PDB is created within the same CDB.
- **Remote clone:** A clone of the PDB is created in a different CDB.
- **Refreshable clone:** A clone of the PDB is created in a different CDB, and you will be able to refresh the cloned PDB.

- Using the Console to Create a Local Clone of a Pluggable Database (PDB)
- Using the Console to Create a Remote Clone of a Pluggable Database (PDB)
- Using the Console to Create a Refreshable Clone of a Pluggable Database (PDB)
- Using the Console to Refresh a Cloned Pluggable Database (PDB)
- Using the Console to Convert a Refreshable Clone to a Regular Pluggable Database (PDB)
- Using the API to clone a pluggable database

Using the Console to Create a Local Clone of a Pluggable Database (PDB)

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

   *Cloud VM clusters (*new resource model*)* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

   On the cloud VM cluster or DB system details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to clone, and then click its name to display details about it.

6. Click **Clone**.
7. In the **Clone PDB** dialog box, enter the following:
   - **Select clone type:** Select **Local clone** to create a copy of the source PDB to the same CDB.
   - **Exadata VM Cluster**: Use the menu to select the cloud VM cluster of the target database.

     > **Note:**
     >
     > The target VM Cluster may be on a different Exadata infrastructure.

   - **Destination database**: This field is disabled.
   - **PDB name**: Provide a name for the new cloned PDB. The name must begin with an alphabetic character and can contain up to 30 characters.
   - **Database TDE wallet password**: *Not applicable for databases using customer-managed keys from the Vault service.* Enter the TDE wallet password for the parent database (CDB) of the source PDB.
   - **Unlock my PDB Admin account**: *Optional.* Select this option to specify a PDB Admin password and configure the PDB to be unlocked at creation.
   - **PDB Admin password**: Create and enter a new PDB Admin password. The password must contain:
     - 9–30 characters
     - At least two uppercase characters
     - At least two lowercase characters
     - At least two special characters. The valid special characters are: underscore (_), a hash sign (#), and a dash (-). You can use two of the same characters or any combination of two of these characters.
     - At least two numeric characters (0-9)
   - **Confirm PDB Admin password**: Enter the PDB Admin password again to confirm.
   - **Take a backup of the PDB immediately after creating it**: You must enable auto-backup on the CDB to back up a PDB immediately after creating it. This check box is checked by default if auto-backup was enabled on the CDB.

     > **Note:**
     >
     > If the checkbox is unchecked, the system displays a warning stating that PDB cannot be recovered until the next daily backup has been successfully completed.

   - **Advanced Options:**
     - **Tags:** Optionally, you can apply tags. If you have permission to create a resource, you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see Resource Tags. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

8. Click **Clone pluggable database**.

Using the Console to Create a Remote Clone of a Pluggable Database (PDB)

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

   *Cloud VM clusters (*new resource model*)* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

   On the cloud VM cluster or DB system details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to clone, and then click its name to display details about it.
6. Click **Clone**.
7. In the **Clone PDB** dialog box, enter the following:
   - **Select clone type:** Select **Remote clone** to create a copy of the source PDB to the same CDB.
   - **Exadata VM Cluster**: Use the menu to select the cloud VM cluster of the target database.

     > **Note:**
     >
     > The target VM Cluster may be on a different Exadata infrastructure.

   - **Destination database**: Use the menu to select an existing database where the PDB will be created. This database can be of the same version as the CDB the source PDB is in or of a higher version.
   - **PDB name**: Provide a name for the new cloned PDB. The name must begin with an alphabetic character and can contain up to 30 characters.
   - **Database TDE wallet password**: *Not applicable for databases using customer-managed keys from the Vault service.* Enter the TDE wallet password for the parent database (CDB) of the source PDB.
   - **Unlock my PDB Admin account**: *Optional.* Select this option to specify a PDB Admin password and configure the PDB to be unlocked at creation.
   - **PDB Admin password**: Create and enter a new PDB Admin password. The password must contain:
     - 9–30 characters
     - At least two uppercase characters
     - At least two lowercase characters

  - At least two special characters. The valid special characters are: underscore (_), a hash sign (#), and a dash (-). You can use two of the same characters or any combination of two of these characters.
  - At least two numeric characters (0-9)
- **Confirm PDB Admin password**: Enter the PDB Admin password again to confirm.
- **Database link**: Enter the user name and password for the database link. Note that the user must be precreated in the source database. The DB link will be created in the destination using that username and password.
- **Take a backup of the PDB immediately after creating it**: You must enable auto-backup on the CDB to back up a PDB immediately after creating it. This check box is checked by default if auto-backup was enabled on the CDB.

> **Note:**
>
> If the checkbox is unchecked, the system displays a warning stating that PDB cannot be recovered until the next daily backup has been successfully completed.

- **Advanced Options:**
  - **Tags:** Optionally, you can apply tags. If you have permission to create a resource, you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see Resource Tags. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

8. Click **Clone pluggable database**.

Using the Console to Create a Refreshable Clone of a Pluggable Database (PDB)

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

   *Cloud VM clusters (*new resource model*)* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

   On the cloud VM cluster or DB system details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to clone, and then click its name to display details about it.
6. Click **Clone**.
7. In the **Clone PDB** dialog box, enter the following:

- **Select clone type:** Select Refreshable clone to create a copy of the source PDB to the same CDB.
  For more information about refreshable clones, see About Refreshable Clone PDBs.
- **Exadata VM Cluster**: Use the menu to select the cloud VM cluster of the target database.

> **Note:**
>
> The target VM Cluster may be on a different Exadata infrastructure.

- **Destination database**: Use the menu to select an existing database where the PDB will be created. This database can be of the same version as the CDB the source PDB is in or of a higher version.
- **PDB name**: Provide a name for the new cloned PDB. The name must begin with an alphabetic character and can contain up to 30 characters.
- **Database TDE wallet password**: *Not applicable for databases using customer-managed keys from the Vault service.* Enter the TDE wallet password for the parent database (CDB) of the source PDB.
- **Unlock my PDB Admin account**: *Optional.* Select this option to specify a PDB Admin password and configure the PDB to be unlocked at creation.
- **PDB Admin password**: Create and enter a new PDB Admin password. The password must contain:
  - 9–30 characters
  - At least two uppercase characters
  - At least two lowercase characters
  - At least two special characters. The valid special characters are: underscore (_), a hash sign (#), and a dash (-). You can use two of the same characters or any combination of two of these characters.
  - At least two numeric characters (0-9)
- **Confirm PDB Admin password**: Enter the PDB Admin password again to confirm.
- **Database link**: Enter the user name and password for the database link. Note that the user must be precreated in the source database. The DB link will be created in the destination using that username and password.
- **Take a backup of the PDB immediately after creating it**: You must enable auto-backup on the CDB to back up a PDB immediately after creating it. This check box is checked by default if auto-backup was enabled on the CDB.

> **Note:**
>
> If the checkbox is unchecked, the system displays a warning stating that PDB cannot be recovered until the next daily backup has been successfully completed.

- **Advanced Options:**
  - **Tags:** Optionally, you can apply tags. If you have permission to create a resource, you also have permission to apply free-form tags to that resource. To apply a

defined tag, you must have permission to use the tag namespace. For more information about tagging, see Resource Tags. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

8. Click **Clone pluggable database**.

Using the Console to Refresh a Cloned Pluggable Database (PDB)

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

   *Cloud VM clusters (*new resource model*)* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

   On the cloud VM cluster or DB system details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to refresh, and then click its name to display details about it.
6. Click **More Actions** and select **Refresh**.
7. In the resulting **Refresh** dialog box, click **Refresh** to confirm.

Using the Console to Convert a Refreshable Clone to a Regular Pluggable Database (PDB)

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

   *Cloud VM clusters (*new resource model*)* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

   On the cloud VM cluster or DB system details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to convert to a regular PDB, and then click its name to display details about it.
6. In the resulting **Convert to regular PDB** dialog box, enter the following:

- **Database TDE wallet password**: *Not applicable for databases using customer-managed keys from the Vault service.* Enter the TDE wallet password for the parent database (CDB) of the source PDB.
- **Take a backup of the PDB immediately after creating it**: You must enable auto-backup on the CDB to back up a PDB immediately after creating it. This check box is checked by default if auto-backup was enabled on the CDB.

> **Note:**
>
> If the checkbox is unchecked, the system displays a warning stating that PDB cannot be recovered until the next daily backup has been successfully completed.

7. Click **Convert**.

Using the API to clone a pluggable database

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.

Use these APIs to clone pluggable databases:

- LocalclonePluggableDatabase
- RemoteclonePluggabledatabase

For the complete list of APIs for the Database service, see Database Service API.

Restoring an Exadata Pluggable Database

You can perfrom in-place and out of place restore of an Exadata pluggable database.

The following types of clones are supported:

- **In place restore:** You can restore a PDB within the same CDB to last known good state or to a specified timestamp.
- **Out of place restore:** You can restore a PDB by creating a database (CDB) from the backup, then selecting a PDB or a subset of them you want to restore on the new database.
- Using the Console to Perform an In-Place Restore of a Pluggable Database (PDB)
- Using the Console to Perform an Out-of-Place Restore of a Pluggable Database (PDB)

Using the Console to Perform an In-Place Restore of a Pluggable Database (PDB)

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

   *Cloud VM clusters (*new resource model*)* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

DB systems Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

On the cloud VM cluster or DB system details page, in the **Databases** table, click the name of the database to display the Database Details page.

4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to restore, and then click its name to display details about it.
6. In the resulting Restore PDB dialog, enter the following:
   - **Restore to latest:** Select this option to restore and recover the database with zero, or least possible, data loss.
   - **Restore to a timestamp:** Select this option to restore and recover the database to the specified timestamp.
7. Click **Restore**.

Using the Console to Perform an Out-of-Place Restore of a Pluggable Database (PDB)

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

   *Cloud VM clusters (*new resource model*)* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

   On the cloud VM cluster or DB system details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to restore, and then click its name to display details about it.
6. Under **Resources**, click **Backups**.
7. From the list of backups, choose a backup, click the Actions menu (three dots), and then select **Create Database**.
8. In the resulting Create database from backup dialog box, select either of these options, **Select all PDBs** or **Specify the PDBs to restore**.

- To create a database by selecting all Pluggable Databases
  Provide the requested information in the Create database from backup page:
- To create a database by specifying a subset of Pluggable Databases
  Provide the requested information in the Create database from backup page:

To create a database by selecting all Pluggable Databases

Provide the requested information in the Create database from backup page:

1. Click **Select all PDBs**.
2. Click **Next**.
3. Select the VM cluster where you want to create the database.
   Click the **Change Compartment** hyperlink to choose your compartment.
4. **Configure Database Home**: Select an existing Database Home or create one as applicable. Note that this field is not available when you create a Database from the Database Home details page.
   - **Select an existing Database Home**: If one or more Database Homes already exist for the database version you have selected, then this option is selected by default. And, you will be presented with a list of Database Homes. Select a Database Home from the list.
   - **Create a new Database Home**: If no Database Homes exist for the database version you have selected, then this option is selected by default.
     a. Enter **Database Home display name**.
     b. Click **Change Database Image** to select your software version.
        Select a **Database Software Image** window is displayed.
     c. Select an **Image Type**, **Oracle Provided Database Software Images**, or **Custom Database Software Images**.
        If you choose **Oracle Provided Database Software Images**, then you can use the **Display all available version** switch to choose from all available PSUs and RUs. The most recent release for each major version is indicated with a **latest** label.

        Note

        For the Oracle Database major version releases available in Oracle Cloud Infrastructure, images are provided for the current version plus the three most recent older versions (N through N - 3). For example, if an instance is using Oracle Database 19c, and the latest version of 19c offered is 19.8.0.0.0, images available for provisioning are for versions 19.8.0.0.0, 19.7.0.0, 19.6.0.0 and 19.5.0.0.
5. **Provide the database name**: Specify a user-friendly name that you can use to identify the database. The database name must contain only the permitted characters.

   Review the following guidelines when selecting a database name.
   - maximum of 8 characters
   - contain only alphanumeric characters
   - begin with an alphabetic character
   - cannot be part of first 8 characters of a `db_unique_name` on the VM cluster
   - unique within a VM cluster
   - **DO NOT** use `grid` because `grid` is a reserved name
   - **DO NOT** use `ASM` because `ASM` is a reserved name
6. **Provide a unique name for the database**: Optionally, specify a unique name for the database. This attribute defines the value of the `db_unique_name` database parameter. The value is case insensitive.

The `db_unique_name` must contain only the permitted characters. Review the following guidelines when selecting a database name.

- maximum of 30 characters
- can contain alphanumeric and underscore (_) characters
- begin with an alphabetic character
- unique across the fleet/tenancy

If a unique name is not provided, then the `db_unique_name` defaults to the following format *<db_name>_<3 char unique string>_<region-name>*.

If you plan to configure the database for backup to a Recovery Appliance backup destination, then the unique database name must match the name that is configured in the Recovery Appliance.

7. **Provide the administration password**: Provide and confirm the Oracle Database administration password. This password is used for administration accounts and functions in the database, including:
   - The password for the Oracle Database `SYS` and `SYSTEM` users.
   - The Transparent Data Encryption (TDE) Keystore password.

   For Oracle Database 12c Release 1 or later releases, the password for the PDB administration user in the first PDB (`PDBADMIN`) must be nine to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters. The special characters must be `_`, `#`, or `-`. In addition, the password must not contain the name of the tenancy or any reserved words, such as `Oracle` or `Table`, regardless of casing.
8. **Enter the source database's TDE wallet or RMAN password:** Password must match the TDE wallet or RMAN password of the source database contained in the backup.
9. Click **Create Backup**.

To create a database by specifying a subset of Pluggable Databases

Provide the requested information in the Create database from backup page:

1. Click **Specify the PDBs to restore**.
2. In the **Specify PDB to restore** field, provide a comma-delimited list of PDBs to restore.
3. Click **Next**.
4. Select the VM cluster where you want to create the database.

   Click the **Change Compartment** hyperlink to choose your compartment.
5. **Configure Database Home**: Select an existing Database Home or create one as applicable. Note that this field is not available when you create a Database from the Database Home details page.
   - **Select an existing Database Home**: If one or more Database Homes already exist for the database version you have selected, then this option is selected by default. And, you will be presented with a list of Database Homes. Select a Database Home from the list.
   - **Create a new Database Home**: If no Database Homes exist for the database version you have selected, then this option is selected by default.
     a. Enter **Database Home display name**.
     b. Click **Change Database Image** to select your software version.

Select a **Database Software Image** window is displayed.

c. Select an **Image Type**, **Oracle Provided Database Software Images**, or **Custom Database Software Images**.
If you choose **Oracle Provided Database Software Images**, then you can use the **Display all available version** switch to choose from all available PSUs and RUs. The most recent release for each major version is indicated with a **latest** label.

Note

For the Oracle Database major version releases available in Oracle Cloud Infrastructure, images are provided for the current version plus the three most recent older versions (N through N - 3). For example, if an instance is using Oracle Database 19c, and the latest version of 19c offered is 19.8.0.0.0, images available for provisioning are for versions 19.8.0.0.0, 19.7.0.0, 19.6.0.0 and 19.5.0.0.

6. **Provide the database name**: Specify a user-friendly name that you can use to identify the database. The database name must contain only the permitted characters.

   Review the following guidelines when selecting a database name.

   - maximum of 8 characters
   - contain only alphanumeric characters
   - begin with an alphabetic character
   - cannot be part of first 8 characters of a `db_unique_name` on the VM cluster
   - unique within a VM cluster
   - **DO NOT** use `grid` because `grid` is a reserved name
   - **DO NOT** use `ASM` because `ASM` is a reserved name

7. **Provide a unique name for the database**: Optionally, specify a unique name for the database. This attribute defines the value of the `db_unique_name` database parameter. The value is case insensitive.

   The `db_unique_name` must contain only the permitted characters. Review the following guidelines when selecting a database name.

   - maximum of 30 characters
   - can contain alphanumeric and underscore (_) characters
   - begin with an alphabetic character
   - unique across the fleet/tenancy

   If a unique name is not provided, then the `db_unique_name` defaults to the following format *<db_name>_<3 char unique string>_<region-name>*.

   If you plan to configure the database for backup to a Recovery Appliance backup destination, then the unique database name must match the name that is configured in the Recovery Appliance.

8. **Provide the administration password**: Provide and confirm the Oracle Database administration password. This password is used for administration accounts and functions in the database, including:

   - The password for the Oracle Database `SYS` and `SYSTEM` users.
   - The Transparent Data Encryption (TDE) Keystore password.

For Oracle Database 12c Release 1 or later releases, the password for the PDB administration user in the first PDB (`PDBADMIN`) must be nine to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters. The special characters must be `_`, `#`, or `-`. In addition, the password must not contain the name of the tenancy or any reserved words, such as `Oracle` or `Table`, regardless of casing.

9. **Enter the source database's TDE wallet or RMAN password:** Password must match the TDE wallet or RMAN password of the source database contained in the backup.
10. Click **Create Backup**.

Changing the Database Passwords

To change the SYS password, or to change the TDE wallet password, use this procedure.

The password that you specify in the **Database Admin Password** field when you create a new Exadata Cloud Infrastructure instance or database is set as the password for the SYS, SYSTEM, TDE wallet, and PDB administrator credentials. Use the following procedures if you need to change passwords for an existing database.

> **Note:**
>
> if you are enabling Data Guard for a database, then the SYS password and the TDE wallet password of the primary and standby databases must all be the same.

> **Note:**
>
> Using the `dbaascli` to change the SYS password will ensure the backup/restore automation can parallelize channels across all nodes in the cluster.

To Change the SYS Password for an Exadata Cloud Infrastructure Database

1. Log onto the Exadata Cloud Infrastructure virtual machine as `opc`.
2. Run the following command:

```
sudo dbaascli database changepassword --dbname database_name --user SYS
```

To Change Database Passwords in a Data Guard Environment

1. Run the following command on the primary database:

```
dbaascli database changePassword –dbName <dbname> --user SYS --
prepareStandbyBlob true --blobLocation <location to create the blob file>
```

2. Copy the blob file created to all the standby databases and update the file ownership to `oracle` user.
3. Run the following command on all the standby databases:

```
dbaascli database changePassword –dbName <dbname> --user SYS --
standbyBlobFromPrimary <location of copies the blob file>
```

To Change the TDE Wallet Password for an Exadata Cloud Infrastructure Database

1. Log onto the Exadata Cloud Infrastructure virtual machine as `opc`.
2. Run the following command:

```
sudo dbaascli tde changepassword --dbname database_name
```

Manage Database Backup and Recovery on Oracle Exadata Database Service on Dedicated Infrastructure

Learn how to work with the backup and recovery facilities provided by Oracle Exadata Database Service on Dedicated Infrastructure.

- Oracle Recommended Options to Perform Backup and Recovery Operations
  Oracle offers the following options for Oracle Database Backup and Recovery operations. These options are mutually exclusive.
- Managing Exadata Database Backups
  Automatic Exadata database backups are managed by Oracle Cloud Infrastructure. You configure this by using the Console or the API.
- Managed Backup Types and Usage Information
  There are two types of automatic Exadata database backups: Autonomous Recovery Service, and Oracle Object Storage.
- Default Backup Channel Allocation
  The default settings for database backup channels when using "Oracle Managed Backup" or "User Configured Backup"
- Prerequisites for Backups on Exadata Cloud Infrastructure
- Using the Console to Manage Backups
- To designate Autonomous Recovery Service as a Backup Destination for an Existing Database
  To designate Autonomous Recovery Service as a Backup Destination for an existing database, use this procedure.
- Recovering an Exadata Database from Backup Destination
  This topic explains how to recover an Exadata database from a backup stored in either Object Storage or Autonomous Recovery Service by using the Console or the API.
- Managing Exadata Database Backups by Using bkup_api
- Using the API to Manage Backup and Recovery
- Alternative Backup Methods
  Learn about alternative backup methods that are available in addition to the OCI Console.
- Recovering a Database Using Oracle Recovery Manager (RMAN)

Oracle Recommended Options to Perform Backup and Recovery Operations

Oracle offers the following options for Oracle Database Backup and Recovery operations. These options are mutually exclusive.

> **Note:**
>
> A hybrid configuration, that is, mixing the options is not supported. Mixing the options will break automation.

Option 1: Oracle Managed Backups

Oracle managed backups are entirely managed by Exadata Cloud Infrastructure (ExaDB-D) or Exadata Cloud@Customer (ExaDB-C@C) based on a one-time configuration. Besides being fully integrated into ExaDB-D or ExaDB-C@C cloud services Control Plane, these backups can also be accessed through OCI APIs. Oracle recommends this approach.

- The `dbaascli database backup` and `dbaascli database recover` commands can be used in conjunction with the automated backups for certain operations. For more information, see `dbaascli database backup` and `dbaascli database recover`.
- Customers are allowed to query RMAN views or issue RMAN restore and recovery commands, for example, table, datafile, or tablespace recovery commands.

> **Note:**
>
> Do not use RMAN configuration to change any of the pre-tuned cloud RMAN settings.

Option 2: User Configured Backups

Customers can also configure backups from the host using the `dbaascli database backup` and `dbaascli database recover` commands. These backups, however, are not synchronized with the Control Plane nor are they integrated with the OCI APIs. Also, neither management nor lifecycle operations on these backups are supported from the service Control Plane console. Hence, this is not a recommended approach.

This approach is useful when direct access to Backup destinations is required to perform certain tasks. Accessing the OSS bucket, for example, to replicate backups across regions or monitor Backup Destinations.

If customers configure backups to Object Storage using RMAN without using the OCI Control Plane or OCI APIs, customers are responsible for manually configuring TDE Wallet backups. By default, Oracle cloud automation cleans up archive log files every 24 hours. When you use RMAN to perform manual backups, there is a risk of the archive logs being deleted. Refer to dbaascli database backup for information on how to configure the archive log cleanup. The recommendation is to use Oracle managed backups.

For more information, see *User Configured Backup*.

Option 3: Backups using RMAN

Backups can be directly taken using RMAN with customer-owned customized scripts. Oracle, however, does not recommend this approach.

It is not recommended to use RMAN backups in conjunction with Oracle Managed Backups or User Configured Backups.

Who can use this option:

- Customers who want to maintain their existing RMAN backup/restore scripts.

- Customers who want to configure backups from Standby database in Data Guard environments to offload the backup workload to Standby.

ExaDB-D:

If you plan to backup using RMAN, then you must unregister the database from backup automation. For more information, see *Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management*.

Related Topics

- dbaascli database backup
- dbaascli database recover
- Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management

Managing Exadata Database Backups

Automatic Exadata database backups are managed by Oracle Cloud Infrastructure. You configure this by using the Console or the API.

For unmanaged backups, see *Managing Exadata Database Backups by Using bkup_api*.

There are two destinations possible for automatic Exadata database backups: Autonomous Recovery Service, or Oracle Object Storage.

> **Note:**
>
> If you previously used `bkup_api` to configure backups and then you switch to using the Console or the API for backups:
>
> - A new backup configuration is created and associated with your database. This means that you can no longer rely on your previously configured unmanaged backups to protect your database.
> - `bkup_api` uses cron jobs to schedule backups. These jobs are not automatically removed when you switch to using managed backups.

Related Topics

- Managing Exadata Database Backups by Using bkup_api

Managed Backup Types and Usage Information

There are two types of automatic Exadata database backups: Autonomous Recovery Service, and Oracle Object Storage.

The database and infrastructure (the VM cluster or DB system) must be in an “Available” state for a backup operation to run successfully. Oracle recommends that you avoid performing actions that could interfere with availability (such as patching operations) while a backup operation is in progress. If an automatic backup operation fails, the Database service retries the operation during the next day’s backup window. If an on-demand full backup fails, you can try the operation again when the Exadata Cloud Infrastructure instance and database availability are restored.

When you enable the Automatic Backup feature, either service creates daily incremental backups of the database to the selected Backup Destination.

If you choose to enable automatic backups, you can control the retention period. The system automatically deletes backups when the assigned retention period is expired.

Object Storage Backup retention period: 7, 15, 30, 45, 60. Default: 30 days.

The automatic backup process starts at any time during your daily backup window. You can optionally specify a 2-hour scheduling window for your database during which the automatic backup process will begin. There are 12 scheduling windows to choose from, each starting on an even-numbered hour (for example, one window runs from 4:00-6:00 AM, and the next from 6:00-8:00 AM). Backups jobs do not necessarily complete within the scheduling window.

The default backup window of 00:00 to 06:00 in the time zone of the Exadata Cloud Infrastructure instance's region is assigned to your database if you do not specify a window. Note that the default backup scheduling window is six hours long, while the windows you specify are two hours long.

Autonomous Recovery Service protection policy:

- **Bronze** :14 days
- **Silver**: 35 days
- **Gold**: 65 days
- **Platinum**: 95 days
- Custom defined by you
- **Default**: Silver - 35 days

The automatic backup process starts at any time or within the assigned window.

> **Note:**
>
> - **Data Guard:** You can enable the Automatic Backup feature on a database with the standby role in a Data Guard association.
> - **Backup Retention Changes:** If you shorten your database's backup retention period or your protection policy in the future, existing backups falling outside the updated retention period are deleted by the system.
> - **Backup Storage Costs:** Automatic backups incur storage usage costs for either Autonomous Recovery Service or Object Storage depending on the backup destination selected.

You can create a full backup of your database at any time using either service.

When you terminate an Exadata Cloud Service instance database, all of its resources are deleted. Managed backups using the Object Storage destination will be deleted, and Managed backups using the Autonomous Recovery Service will be deleted according to the deletion option selected. Standalone backups created in Object Storage will remain after the database is terminated and must be manually deleted. You can use a standalone backup to create a new database.

To align with the Oracle recommended practice of using SYSBACKUP administrative privilege for Backup and Recovery operations, cloud automation creates a common administrative user C##DBLCMUSER with SYSBACKUP role at the CDB$ROOT container level. Backup and Recovery operations are therefore performed with the user having the least required privileges. Credentials for this user are randomly generated and securely managed by cloud automation.

If the user is not found or is LOCKED and EXPIRED, then cloud automation will recreate or unlock this user during the backup or recovery operation. This change in the cloud automation is made starting with *dbaastools version 21.4.1.1.0*.

Related Topics

- Release 21.4.1.1.0 (220209)
- To terminate a database

Default Backup Channel Allocation

The default settings for database backup channels when using "Oracle Managed Backup" or "User Configured Backup"

When a database is configured for backup using "Oracle Managed Backup" or "User Configured Backup", the tooling uses "default" for the backup channels. When default is used, dbaas will determine the number of channels to allocate at the time the backup or restore command is executed. The number of channels allocated is determined by the OCPU count of the node. The following table provides the values used and the OCPU range, both the OCPU and the channel values are per node. Restore operations are prioritized. The cluster-wide total channel count is the per node value multiplied by the number of nodes. The automation uses the SCAN to distribute RMAN channels across all nodes in the cluster.

OCPUs Per Node	Formula	Backup Channels Allocation Per Node	Restore Channels Allocation Per Node
Less than or equal to 12	OCPU <= 12	2	4
Greater than 12 and less than or equal to 24	OCPU > 12 and OCPU <= 24	4	8
Greater than 24	OCPU > 24	8	16

If needed, a static per node value can be set by using the DBAASCLI getConfig/configure to generate a bckup cfg and setting the parameter `bkup_channels_node` to the number of channels per node desired.

Valid values are 1 - 32: The total channel count will be the value times the number of nodes. This value cannot exceed the limit of 255 channels. A value of `default` for `bkup_channels_node` sets OCPU channel based allocation.

Prerequisites for Backups on Exadata Cloud Infrastructure

Recovery Service

Ensure that your tenancy is configured to use Recovery Service.

Table 5-5 Review the prerequisite tasks before you use Recovery Service as the automatic backup destination

Task	More Information	Required or Optional
Create IAM policies	Policies to Enable Access to Recovery Service and Related Resources	Required
Configure network resources and register a Recovery Service subnet	Creating a Recovery Service Subnet in the Database VCN	Required

Table 5-5 (Cont.) Review the prerequisite tasks before you use Recovery Service as the automatic backup destination

Task	More Information	Required or Optional
Create protection policies	Review Protection Policies for Database Backup Retention	Optional

For more information about Recovery Service, see Overview of Oracle Database Autonomous Recovery Service.

Object Storage

- The Exadata Cloud Service instance requires access to the Oracle Cloud Infrastructure Object Storage. Oracle recommends using a service gateway with the VCN to enable this access. For more information, see Network Setup for Exadata Cloud Infrastructure Instances. In that topic, pay particular attention to:
  - Service Gateway for the VCN
  - Node Access to Object Storage: Static Route
  - *Backup egress rule: Allows access to Object Storage*
  - Subnet Size Requirements and Security Rules for Recovery Service Subnet
- An existing Object Storage bucket to use as the backup destination. You can use the Console or the Object Storage API to create the bucket. For more information, see Managing Buckets.
- An auth token generated by Oracle Cloud Infrastructure. You can use the Console or the IAM API to generate the password. For more information, see Working with Auth Tokens.
- The user name specified in the backup configuration file must have tenancy-level access to Object Storage. An easy way to do this is to add the user name to the Administrators group. However, that allows access to all of the cloud services. Instead, an administrator should create a policy like the following that limits access to only the required resources in Object Storage for backing up and restoring the database:

  ```
  Allow group <group_name> to manage objects in compartment
  <compartment_name> where target.bucket.name = '<bucket_name>'
  Allow group <group_name> to read buckets in compartment <compartment_name>
  ```

  For more information about adding a user to a group, see Managing Groups. For more information about policies, see Getting Started with Policies.

Related Topics

- Auth Token

Using the Console to Manage Backups

You can use the Console to enable automatic incremental backups, create full backups on demand, and view the list of managed backups for a database. You can also use the Console to delete manual (on-demand) backups.

> **Note:**
>
> - The list of backups you see in the Console does not include any unmanaged backups (backups created directly by using `bkup_api`).
> - All backups are encrypted with the same master key used for Transparent Data Encryption (TDE) wallet encryption.
> - Backups for a particular database are listed on the details page for that database. The **Encryption Key** column displays either **Oracle-Managed Key** or a key name if you are using your own encryption keys to protect the database. See Backing Up Vaults and Keys for more information.

> **Note:**
>
> Do not delete any necessary encryption keys from the vault because this causes databases and backups protected by the key to become unavailable.

To configure automatic backups for a database

When you create an Exadata Cloud Infrastructure instance, you can optionally enable automatic backups for the initial database. Use this procedure to enable or disable automatic backups after the database is created.

> **Note:**
>
> Databases in a *security zone compartment* must have automatic backups enabled. See the *Security Zone Policies* topic for a full list of policies that affect Database service resources.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata on Oracle Public Cloud**.
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system containing the database you want to configure:
   Cloud VM clusters (The New Exadata Cloud Infrastructure Resource Model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

DB systems: Under Oracle Base Database, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

4. In the list of databases, find the database for which you want to enable or disable automatic backups, and click its name to display database details. The details indicate whether automatic backups are enabled.
5. Click **Configure Automatic Backups**.
6. In the Configure Automatic Backups dialog, enter the following details:
   - **Backup Destination**: Your choices are **Autonomous Recovery Service** (default) or **Object Storage**.
     - **Scenario 1:** The customer enables automatic backups AND has available limits AND there is available capacity in the region for Autonomous Recovery Service.

       Backup Destination: Your choices are Autonomous Recovery Service (default) or Object Storage. You can switch the backup destination from Autonomous Recovery Service to Object Storage.
     - **Scenario 2:** Customer enables automatic backups AND has exhausted the default limits for the Recovery Service AND there is available capacity in the region for Autonomous Recovery Service.

       Backup Destination: You can only use Object Storage. However, you can make an additional limits request and then use Autonomous Recovery Service.

       The system displays the following message with a link to request an increase to the limits.

       ```
       Tenancy has reached the limit for Autonomous Recovery Service. View your service limits and request an update.
       ```
     - **Scenario 3:** Customer enables automatic backups AND there is no available capacity in the region for Autonomous Recovery Service.

       Backup Destination: You can only use Object Storage. You can transition to Autonomous Recovery Service when there is sufficient capacity.

       The system displays the following message

       ```
       Autonomous Recovery Service has no available capacity in this region. Select Object Storage as your backup destination. You can transition from Object Storage to Autonomous Recovery Service when there is sufficient capacity.
       ```

       Proactively check if Autonomous Recovery Service capacity is available. If the required capacity becomes available and if you had chosen Object Storage, then you can transition to Autonomous Recovery Service.
   - **Backup Scheduling**:
     - **Object Storage (L0)**:
       - **Full backup scheduling day**: Choose a day of the week for the initial and future L0 backups to start.
       - **Full backup scheduling time (UTC)**: Specify the time window when the full backups start when the automatic backup capability is selected.
       - **Take the first backup immediately**: A full backup is an operating system backup of all datafiles and the control file that constitute an Oracle Database. A full backup should also include the parameter file(s) associated with the database. You can take a full database backup when the database is shut

down or while the database is open. You should not normally take a full backup after an instance failure or other unusual circumstances.

If you choose to defer the first full backup your database may not be recoverable in the event of a database failure.

- **Object Storage (L1)**:
  - **Incremental backup scheduling time (UTC)**: Specify the time window when the incremental backups start when the automatic backup capability is selected.
- **Autonomous Recovery Service (L0)**:
  - **Scheduled day for initial backup**: Choose a day of the week for the initial backup.
  - **Scheduled time for initial backup (UTC)**: Select the time window for the initial backup.
  - **Take the first backup immediately**: A full backup is an operating system backup of all datafiles and the control file that constitute an Oracle Database. A full backup should also include the parameter file(s) associated with the database. You can take a full database backup when the database is shut down or while the database is open. You should not normally take a full backup after an instance failure or other unusual circumstances.
    If you choose to defer the first full backup your database may not be recoverable in the event of a database failure.
- **Autonomous Recovery Service (L1)**:
  - **Scheduled time for daily backup (UTC)**: Specify the time window when the incremental backups start when the automatic backup capability is selected.
- **Deletion options after database termination**: Options that you can use to retain protected database backups after the database is terminated. These options can also help restore the database from backups in case of accidental or malicious damage to the database.
  - **Retain backups for the period specified in your protection policy or backup retention period**: Select this option if you want to retain database backups for the entire period defined in the Object Storage Backup retention period or Autonomous Recovery Service protection policy after the database is terminated.
  - **Retain backups for 72 hours, then delete**: Select this option to retain backups for a period of 72 hours after you terminate the database.

- **Enable Real-Time Data Protection**: Real-time protection is the continuous transfer of redo changes from a protected database to **Autonomous Recovery Service**. This reduces data loss and provides a recovery point objective (RPO) near 0. This is an extra cost option.

7. Click **Save Changes**.
   The Database Details page displays the configuration details, **Health**, **Real-Time Data Protection**, and **Policy information** in the **Backup** section.

Related Topics

- security zone compartment
- Security Zone Policies

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

To create an on-demand backup of a database

> **Note:**
>
> Object Storage creates a full backup of the database while Recovery Service creates an incremental backup.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system containing the database you want to back up:

   Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems: Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. In the list of databases, find the database for which you want to create an on-demand full backup and click its name to display database details.
5. Under **Resources**, click **Backups**.

   A list of backups is displayed.
6. Click **Create Backup**.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

To view backup status

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster containing the database backup you want to view.
4. Click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
5. In the list of databases, find the database you are interested in and click its name to display database details.
6. Under **Resources**, click **Backups**.
   A list of backups is displayed. The state column displays the status of the backup: **Active**, **Creating**, **Canceled**, **Canceling**, or **Failed**.

To cancel a backup

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata on Oracle Public Cloud**.
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster containing the database backup you want to view:
4. Click **Exadata VM Clusters**.
   In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
5. In the list of databases, find the database you are interested in and click its name to display database details.
6. Under **Resources**, click **Backups**.
   A list of backups is displayed. The state column displays the status of the backup: **Active**, **Creating**, **Canceled**, **Canceling**, or **Failed**.
7. A backup in the Creating state may be canceled by clicking the Actions icon (three dots) on the right of the backup row and clicking **Cancel Backup**.
   A Cancel Backup confirmation dialog will appear.
8. Enter the name of the backup, and click **Cancel Backup**.
   The state changes to **Canceling**.

   The Cancel backup Work request can be viewed, by clicking **Work requests** under **Resources**.

If the Cancel backup fails:

- In the Work requests pane under Resources, you will see a line item called "**Cancel Database Backup**" with a state of "**Failed**". There will also be a work request for the backup "**Create Database Backup**" that will reflect the state of the Backup operation.

To delete full backups from Object Storage

> **Note:**
>
> You cannot explicitly delete automatic backups. Unless you terminate the database, automatic backups remain in Recovery Service and Object Storage for the number of days specified by the user, after which time they are automatically deleted.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system containing the database backup you want to delete:

   Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

DB systems: Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

4. In the list of databases, find the database you are interested in and click its name to display database details.
5. Under **Resources**, click **Backups**.

   A list of backups is displayed.
6. Click the Actions icon (three dots) for the backup you are interested in, and then click **Delete**.
7. Confirm when prompted.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

To delete standalone backups from Object Storage

1. Open the navigation menu. Click **Oracle Database**, then click **Standalone Backups** under **Resources**.
2. In the list of standalone backups, find the backup you want to use to delete.
3. Click the Actions menu for the backup you are interested in, and then click **Delete**.
4. In the **Delete** dialog, click **Delete** to confirm the backup deletion.

To designate Autonomous Recovery Service as a Backup Destination for an Existing Database

To designate Autonomous Recovery Service as a Backup Destination for an existing database, use this procedure.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata on Oracle Public Cloud**.
2. Choose your **Compartment**.
3. Navigate to the database:
   Cloud VM clusters (The New Exadata Cloud Infrastructure Resource Model): Under **Exadata on Oracle Public Cloud**, click **Exadata VM Clusters**.

   In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems: Under Oracle Base Database, click **DB Systems**.

   In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

   On the cloud **VM cluster** or **DB system** details page, in the Databases table, click the name of the database to display the **Database Details** page.
4. Click **Configure automatic backups**.
5. In the resulting window, provide the following details:

- **Enable automatic backup**: Check the check box to enable automatic incremental backups for this database. If you are creating a database in a security zone compartment, you must enable automatic backups.
- **Backup Destination**: Select **Autonomous Recovery Service**.
- **Backup Scheduling**: If you enable automatic backups, you can choose a two-hour scheduling window to control when backup operations begin. If you do not specify a window, then a six-hour default window of 00:00 to 06:00 (in the time zone of the DB system's region) is used for your database.
- **Protection Policy**: If you choose to enable automatic backups, you can choose a policy with one of the following preset retention periods, or a Custom policy.

  Object Storage Backup retention period: 7, 15, 30, 45, 60. Default: 30. The system automatically deletes your incremental backups at the end of your chosen retention period.

  Autonomous Recovery Service protection policy:

  - **Bronze:** 14 days
  - **Silver:** 35 days
  - **Gold:** 65 days
  - **Platinum:** 95 days
  - Custom defined by you
  - **Default:** Silver - 35 days
- **Enable Real-Time Data Protection**: Real-time protection is the continuous transfer of redo changes from a protected database to **Autonomous Recovery Service**. This reduces data loss and provides a recovery point objective (RPO) near 0. This is an extra cost option.

6. Click **Save Changes**.

Recovering an Exadata Database from Backup Destination

This topic explains how to recover an Exadata database from a backup stored in either Object Storage or Autonomous Recovery Service by using the Console or the API.

- Object Storage service is a secure, scalable, on-demand storage solution in Exadata Cloud Infrastructure.
- OracleDatabase Autonomous Recovery Service is a centralized, fully managed, and standalone backup solution for Oracle Cloud Infrastructure (OCI) databases.

For more information about backing up your databases to Object Storage, see *Managing Exadata Database Backups*.

- Using the Console to restore a database
  You can use the Console to restore the database from a backup in a backup destination that was created by using the Console.

Related Topics

- Managing Exadata Database Backups
  Automatic Exadata database backups are managed by Oracle Cloud Infrastructure. You configure this by using the Console or the API.

Using the Console to restore a database

You can use the Console to restore the database from a backup in a backup destination that was created by using the Console.

You can restore to:

- Restore to latest
- Restore to a timestamp
- Restore to SCN

> **Note:**
>
> The list of backups you see in the Console does not include any unmanaged backups (backups created directly by using `bkup_api`).

- To restore a database

To restore a database

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system containing the database you want to restore:
   Cloud VM clusters (The New Exadata Cloud Infrastructure Resource Model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems: Under **Oracle Base Database**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. In the list of databases, find the database you want to restore, and click its name to display details about it.
5. Click **Restore**.
6. Select one of the following options, and click **Restore Database**:
   - **Restore to the latest**: Restores the database to the last known good state with the least possible data loss.
   - **Restore to the timestamp**: Restores the database to the timestamp specified.
   - **Restore to System Change Number (SCN)**: Restores the database using the SCN specified. This SCN must be valid.

     > **Note:**
     >
     > You can determine the SCN number to use either by accessing and querying your database host, or by accessing any online or archived logs.

7. Confirm when prompted.
If the restore operation fails, the database will be in a "**Restore Failed**" state. You can try restoring again using a different restore option. However, Oracle recommends that you review the `RMAN` logs on the host and fix any issues before reattempting to restore the database. These log files can be found in subdirectories of the `/var/opt/oracle/log` directory.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

Managing Exadata Database Backups by Using bkup_api

You can use Exadata's backup utility, `bkup_api`, to back up databases on an Exadata Cloud Infrastructure instance to an existing bucket in the Oracle Object Storage service and to the local disk Fast Recovery Area.

> **Note:**
>
> `bkup_api` is deprecated equivalent `dbaascli` commands are listed for each `bkup_api` command

For backups managed by Oracle Cloud Infrastructure, see Managing Exadata Database Backups.

This topic explains how to:

- Create a backup configuration file that indicates the backup destination, when the backup should run, and how long backups are retained. If the backup destination is Object Storage, the file also contains the credentials to access the service.
- Associate the backup configuration file with a database. The database will be backed up as scheduled, or you can create an on-demand backup.

> **Note:**
>
> You must update the cloud-specific tooling on all the compute nodes in your Exadata Cloud Infrastructure instance before performing the following procedures. For more information, see Updating an Exadata Cloud Service Instance.

- Default Backup Configuration
Description of the default backup configuration that follows Oracle best practice guidelines.
- To create a backup configuration file
- To create an on-demand backup
- To remove the backup configuration
- To delete a local backup
- To delete a backup in Object Storage

Default Backup Configuration

Description of the default backup configuration that follows Oracle best practice guidelines.

The backup configuration follows a set of Oracle best-practice guidelines:

- Full (level 0) backup of the database followed by rolling incremental (level 1) backups on a seven-day cycle (a 30-day cycle for the Object Storage destination).
- Full backup of selected system files.
- Automatic backups daily at a specific time set during the database deployment creation process.

Retention period:

- Both Object Storage and local storage: 30 days, with the 7 most recent days' backups available on local storage.
- Object Storage only: 30 days.
- Local storage only: Seven days.

Encryption:

- Both Object Storage and local storage: All backups to cloud storage are encrypted.
- Object Storage only: All backups to cloud storage are encrypted.

To create a backup configuration file

> **Note:**
>
> `bkup_api` is deprecated, use `dbaascli database backup` and its options instead

> **Note:**
>
> The following procedure must be performed on the first compute node in the Exadata Cloud Infrastructure VM cluster or DB system resource. To determine the first compute node, connect to any compute node as the `grid` user and execute the following command:
>
> ```
> $ $ORACLE_HOME/bin/olsnodes -n
> ```
>
> The first node has the number 1 listed beside the node name.

1. SSH to the first compute node in the VM cluster or DB system resource.

   ```
   ssh -i <private_key_path> opc@<node_1_ip_address>
   ```

2. Log in as opc and then sudo to the root user.

   ```
   login as: opc
   ```

```
[opc@dbsys ~]$ sudo su -
```

3. Use the `bkup_api get config` command to generate a file containing the current backup settings for the database deployment:

```
# /var/opt/oracle/bkup_api/bkup_api get config [--file=<file_name>] --
dbname=<db_name>
```

4. Use the following command to install the backup configuration, configure the credentials, schedule the backup, and associate the configuration with a database name.

```
[root@dbsys bkup]# /var/opt/oracle/ocde/assistants/bkup/bkup -cfg bkup.cfg
--dbname=<database_name>
```

   The backup is scheduled via cron and can be viewed at `/etc/crontab`.

When the scheduled backup runs, you can check its progress with the following command.

```
[root@dbsys bkup]# /var/opt/oracle/bkup_api/bkup_api bkup_statusThe backup
configuration file parameters are described in the following table:
```

Parameter	Description
bkup_disk=[yes|no]	Whether to back up locally to disk (Fast Recovery Area).
bkup_oss=[yes|no]	Whether to back up to Object Storage. If yes, you must also provide the parameters bkup_oss_url, bkup_oss_user, bkup_oss_passwd, and bkup_oss_recovery_window.
bkup_oss_url=<swift_url>	Required if bkup_oss=yes. The Object Storage URL including the tenant and bucket you want to use. The URL is: `https://swiftobjectstorage.<region_name>.oraclecloud.com/v1/<tenant>/<bucket>` where <tenant> is the *lowercase* tenant name (even if it contains uppercase characters) that you specify when signing in to the Console and <bucket> is the name of the existing bucket you want to use for backups.

Parameter	Description
bkup_oss_user=<oci_user_name>	Required if bkup_oss=yes. The user name for the Oracle Cloud Infrastructure user account. This is the user name you use to sign in to the Oracle Cloud Infrastructure Console. For example, `jsmith@example.com` for a local user or `<identity_provider>/jsmith@example.com` for a federated user. To determine which type of user you have, see the following topics: • Managing Users (for information on local users) • Federating with Identity Providers (for information on federated users) Note that the user must be a member of the Administrators group, as described in Prerequisites.
bkup_oss_passwd=<auth_token>	Required if bkup_oss=yes. The auth token generated by using the Console or IAM API, as described in Prerequisites. This is **not** the password for the Oracle Cloud Infrastructure user.
bkup_oss_recovery_window=n	Required if bkup_oss=yes. The number of days for which backups and archived redo logs are maintained in the Object Storage bucket. Specify 7 to 90 days.
bkup_daily_time=hh:mm	The time at which the daily backup is scheduled, specified in hours and minutes (hh:mm), in 24-hour format.
bkup_archlog_cron_entry=[yes|no]	When no backups are configured using dbaastools, setting `bkup_archlog_cron_entry=no` will remove the archive log cleanup job from crontab. The default value is “yes”.

Related Topics

- dbaascli database backup
  To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

To create an on-demand backup

> **Note:**
>
> `bkup_api` is deprecated. Use `dbaascli database backup` and its options instead.

You can use the `bkup_api` utility to create an on-demand backup of a database.

1. SSH to the first compute node in the Exadata VM cluster or DB system resource.

   ```
   ssh -i <private_key_path> opc@<node_1_ip_address>
   ```

To determine the first compute node, connect to any compute node as the `grid` user and execute the following command:

```
$ $ORACLE_HOME/bin/olsnodes -n
```

The first node has the number 1 listed beside the node name.

2. Log in as opc and then sudo to the root user.

```
login as: opc

[opc@dbsys ~]$ sudo su -
```

3. You can let the backup follow the current retention policy, or you can create a long-term backup that persists until you delete it:
   - To create a backup that follows the current retention policy, enter the following command:

     ```
     # /var/opt/oracle/bkup_api/bkup_api bkup_start --dbname=<database_name>
     ```

   - To create a long-term backup, enter the following command:

     ```
     # /var/opt/oracle/bkup_api/bkup_api bkup_start --keep --
     dbname=<database_name>
     ```

4. Exit the root-user command shell and disconnect from the compute node:

```
# exit
$ exit
```

By default, the backup is given a timestamp-based tag. To specify a custom backup tag, add the `--tag` option to the `bkup_api` command; for example, to create a long-term backup with the tag "monthly", enter the following command:

```
# /var/opt/oracle/bkup_api/bkup_api bkup_start --keep --tag=monthly
```

After you enter a `bkup_api bkup_start` command, the `bkup_api` utility starts the backup process, which runs in the background. To check the progress of the backup process, enter the following command:

```
# /var/opt/oracle/bkup_api/bkup_api bkup_status --dbname=<database_name>
```

Related Topics

- dbaascli database backup
  To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

To remove the backup configuration

> **Note:**
>
> `bkup_api` is deprecated. Use `dbaascli database backup` with its option instead.

A backup configuration can contain the credentials to access the Object Storage bucket. For this reason, you might want to remove the file after successfully configuring the backup.

```
[root@dbsys bkup]# rm bkup.cfg
```

Related Topics

- dbaascli database backup
  To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

To delete a local backup

> **Note:**
>
> `bkup_api` is deprecated. Use `dbaascli database backup` and its option instead.

To delete a backup of a database deployment on the Exadata Cloud Infrastructure instance, use the `bkup_api` utility.

1. Connect to the first compute node in your Exadata VM cluster or DB system resource as the `opc` user.

   To determine the first compute node, connect to any compute node as the `grid` user and execute the following command:

   ```
   $ $ORACLE_HOME/bin/olsnodes -n
   ```

   The first node has the number 1 listed beside the node name.
2. Start a root-user command shell:

   ```
   $ sudo -s#
   ```

3. List the available backups:

   ```
   # >/var/opt/oracle/bkup_api/bkup_api recover_list --dbname=<database_name>
   ```

   where `dbname` is the database name for the database that you want to act on.

   A list of available backups is displayed.

4. Delete the backup you want:

```
# /var/opt/oracle/bkup_api/bkup_api bkup_delete --bkup=<backup-tag> --
dbname=<database_name>
```

where `backup-tag` is the tag of the backup you want to delete.

5. Exit the root-user command shell:

```
# exit
$
```

Related Topics

- dbaascli database backup
  To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

To delete a backup in Object Storage

Use the RMAN `delete backup` command to delete a backup from the Object Store.

Using the API to Manage Backup and Recovery

- Using the API to manage backups

Using the API to manage backups

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.

Use these API operations to manage database backups:

- ListBackups
- GetBackup
- CreateBackup
- DeleteBackup
- UpdateDatabase - To enable and disable automatic backups.
- RestoreDatabase

For the complete list of APIs for the Database service, see Database Service API.

Alternative Backup Methods

Learn about alternative backup methods that are available in addition to the OCI Console.

Backup for databases on Exadata Cloud Infrastructure can be accomplished through several methods in addition to the automatic backups configured in the console. Generally, the console (or the OCI API / CLI that correspond to it) is the preferred method as it provides the simplest and most automated method. In general, it is preferable to leverage the OCI Console, OCI API, or OCI Command-Line over alternative management methods. However, if required actions

cannot be completed through the preferred methods, two other options are available to manually configure backups: `bkup_api` and Oracle Recovery Manager (RMAN).

> **Note:**
>
> `bkup_api` will be deprecated in a future release. Use the `dbaascli database backup`, `dbaascli pdb backup`, `dbaascli database recover`, and `dbaascli pdb recover` commands to backup and recover container databases and pluggable databases. For more information, see *User Configured Backup*.

RMAN is the backup tool included with the Oracle Database. For information about using RMAN, see the *Oracle Database Backup and Recovery User's Guide for Release 19*. Using RMAN to back up databases on Exadata Cloud Infrastructure provides the most flexibility in terms of backup options, but also the most complexity.

> **Note:**
>
> While using RMAN for restoring databases backed up through any method described herein is considered safe, RMAN should NEVER be used to set up backups in conjunction with either console (and OCI API / CLI), nor in conjunction with `bkup_api`. If you choose to orchestrate backups manually leveraging RMAN, you should not use either console automated backups, nor should you use `bkup_api`. You must first completely disable console based automated backups. For more information, see *Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management*.

The `bkup_api` method offers a middle ground between RMAN and console automated backups in terms of flexibility and simplicity. Use `bkup_api` if needed functionality is not supported with console automated backups, but when you wish to avoid complexity of using RMAN directly. In certain cases, `bkup_api` can be used to modify the console automated backup configuration, but this is not generally the case. Generally, `bkup_api` must be used instead of enabling backups in the console.

- Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management

Related Topics

- Managing Exadata Database Backups by Using bkup_api
- Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management

Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management

Backups, configured in the Exadata Cloud Infrastructure console, API or `bkup_api` work for a variety of backup and recovery use cases. If you require use cases not supported by the cloud-managed backups, then you can manage database backup and recovery manually, using the Oracle Recovery Manager (RMAN) utility. For information about using RMAN, see the *Oracle Database Backup and Recovery User's Guide for Release 19*.

Managing backup and recovery, using RMAN, on Exadata Cloud Infrastructure requires taking full ownership of both database and archive log backups, and the cloud-managed backups

should no longer be used. Before manual backups are started, the cloud-managed backup functionality should be disabled. This is needed so the cloud backup jobs do not purge archive logs before they are manually backed up and do not conflict with the manual backups.

You can use the `bkup_api` utility to disable cloud-managed backups, including disabling the automatic archive log purge job, by following this procedure:

> **Note:**
>
> If you execute these steps, then the automation will no longer purge/backup the archive logs in the FRA for the database.

1. Connect as the `opc` user to the first compute node.

   For detailed instructions, see *Connecting to a Compute Node with SSH*.

2. Start a root-user command shell:

   ```
   sudo -s
   ```

3. Use the `bkup_api` get config command to generate a file containing the current backup settings for the database deployment:

   ```
   /var/opt/oracle/bkup_api/bkup_api get config [--file=filename] --
   dbname=dbname
   ```

   Where:

   - *`filename`* is an optional parameter used to specify a name for the file that is generated
   - *`dbname`* is the database name for the database that you want to act on

4. Edit the parameter values in the generated file to change the following parameters.

   This will remove the backup crontab entries and disable all automatic backups. If the values are set to `yes`, then set to `no`.

   ```
   bkup_cron_entry=no
   bkup_archlog_cron_entry=no
   bkup_nfs=no
   bkup_oss=no
   bkup_local=no
   ```

5. Use the `bkup_api set config` command to update the backup settings using the file containing your updated backup settings:

   ```
   /var/opt/oracle/bkup_api/bkup_api set config --file=filename --
   dbname=dbname
   ```

   Where:

   - *`filename`* is an optional parameter used to specify a name for the file that is generated
   - *`dbname`* is the database name for the database that you want to act on

   The job to set the configuration will take several minutes to complete.

6. You can use the `bkup_api configure_status` command to check the status of the configuration update:

   ```
   /var/opt/oracle/bkup_api/bkup_api configure_status --dbname=dbname
   ```

   Where:

   - *dbname* is the database name for the database that you want to act on

   The **Configure backup status** starts as **running** and then moves to **finished** when complete.

7. Run the `bkup_api get config` command again and verify the settings listed above are set to `no`.

   ```
   /var/opt/oracle/bkup_api/bkup_api get config [--file=filename] --
   dbname=dbname
   ```

   Where:

   - *filename* is an optional parameter used to specify a name for the file that is generated
   - *dbname* is the database name for the database that you want to act on

   > **Note:**
   >
   > After making these changes, no backups, including archive log backups, are made by the cloud automation. Ensure that manual RMAN backups are in place to avoid filling the archive log location.

   > **Note:**
   >
   > Changes made using the `bkup_api` command are not reflected in the Oracle Exadata Database Service on Dedicated Infrastructure console.

8. Exit the root-user command shell:

   ```
   exit
   ```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- Oracle Database Backup and Recovery User's Guide for Release 19

Recovering a Database Using Oracle Recovery Manager (RMAN)

If you backed up your database using `bkup_api`, then you can manually restore that database backup by using the Oracle Recovery Manager (RMAN) utility. For information about using RMAN, see the *Oracle Database Backup and Recovery User's Guide for Release 19*.

> **Note:**
>
> While recovering using RMAN is safe, you must not use RMAN to initiate backups or edit backup setting in conjunction with either `backup_api` usage or in conjunction with automated console backups. Doing so could result in conflicting conditions or over-writes of settings, and backups may not execute successfully.

Related Topics

- Oracle Database Backup and Recovery User's Guide for Release 19

Patch and Update an Exadata Cloud Infrastructure System

- User-Managed Maintenance Updates
  Maintaining a secure Exadata Cloud Infrastructure instance in the best working order requires you to perform the following tasks regularly:
- Patching and Updating an Exadata Cloud Infrastructure System
  Learn how to perform patching operations on Exadata database virtual machines and Database Homes by using the Console, API, or the CLI.
- Patching and Updating an Exadata Cloud Infrastructure System Manually
  This topic describes the procedures for patching and updating various components in Exadata Cloud Service outside of the cloud automation.

User-Managed Maintenance Updates

Maintaining a secure Exadata Cloud Infrastructure instance in the best working order requires you to perform the following tasks regularly:

- Patching the Oracle Grid Infrastructure and Oracle Database software on the VM Cluster virtual machines. For information and instructions, see *Patching and Updating VM Cluster's GI and Database Homes*.
- Updating the operating system on the VM Cluster virtual machines. See *Updating an Exadata Cloud VM Cluster Operating System* for information and instructions.

Related Topics

- Patching and Updating VM Cluster's GI and Database Homes
  This topic explains how to perform patching operations on Exadata Cloud Infrastructure resources by using the Console, API, or the CLI.
- Oracle Clusterware Configuration and Administration
- Updating an Exadata Cloud VM Cluster Operating System
  Exadata VM cluster image updates allow you to update the OS image on your Exadata cloud VM cluster nodes in an automated manner from the OCI console and APIs.
- Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli
  Learn to use the `dbaascli` utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata Cloud Infrastructure system.

Patching and Updating an Exadata Cloud Infrastructure System

Learn how to perform patching operations on Exadata database virtual machines and Database Homes by using the Console, API, or the CLI.

For information and instructions on patching the system by using the `dbaascli` utility, see *Patching and Updating an Exadata Cloud Infrastructure System Manually*.

For more information and examples for applying database quarterly patches on Exadata Cloud Infrastructure refer to My Oracle Support note: *How to Apply Database Quarterly Patch on Exadata Cloud Service and Exadata Cloud at Customer Gen 2 (Doc ID 2701789.1)*.

For more guidance on achieving continuous service during patching operations, see the *Application Checklist for Continuous Service for MAA Solutions* white paper.

- Patching and Updating VM Cluster's GI and Database Homes
  This topic explains how to perform patching operations on Exadata Cloud Infrastructure resources by using the Console, API, or the CLI.
- Updating an Exadata Cloud VM Cluster Operating System
  Exadata VM cluster image updates allow you to update the OS image on your Exadata cloud VM cluster nodes in an automated manner from the OCI console and APIs.
- Upgrading Exadata Grid Infrastructure
  This topic describes how to upgrade the Oracle Grid Infrastructure (GI) on an Exadata cloud VM cluster using the Oracle Cloud Infrastructure Console or API.
- Upgrading Exadata Databases
  This topic describes the procedures to upgrade an Exadata database instance to Oracle Database 19c (Long Term Release) by using the Console and the API. The upgrade is accomplished by moving the Exadata database to a Database Home that uses the target software version.

Related Topics

- Patching and Updating an Exadata Cloud Infrastructure System Manually
  This topic describes the procedures for patching and updating various components in Exadata Cloud Service outside of the cloud automation.
- https://support.oracle.com/epmos/faces/DocContentDisplay?id=2701789.1
- Application Checklist for Continuous Service for MAA Solutions

Patching and Updating VM Cluster's GI and Database Homes

This topic explains how to perform patching operations on Exadata Cloud Infrastructure resources by using the Console, API, or the CLI.

> **Note:**
>
> Oracle recommends patching databases by moving them to a Database Home that uses the target patching level. See To patch a database by moving it to another Database Home for instructions on this method of database patching.

For information and instructions on patching the system by using the `dbaascli` utility, see Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli.

- About Patching and Updating VM Cluster's GI and Database Homes
  This topic describes the types of patching performed on an Exadata Cloud Infrastructure instances and provides instructions for completing the patching operations.
- Prerequisites for Patching and Updating an Exadata Cloud Infrastructure System
  The Exadata Cloud Infrastructure instance requires access to the Oracle Cloud Infrastructure Object Storage service, including connectivity to the applicable Swift endpoint for Object Storage
- Using the Console to Patch and Update Exadata Cloud Infrastructure Instances
  You can use the Console to view the history of patch operations on Exadata Cloud Infrastructure instances, apply patches, and monitor the status of patch operations.
- Using the API to Patch an Exadata Cloud Infrastructure Instance
  Use these API operations to manage patching the following Exadata resources: cloud VM clusters, DB systems, databases, and Database Homes.

About Patching and Updating VM Cluster's GI and Database Homes

This topic describes the types of patching performed on an Exadata Cloud Infrastructure instances and provides instructions for completing the patching operations.

- Oracle Grid Infrastructure (GI) Patching
  Patching an Exadata Cloud Infrastructure instance updates the components on all the compute nodes in the instance. A VM cluster or DB system patch updates the Oracle Grid Infrastructure (GI) on the resource.
- Database Home Patching
  A Database Home patch updates the Oracle Database software shared by the databases in that home.
- Best Practices for Patching Exadata Cloud Infrastructure Components

Oracle Grid Infrastructure (GI) Patching

Patching an Exadata Cloud Infrastructure instance updates the components on all the compute nodes in the instance. A VM cluster or DB system patch updates the Oracle Grid Infrastructure (GI) on the resource.

> **Note:**
>
> The cloud Exadata resource model the instance is using determines whether you patch the Grid Infrastructure on a DB system resource or a cloud VM cluster resource. VM clusters are used by the The New Exadata Cloud Infrastructure Resource Model DB systems can be easily migrated to the new resource model with no system downtime. Switch an Exadata DB system to the new Exadata resource model .

Database Home Patching

A Database Home patch updates the Oracle Database software shared by the databases in that home.

Thus, you patch a database by either of the following methods:

- Move the database to a Database Home that has the correct patch version. This affects only the database being moved.
- Patching the Database Home the database is currently in. This affects all databases located in the Database Home being patched.

When patching a Database Home, you can use an Oracle-provided database software image to apply a generally-available Oracle Database software update, or you can use a custom database software image created by your organization to apply a specific set of patches required by your database. See Oracle Database Software Images for more information on creating and using custom images.

For instructions on performing patching operations, see To patch the Oracle Database software in a Database Home (cloud VM cluster). For Exadata Cloud Infrastructure instances using the older DB system resource model, see To patch the Oracle Database software in a Database Home (DB system).

Best Practices for Patching Exadata Cloud Infrastructure Components

Consider the following best practices:

- Back up your databases before you apply any patches. For information about backing up the databases, see Managing Exadata Database Backups .
- Patch a VM cluster or an Exadata DB system before you patch the Databases Homes and databases on that resource.
- Before you apply any patch, run the precheck operation to ensure your VM cluster, Exadata DB system, or Database Home meets the requirements for that patch.
- To patch a database to a version other than the database version of the current home, move the database to a Database Home running the target version. This technique requires less downtime and allows you to easily roll back the database to the previous version by moving it back to the old Database Home. See To move a database to another Database Home To patch a database by moving it to another Database Home.
- For the Oracle Database and Oracle Grid Infrastructure major version releases available in Oracle Cloud Infrastructure, patches are provided for the current version plus the three most recent older versions (*N* through *N - 3*). For example, if an instance is using Oracle Database 19c, and the latest version of 19c offered is 19.8.0.0.0, patches are available for versions 19.8.0.0.0, 19.7.0.0, 19.6.0.0 and 19.5.0.0.
- dbaascli database runDatapatch
  To patch an Oracle Database, use the `dbaascli database runDatapatch` command.
- Customer-Managed Keys in Exadata Cloud Infrastructure
  Customer-managed keys for Exadata Cloud Infrastructure is a feature of Oracle Cloud Infrastructure (OCI) Vault service that enables you to encrypt your data using encryption keys that you control.
- dbaascli database addInstance
  To add the database instance on the specified node, use the `dbaascli database addInstance` command.
- dbaascli database convertToPDB
  To convert the specified non-CDB database to PDB, use the `dbaascli database convertToPDB` command.
- dbaascli database getDetails
  This command shows the detailed information of a given database e.g. dbname, node information, pluggable databases information etc.

- dbaascli database modifyParameters
  To modify or reset initialization parameters for an Oracle Database, use the `dbaascli database modifyParameters` command.
- dbaascli database upgrade
  To upgrade an Oracle Database, use the `dbaascli database upgrade` command.

dbaascli database runDatapatch

To patch an Oracle Database, use the `dbaascli database runDatapatch` command.

Prerequisites

- Before performing a `runDatapatch` operation, ensure that all of the database instances associated with the database are up and running.
- Run the command as the `root` user.

Syntax

```
dbaascli database runDatapatch --dbname
[--resume]
    [--sessionID]
[--skipPdbs | --pdbs]
[--executePrereqs]
[--patchList]
[--skipClosedPdbs]
[--rollback]
```

Where:

- `--dbname` specifies the name of the database
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--skipPdbs` skips running the datapatch on a specified comma-delimited list of PDBs. For example: *pdb1*,*pdb2*...
- `--pdbs` runs the datapatch only on a specified comma-delimited list of PDBs. For example: *pdb1*,*pdb2*...
- `--executePrereqs` runs prerequisite checks
- `--patchList` applies or rolls back the specified comma-delimited list of patches. For example: *patch1*,*patch2*...
- `--skipClosedPdbs` skips running the datapatch on closed PDBs
- `--rollback` rolls back the patches applied

```
dbaascli database runDatapatch --dbname db19
```

Customer-Managed Keys in Exadata Cloud Infrastructure

Customer-managed keys for Exadata Cloud Infrastructure is a feature of Oracle Cloud Infrastructure (OCI) Vault service that enables you to encrypt your data using encryption keys that you control.

The OCI Vault service provides you with centralized key management capabilities that are highly available and durable. This key-management solution also offers secure key storage

using isolated partitions (and a lower-cost shared partition option) in FIPS 140-2 Level 3-certified hardware security modules, and integration with select Oracle Cloud Infrastructure services. Use customer-managed keys when you need security governance, regulatory compliance, and homogenous encryption of data, while centrally managing, storing, and monitoring the life cycle of the keys you use to protect your data.

You can:

- Enable customer-managed keys when you create databases in Exadata Cloud Infrastructure
- Switch from Oracle-managed keys to customer-managed keys
- Rotate your keys to maintain security compliance

Requirements

To enable management of customer-managed encryption keys, you must create a policy in the tenancy that allows a particular dynamic group to do so, similar to the following: `allow dynamic-group dynamic_group_name to manage keys in tenancy`.

Another policy is needed if the Vault being used by the customer is replicated (https://docs.oracle.com/en-us/iaas/Content/KeyManagement/Tasks/replicatingvaults.htm). For vaults that are replicated, this policy is needed: `allow dynamic-group dynamic_group_name to read vaults in tenancy`

Limitations

To enable Data Guard on Exadata Cloud Infrastructure databases that use customer-managed keys, the primary and standby databases must be in the same realm.

Task 1. Create a Vault and a Master Encryption Key

Create a vault in the Vault service by following the instructions in To create a new vault in Oracle Cloud Infrastructure Documentation. When following these instructions, Oracle recommends that you create the vault in a compartment created specifically to contain the vaults containing customer-managed keys, as described in Before You Begin: Compartment Hierarchy Best Practice.

After creating the vault, create at least one master encryption key in the vault by following the instructions in To create a new master encryption key in Oracle Cloud Infrastructure Documentation. When following these instructions, make these choices:

- **Create in Compartment**: Oracle recommends that you create the master encryption key in the same compartment as its vault; that is, the compartment created specifically to contain the vaults containing customer-managed keys.
- **Protection Mode**: Choose an appropriate value from the drop-down list:
  - **HSM** to create a master encryption key that is stored and processed on a hardware security module (HSM).
  - **Software** to create a master encryption key that is stored in a software file system in the Vault service. Software-protected keys are protected at rest using an HSM-based root key. You may export software keys to other key management devices or to a different OCI cloud region. Unlike HSM keys, software-protected keys are free of cost.
- **Key Shape Algorithm**: AES
- **Key Shape Length**: 256 bits

Oracle strongly recommends that you create a separate master encryption key for each of your container databases (CDBs). Doing so makes management of key rotation over time much simpler.

Task 2. Create a Service Gateway, a Route Rule, and an Egress Security Rule

Create a service gateway in the VCN (Virtual Cloud Network) where your Oracle Exadata Database Service on Dedicated Infrastructure resources reside by following the instructions in Task 1: Create the service gateway in Oracle Cloud Infrastructure Documentation.

After creating the service gateway, add a route rule and an egress security rule ***to each subnet*** (in the VCN) where Oracle Exadata Database Service on Dedicated Infrastructure resources reside so that these resources can use the gateway to access the Vault service:

1. Go to the **Subnet Details** page for the subnet.
2. In the **Subnet Information** tab, click the name of the subnet's **Route Table** to display its **Route Table Details** page.
3. In the table of existing **Route Rules**, check whether there is already a rule with the following characteristics:
   - **Destination**: All IAD Services In Oracle Services Network
   - **Target Type**: Service Gateway
   - **Target**: The name of the service gateway you just created in the VCN

   If such a rule does not exist, click **Add Route Rules** and add a route rule with these characteristics.
4. Return to the **Subnet Details** page for the subnet.
5. In the subnet's **Security Lists** table, click the name of the subnet's security list to display its **Security List Details** page.
6. In the side menu, under **Resources**, click **Egress Rules**.
7. In the table of existing **Egress Rules**, check whether there is already a rule with the following characteristics:
   - **Stateless**: No
   - **Destination**: All IAD Services In Oracle Services Network
   - **IP Protocol**: TCP
   - **Source Port Range**: All
   - **Destination Port Range**: 443

   If such a rule does not exist, click **Add Egress Rules** and add an egress rule with these characteristics.

Task 3. Create a Dynamic Group and a Policy Statement

To grant your Oracle Exadata Database Service on Dedicated Infrastructure resources permission to access customer-managed keys, you create an IAM dynamic group that identifies these resources and then create an IAM policy that grants this dynamic group access to the master encryption keys you created in the Vault service.

When defining the dynamic group, you identify your Oracle Exadata Database Service on Dedicated Infrastructure resources by specifying the OCID of the compartment containing your Exadata Infrastructure resource.

1. Copy the OCID of the compartment containing your Exadata Infrastructure resource. You can find this OCID on the **Compartment Details** page of the compartment.
2. Create a dynamic group by following the instructions in To create a dynamic group in Oracle Cloud Infrastructure Documentation. When following these instructions, enter a matching rule of this format:

   ```
   ALL {resource.compartment.id ='<compartment-ocid>'}
   ```

   where *<compartment-ocid>* is the OCID of the compartment containing your Exadata Infrastructure resource.

After creating the dynamic group, navigate to (or create) an IAM policy in a compartment higher up in your compartment hierarchy than the compartment containing your vaults and keys. Then, add a policy statement of this format:

```
allow dynamic-group <dynamic-group-name>
to manage keys
in compartment <vaults-and-keys-compartment>
where all {
target.key.id='<key_ocid>',
request.permission!='KEY_DELETE',
request.permission!='KEY_MOVE',
request.permission!='KEY_IMPORT',
request.permission!='KEY_BACKUP'
}
```

If you are using a replicated virtual private vault for the Oracle Data Guard deployment, add an additional policy statement in this format:

```
allow dynamic-group <dynamic-group>
to read vaults
in tenancy | compartment <vaults-and-keys-compartment>
```

where *<dynamic-group>* is the name of the dynamic group you created and *<vaults-and-keys-compartment>* is the name of the compartment in which you created your vaults and master encryption keys.

Related Topics

- To create a database in an existing Exadata Cloud Infrastructure instance
  This topic covers creating your first or subsequent databases.
- To administer Vault encryption keys
  Use this procedure to rotate the Vault encryption key or or change the encryption management configuration.
- Known Issues for Exadata Cloud Infrastructure and Data Guard
  Possible TDE key replication issue, and MRP and DG LCM operation failures.
- To integrate customer-managed key management into Exadata Cloud Infrastructure
  If you choose to encrypt databases in an Exadata Cloud Infrastructure instance using encryption keys that you manage, then you may update the following two packages (using Red Hat Package Manager) to enable DBAASTOOLS to interact with the APIs that customer-managed key management uses.

dbaascli database addInstance

To add the database instance on the specified node, use the `dbaascli database addInstance` command.

Prerequisite

- Run the command as the `root` user.

Syntax

```
dbaascli database addInstance --dbname <value> --node <value> [--newNodeSID
<value>]
```

Where:

- `--dbname` specifies Oracle Database name
- `--node` specifies the node name for the database instance
  - `--newNodeSID` specifies SID for the instance to add in the new node

dbaascli database convertToPDB

To convert the specified non-CDB database to PDB, use the `dbaascli database convertToPDB` command.

Syntax

```
dbaascli database convertToPDB --dbname <value> [--cdbName <value>] [--
executePrereqs]
        {
            [--copyDatafiles [--keepSourceDB]]|[backupPrepared]
        }
        [--targetPDBName <value>] [--waitForCompletion <value>] [--resume [--
sessionID <value>]]
```

Where:

- `--dbname` specifies the name of Oracle Database
- `--cdbName` specifies the name of the target CDB in which the PDB will be created. If the CDB does not exist, then it will be created in the same Oracle home as the source non-CDB
- `--executePrereqs` specifies to run only the pre-conversion checks
- `--copyDatafiles` specifies to create a new copy of the data files instead of using the ones from the source database
  `--keepSourceDB` - to preserve the source database after completing the operation.
- `--backupPrepared` - flag to acknowledge that a proper database backup is in place for the non CDB prior to performing the conversion to PDB.
- `--backupPrepared` flag to acknowledge that a proper database backup is in place for the non-CDB prior to performing the conversion to PDB
- `--targetPDBName` specifies the name of the PDB that will be created as part of the operation
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true|false`

- `--resume` specifies to resume the previous execution
  - `--sessionID` specifies to resume a specific session ID

Example 5-3 dbaascli database convertToPDB

To run pre-conversion prechecks:

```
dbaascli database convertToPDB --dbname ndb19 --cdbname cdb19 --
backupPrepared --executePrereqs
```

To run a full conversion with a copy of the data files from the non-CDB:

```
dbaascli database convertToPDB --dbname tst19 --cdbname cdb19 --copyDatafiles
```

dbaascli database getDetails

This command shows the detailed information of a given database e.g. dbname, node information, pluggable databases information etc.

Prerequisites

Run the command as the `root` user or the `oracle` user

Syntax

```
dbaascli database getDetails --dbname <value>
```

Where :

- `--dbname` - Oracle database name.

dbaascli database modifyParameters

To modify or reset initialization parameters for an Oracle Database, use the `dbaascli database modifyParameters` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli database modifyParameters --dbname <value> --setParameters <values>|
--resetParameters <values> | --responseFile
[--backupPrepared]
[--instance]
[--allowBounce]
```

Where:

- `--dbname` specifies the name of the database.
- `--setParameters` specifies a comma-delimited list of parameters to modify with new values. For example: `parameter1=valueA`,`parameter2=valueB`, and so on. For blank values use parameter1=valueA,parameter2='',etc.
- `--resetParameters` specifies a comma-delimited list of parameters to be reset to their corresponding default values. For example, `parameter1`,`parameter2`, and so on.

- `--responseFile` specifies the absolute location of the response JSON file to modify the database parameters
- `--backupPrepared` acknowledges that a proper database backup is in place prior to modifying critical or sensitive parameters.
- `--instance` specifies the name of the instance on which the parameters will be processed. If not specified, then the operation will be performed at the database level.
- `--allowBounce` grants permission to bounce the database in order to reflect the changes on applicable static parameters.

Example 5-4 dbaascli database modifyParameters

```
dbaascli database modifyParameters --dbname dbname --setParameters
"log_archive_dest_state_17=ENABLE"
```

dbaascli database upgrade

To upgrade an Oracle Database, use the `dbaascli database upgrade` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli database upgrade --dbname <value>
{--targetHome <value> | --targetHomeName <value>}
{ [--executePrereqs | --postUpgrade | --rollback]}
{[--standBy | --allStandbyPrepared]}
{[--upgradeOptions <value>]  | [--standBy]}
[--removeGRP]
[--increaseCompatibleParameter]
[--resume [--sessionID <value>]]
[--waitForCompletion <value>]
```

Where:

- `--dbname` (mandatory) specifies the name of the database.
- `--targetHome` specifies the target Oracle home location
- `--targetHomeName` specifies the name of the target Oracle Database home
- `--standBy` use this option to upgrade standby databases in Data Guard configurations
- `--allStandbyPrepared` required for Data Guard configured primary databases. Flags to acknowledge that all the required operations are performed on the standby databases prior to upgrading primary database
- `--removeGRP` automatically removes the Guaranteed Restore Point (GRP) backup only if the database upgrade was successful
- `--increaseCompatibleParameter` automatically increases the compatible parameter as part of the database upgrade. The parameter will get increased only if the database upgrade was successful
- `--executePrereqs` runs only the preupgrade checks

- `--postUpgrade` use this option if postupgrade fails and needs to rerun the postupgrade steps
- `--rollback` reverts an Oracle Database to its original Oracle home
- `--upgradeOptions` use this option to pass DBUA-specific arguments to perform the Oracle Database upgrade. Refer to the corresponding Oracle documentation for the supported arguments and options.
  `--standby`
- `--resume` to resume the previous execution
- `--sessionID` to resume a specific session id.
- `--waitForCompletion` specify false to run the operation in background. Valid values : true| false.

Example 5-5 dbaascli database upgrade pre-upgrade requisite checks

```
dbaascli database upgrade --dbbname dbname --targetHome Target Oracle home
location --executePrereqs
```

Prerequisites for Patching and Updating an Exadata Cloud Infrastructure System

The Exadata Cloud Infrastructure instance requires access to the Oracle Cloud Infrastructure Object Storage service, including connectivity to the applicable Swift endpoint for Object Storage

Oracle recommends using a service gateway with the VCN to enable this access. For more information, see these topics:

- Network Setup for Exadata Cloud Infrastructure Instances: For information about setting up your VCN for the Exadata Cloud Service instance, including the service gateway.
- Object Storage FAQ

> **Note:**
>
> Ensure that the following conditions are met to avoid patching failures:
>
> - The `/u01` directory on the database host file system has at least 15 GB of free space for the execution of patching processes.
> - The Oracle Clusterware is up and running on the VM cluster.
> - All nodes of the VM cluster are up and running.

Using the Console to Patch and Update Exadata Cloud Infrastructure Instances

You can use the Console to view the history of patch operations on Exadata Cloud Infrastructure instances, apply patches, and monitor the status of patch operations.

- Patching Exadata Instances That Use the DB System Resource Model
  The tasks in this section describe how to apply patches and monitor the status of patch operations on Exadata DB systems and their Database Homes.

- Patching Exadata Instances That use the New Resource Model
  The tasks in this section describe how to apply patches and monitor the status of patch operations on cloud VM clusters and their Database Homes.
- Patching Individual Oracle Databases in an Exadata Cloud Infrastructure Instance
  This task explains how to patch a single Oracle Database in your Exadata Cloud Infrastructure instance by moving it to another Database Home.
- Viewing Patch History
  Each patch history entry represents an attempted patch operation and indicates whether the operation was successful or failed. You can retry a failed patch operation. Repeating an operation results in a new patch history entry.

Patching Exadata Instances That Use the DB System Resource Model

The tasks in this section describe how to apply patches and monitor the status of patch operations on Exadata DB systems and their Database Homes.

- To patch the Oracle Grid Infrastructure on an Exadata DB system
  How to apply patches and monitor the status of patch operations on Exadata DB systems
- To patch the Oracle Database software in a Database Home (DB system)
  How to apply patches and monitor the status of patch operations on Exadata Database Homes for DB Systems.

To patch the Oracle Grid Infrastructure on an Exadata DB system

How to apply patches and monitor the status of patch operations on Exadata DB systems

1. Open the navigation menu. Click **Oracle Database**, then click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
3. In the list of DB systems, click the name of the Exadata DB system you want to patch to display the DB system details.
4. Under **DB System Version**, click the **View** link beside the **Latest Patch Available** field.
5. Review the list of available patches for the DB system.
6. Click the Actions menu for the patch you are interested in, and then click one of the following actions:Run Precheck: Check for any prerequisites to make sure that the patch can be successfully applied.Apply: Applies the selected patch. Oracle highly recommends that you run the precheck operation for a patch before you apply it.
   - **Run Precheck**: Check for any prerequisites to make sure that the patch can be successfully applied.
   - **Apply**: Applies the selected patch. Oracle highly recommends that you run the precheck operation for a patch before you apply it.
7. Confirm when prompted.

The patch list displays the status of the operation. While a patch is being applied, the patch's status displays as **Patching** and the DB system's status displays as **Updating**. Lifecycle operations on the DB system and its resources might be temporarily unavailable. If patching completes successfully, the patch's status changes to **Applied** and the status of the DB system changes to **Available**. You can view more details about an individual patch operation by clicking **Patch History**.

To patch the Oracle Database software in a Database Home (DB system)

How to apply patches and monitor the status of patch operations on Exadata Database Homes for DB Systems.

> **Note:**
>
> This patching procedure updates the Oracle Database software for all databases located in the Database Home. To patch an individual database, you can move it to another Database Home that uses the desired Oracle Database software configuration.

1. Open the navigation menu. Click **Oracle Database**, then click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
3. In the list of DB systems, click the name of the Exadata DB system with the Database Home you want to patch to display the DB system details.
4. Under **Resources**, click **Database Homes**.
5. Click the name of the Database Home you want to patch to display the Database Home details.
6. Under **Database Software Version**, locate the **Latest Patch Available** field and click **View**.
7. Review the available patches for the Database Home. You can choose to patch using an Oracle-provided software image or a custom software image. Oracle-provide images are generally available release updates. Custom software images are created by your organization with a specified set of patches. See Oracle Database Software Images for information on creating custom software images. The image you use to patch must be based on either the latest version of the Oracle Database software release or one of the three prior versions of the release.
8. Click the Actions menu at the end of the table row that lists the patch you are interested in, and then click one of the following actions:
   - **Precheck:** Check for any prerequisites to make sure that the patch can be successfully applied.
   - **Apply:** Applies the selected patch. Oracle highly recommends that you run the precheck operation for a patch before you apply it.
9. Confirm when prompted.

The patch list displays the status of the operation. While a patch is being applied, the status of the patch displays as **Patching** and the status of the Database Home and the databases in it display as **Updating**. During the operation, each database in the home is stopped and then restarted. If patching completes successfully, the patch's status changes to **Applied** and the Database Home's status changes to **Available**. You can view more details about an individual patch operation by clicking **Patch History**.

Patching Exadata Instances That use the New Resource Model

The tasks in this section describe how to apply patches and monitor the status of patch operations on cloud VM clusters and their Database Homes.

- To patch the Oracle Grid Infrastructure on an Exadata cloud VM cluster
  How to apply patches and monitor the status of patch operations on cloud VM clusters.

- To patch the Oracle Database software in a Database Home

To patch the Oracle Grid Infrastructure on an Exadata cloud VM cluster

How to apply patches and monitor the status of patch operations on cloud VM clusters.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters**.
4. In the list of cloud VM clusters, click the name of the cluster you want to patch to display the cluster details.
5. Under **Version**, click the **View Patches** link beside the **Updates Available**field.
6. Review the list of available patches for the cloud VM cluster.
7. Click the Actions menu for the patch you are interested in, and then click one of the following actions:
   - **Run Precheck:** Check for any prerequisites to make sure that the patch can be successfully applied.
   - **Update Grid Infrastructure:** Applies the selected patch. Oracle highly recommends that you run the precheck operation for a patch before you apply it.
8. Confirm when prompted.

The patch list displays the status of the operation. While a patch is being applied, the patch's status displays as **Patching** and the cloud VM cluster's status displays as **Updating**. Lifecycle operations on the cluster and its resources might be temporarily unavailable. If patching completes successfully, the patch's status changes to **Applied** and the status of the cluster changes to **Available**. You can view more details about an individual patch operation by clicking **Update History**.

To patch the Oracle Database software in a Database Home

> **Note:**
>
> This patching procedure updates the Oracle Database software for all databases located in the Database Home. To patch an individual database, you can To move a database to another Database Home that uses the desired Oracle Database software configuration.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters**.
4. In the list of cloud VM clusters, click the name of the cluster you want to patch to display the cluster details.
5. Under **Resources**, click **Database Homes**.
6. Click the name of the Database Home you want to patch to display the Database Home details.
7. Under **Database Software Version**, locate the **Latest Patch Available** field and click **View**.

8. Review the available patches for the Database Home. You can choose to patch using an Oracle-provided software image or a custom software image. Oracle-provide images are generally available release updates. Custom software images are created by your organization with a specified set of patches. See Oracle Database Software Images for information on creating custom software images. The image you use to patch must be based on either the latest version of the Oracle Database software release or one of the three prior versions of the release.
9. Click the Actions menu at the end of the table row that lists the patch you are interested in, and then click one of the following actions:
   - **Precheck:** Check for any prerequisites to make sure that the patch can be successfully applied.
   - **Apply:** Applies the selected patch. Oracle highly recommends that you run the precheck operation for a patch before you apply it.
10. Confirm when prompted.

    The patch list displays the status of the operation. While a patch is being applied, the status of the patch displays as **Patching** and the status of the Database Home and the databases in it display as **Updating**. During the operation, each database in the home is stopped and then restarted. If patching completes successfully, the patch's status changes to **Applied** and the Database Home's status changes to **Available**. You can view more details about an individual patch operation by clicking **Update History**.

Patching Individual Oracle Databases in an Exadata Cloud Infrastructure Instance

This task explains how to patch a single Oracle Database in your Exadata Cloud Infrastructure instance by moving it to another Database Home.

For information on patching Database Homes, see To patch the Oracle Database software in a Database Home (cloud VM cluster)

- To move a database to another Database Home
  This task explains how to patch a single Oracle Database in your Exadata Cloud Infrastructure instance by moving it to another Database Home.

To move a database to another Database Home

This task explains how to patch a single Oracle Database in your Exadata Cloud Infrastructure instance by moving it to another Database Home.

You can move a database to any Database Home that meets at either of the following criteria:

- The target Database Home uses the same Oracle Database software version (including patch updates) as the source Database Home
- The target Database Home is based on either the latest version of the Oracle Database software release used by the database, or one of the three prior versions of the release

Moving a database to a new Database Home brings the database up to the patch level of the target Database Home. For information on patching Database Homes, see Database Home Patching and .

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the database you want to move.:
   Cloud VM clusters (The New Exadata Cloud Infrastructure Resource Model *):* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM**

Clusters. In the list of VM clusters, click the name of the VM cluster that contains the database you wan to move.

DB systems: Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find you want to access, and then click the name of the Exadata DB system that contains the database you want to move..

4. Click **More Actions**, then click **Move to Another Home**.
5. Select the target Database Home.
6. Click **Move Database**.
7. Confirm the move operation.

   The database is moved in a rolling fashion. The database instance will be stopped, node by node, in the current home and then restarted in the destination home. While the database is being moved, the Database Home status displays as **Moving Databse**. When the operation completes, Database Home is updated with the current home. Datapatch is executed automatically, as part of the database move, to complete post-patch SQL actions for all patches, including one-offs, on the new Database Home. If the database move operation is unsuccessful, then the status of the database displays as `Failed`, and the Database Home field provides information about the reason for the failure.

Viewing Patch History

Each patch history entry represents an attempted patch operation and indicates whether the operation was successful or failed. You can retry a failed patch operation. Repeating an operation results in a new patch history entry.

Patch history views in the Console do not show patches that were applied by using command line tools such as `dbaascli`.

If your service instance uses the new resource model, the patch history available by navigating to the VM Cluster Details page. If your service instance uses the DB system resource model, the patch history is available by navigating to the DB System Details page.

- To view the patch history of a cloud VM cluster
  Each patch history entry represents an attempted patch operation and indicates whether the operation was successful or failed.
- To view the patch history of a DB system
  Each patch history entry represents an attempted patch operation and indicates whether the operation was successful or failed. You can retry a failed patch operation. Repeating an operation results in a new patch history entry. For a service instance using the DB system resource model, the patch history is available by navigating to the DB System Details page
- To view the patch history of a Database Home
  Each patch history entry represents an attempted patch operation and indicates whether the operation was successful or failed. You can retry a failed patch operation. Repeating an operation results in a new patch history entry. If your service instance uses the new resource model, the patch history available by navigating to the VM Cluster Details page.

To view the patch history of a cloud VM cluster

Each patch history entry represents an attempted patch operation and indicates whether the operation was successful or failed.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.

3. Click **Exadata VM Clusters**.
4. In the list of cloud VM clusters, click the name of the cluster you want to patch to display the cluster details.
5. Under **Version**, click the **View Patches** link beside the **Updates Available** field.
6. Click **Update History**.

   The Update History page displays the history of patch operations for that cloud VM cluster and for the Database Homes on that cloud VM cluster.

To view the patch history of a DB system

Each patch history entry represents an attempted patch operation and indicates whether the operation was successful or failed. You can retry a failed patch operation. Repeating an operation results in a new patch history entry. For a service instance using the DB system resource model, the patch history is available by navigating to the DB System Details page

1. Open the navigation menu. Click **Oracle Database**, then click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
3. In the list of DB systems, click the name of the Exadata DB system with the information you want to view to display the DB system details.
4. Under **DB System Version**, click the **View** beside the **Latest Patch Available** field.
5. Click **Patch History**.

The Patch History page displays the history of patch operations for that DB system and for the Database Homes on that DB system.

To view the patch history of a Database Home

Each patch history entry represents an attempted patch operation and indicates whether the operation was successful or failed. You can retry a failed patch operation. Repeating an operation results in a new patch history entry. If your service instance uses the new resource model, the patch history available by navigating to the VM Cluster Details page.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system that contains the Database Home.
   - *Cloud VM clusters (*The Exadata Cloud Infrastructure Resource Model*)* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
   - *DB systems* Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. Under **Resources**, click **Database Homes**.
5. Click the name of the Database Home you want to view to display the Database Home details.
6. Under **Database Software Version**, click **View** by the **Latest Patch Available** field.
7. Click **Patch History** (DB systems) or **Update History** (cloud VM clusters).

   The history page displays the history of patch operations for that Database Home and for the cloud VM cluster or DB system to which it belongs.

Using the API to Patch an Exadata Cloud Infrastructure Instance

Use these API operations to manage patching the following Exadata resources: cloud VM clusters, DB systems, databases, and Database Homes.

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.

Cloud VM clusters *(for systems using* the new resource model*)*:

- ListCloudVmClusterUpdates
- ListCloudVmClusterUpdateHistoryEntries
- GetCloudVmClusterUpdate
- GetCloudVmClusterUpdateHistoryEntry
- UpdateVmCluster

DB systems:

- ListDbSystemPatches
- ListDbSystemPatchHistoryEntries
- GetDbSystemPatch
- GetDbSystemPatchHistoryEntry
- UpdateDbSystem

Databases:

- UpdateDatabase - Use this operation to patch a database by moving it to another Database Home

Database Homes:

- ListDbHomePatches
- ListDbHomePatchHistoryEntries
- GetDbHomePatch
- GetDbHomePatchHistoryEntry
- UpdateDbHome

For the complete list of APIs for the Database service, see Database Service API.

Updating an Exadata Cloud VM Cluster Operating System

Exadata VM cluster image updates allow you to update the OS image on your Exadata cloud VM cluster nodes in an automated manner from the OCI console and APIs.

This automated feature simplifies and speeds up VM cluster patching, makes patching less error-prone, and eliminates the need to use Patch Manager.

When you apply a patch, the system runs a precheck operation to ensure your cloud VM cluster, Exadata DB system, or Database Home meets the requirements for that patch. If the precheck is not successful, the patch is not applied, and the system displays a message that the patch cannot be applied because the precheck failed. A separate precheck operation that you can run in advance of the planned update is also available.

- Supported Software Versions and Update Restrictions
  Minimum requirements for updating to Exadata image release 23.1.0.0.0 (Oracle Linux 8-based image):
- Updating the Operating System using the Console

Supported Software Versions and Update Restrictions

Minimum requirements for updating to Exadata image release 23.1.0.0.0 (Oracle Linux 8-based image):

> **Note:**
>
> These are just the minimum requirements. If you want to update Grid Infrastructure and/or Oracle Database to meet the Exadata 23.1 requirements, then the recommendation is to update to the latest available versions of Grid Infrastructure and Oracle Database, and not to the minimum.

- **Exadata Image (Guest OS):** Exadata image release 22.1.0 (May 2022) or 21.2.10 (March 2022). Systems running versions older than 21.2.10 will first need to upgrade to at least 22.1.0 (May 2022) or 21.2.10 (March 2022) before updating to 23.1.0.0.0. This applies to both storage and database servers.
  - In addition to performing minor version updates to the Exadata VM Cluster images, you can update to a new major version if the currently installed version is 19.2 or higher. For example, if the VM cluster is on version 20, then you can update it to version 21.
  - The latest 4 (N to N-3) or more minor versions of each major version of the VM Cluster images are available through the console to apply.
- **Oracle Grid Infrastructure:** Exadata image release 23.1.0.0.0 supports the following minimum or newer Oracle Grid Infrastructure versions.
  - Release 19c: Version 19.15, April 2022 Release Update (RU) and newer (Default)
  - Release 21c: Version 21.6, April 2022 Release Update (RU) and newer
- **Oracle Database:** Exadata System Software 23.1 supports the following minimum versions or newer for new database installations.
  - Release 19c: Version 19.15, April 2022 Release Update (RU) and newer (Default)
  - Additional supported database releases under Market Driven Support or Quarterly Updates exception approval:
    * Release 12.2.0.1, Release Update (RU) 12.2.0.1.220118 (Jan 2022)
    * Release 12.1.0.2, Bundle Patch 12.1.0.2.220719 (Jul 2022) - requires patch 30159782
    * Release 11.2.0.4, Bundle Patch 11.2.0.4.210119 (Jan 2021) - requires patch 30159782, patch 33991024
- If you have an Exadata infrastructure maintenance operation scheduled to start within the next 24 hours, then the Exadata Image update feature is not available.
- Once the VM cluster is upgraded to Exadata Database Service Guest VM OS 23.1, you will be able to add a new VM or a new database server to this VM cluster if Exadata Cloud Infrastructure is running an Exadata System Software version 22.1.16 and later.

> **Note:**
>
> Upgrade to Exadata System Software 23.1 for Exadata Cloud Infrastructure will be available with February 2024 update cycle.

Updating the Operating System using the Console

> **Note:**
>
> Once the VM cluster is upgraded to Exadata Database Service Guest VM OS 23.1, you will be able to add a new VM or a new database server to this VM cluster if Exadata Cloud Infrastructure is running an Exadata System Software version 22.1.16 and later.
> Upgrade to Exadata System Software 23.1 for Exadata Cloud Infrastructure will be available with February 2024 update cycle.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.
3. In the list of cloud VM clusters, click the name of the cluster that you want to patch to display the details page.
4. In the **Version** section, to the right of the **Updates Available**, click **View Updates** to display the **Updates** page.
5. Review the list of available software updates and locate the OS patch you are applying.
6. Click the Actions icon (three dots) at the end of the row listing the patch you are interested in, and then click one of the following actions:
   - **Run Precheck.** Precheck checks the prerequisites to ensure that the patch can be successfully applied. Oracle highly recommends that you run the precheck operation before you apply a patch. The reason is that things can change in a database at any time, and the precheck you run just before running a patch may find errors that the previous precheck did not find

     > **Note:**
     >
     > If the precheck fails, the system displays a message in the **Apply Exadata OS Image Update** dialog that the last precheck has failed. Oracle recommends that you run the precheck again. Click the Actions icon (three dots) at the end of the row listing the OS patch to view the dialog.

   - **Apply Exadata OS Image Update**. This link displays the Apply Exadata Image Update dialog that you use to apply the patch. The dialog shows the name of the database system you are patching, the current version of the database, and the new version of the database after the patch is applied. To start the process, click **Apply Exadata OS Image Update**.

- **Copy OCID.** This copies the Oracle Cloud ID. This can be used when troubleshooting a patch or to give to Support when contacting them.

> **Note:**
>
> While the patch is running:
>
> – Run Precheck and Apply OS Image Update are not available. When the patch has completed, these actions are available again.
>
> – If the Exadata infrastructure containing this VM cluster is scheduled for maintenance that conflicts with the patching operation, the patch fails and the system displays a message explaining why. After the infrastructure maintenance is complete, run the patch operation again.

7. Confirm when prompted.

The patch list displays the status of the operation in the Version section of the database details page. Click **View Updates** to view more details about an individual patch status and to display any updates that are available to run. If no new updates are available, the system displays a message that says **No Updates Available**.

Upgrading Exadata Grid Infrastructure

This topic describes how to upgrade the Oracle Grid Infrastructure (GI) on an Exadata cloud VM cluster using the Oracle Cloud Infrastructure Console or API.

Upgrading allows you to provision Oracle Database Homes and databases that use the most current Oracle Database software. For more information on Exadata cloud VM clusters and the new Exadata resource model, see Overview of X8M and X9M Scalable Exadata Infrastructure .

- Prerequisites for Upgrading Exadata Grid Infrastructure
  To upgrade your GI to Oracle Database 19c, you must be using the Oracle Linux 7 operating system for your VM cluster.
- About Upgrading Oracle Grid Infrastructure
  Upgrading the Oracle Grid Infrastructure (GI) on a VM cluster involves upgrading all the compute nodes in the instance. The upgrade is performed in a rolling fashion, with only one node being upgraded at a time.
- Using the console to upgrade your Grid Infrastructure
  You can use the Console to perform a precheck prior to upgrading your Oracle Grid Infrastructure (GI), and to perform the GI upgrade operation.
- Using the API to Upgrade the Grid Infrastructure in a VM Cluster

Prerequisites for Upgrading Exadata Grid Infrastructure

To upgrade your GI to Oracle Database 19c, you must be using the Oracle Linux 7 operating system for your VM cluster.

For more information on upgrading the operating system, see the following document:

- How to update the Exadata System Software (DomU) to 19 from 18 on the Exadata Cloud Service in OCI(My Oracle Support Doc ID 2521053.1).

About Upgrading Oracle Grid Infrastructure

Upgrading the Oracle Grid Infrastructure (GI) on a VM cluster involves upgrading all the compute nodes in the instance. The upgrade is performed in a rolling fashion, with only one node being upgraded at a time.

- Oracle recommends running an upgrade precheck to identify and resolve any issues that would prevent a successful upgrade.
- You can monitor the progress of the upgrade operation by viewing the associated *work requests*.
- If you have an Exadata infrastructure maintenance operation scheduled to start within the next 24 hours, then the GI upgrade feature is not available.
- During the upgrade, you cannot perform other management operations such as starting, stopping, or rebooting nodes, scaling CPU, provisioning or managing Database Homes or databases, restoring a database, or editing IORM settings. The following Data Guard operations are not allowed on the VM cluster undergoing a GI upgrade:
  - Enable Data Guard
  - Switchover
  - Failover to the database using the VM cluster (a failover operation to standby on another VM cluster is possible)

Related Topics

- Work Requests Integration

Using the console to upgrade your Grid Infrastructure

You can use the Console to perform a precheck prior to upgrading your Oracle Grid Infrastructure (GI), and to perform the GI upgrade operation.

- To Precheck your Cloud VM Cluster prior to Upgrading
- To Upgrade the Oracle Grid Infrastructure of a Cloud VM Cluster
  Procedure for upgrading the Oracle Grid Infrastructure of a Cloud VM Cluster.

To Precheck your Cloud VM Cluster prior to Upgrading

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters**.
4. In the list of cloud VM clusters, click the name of the cluster you want to patch to display the cluster details.
5. Under **Version**, click the **View Patches** link beside the **Updates Available**field.
6. Click **Updates** to view the list of available patches and upgrades.
7. Click the Actions icon (three dots) at the end of the row listing the Oracle Grid Infrastructure (GI) upgrade, then click **Run Precheck**.
8. In the **Confirm** dialog, confirm you want to upgrade to begin the precheck operation.

To Upgrade the Oracle Grid Infrastructure of a Cloud VM Cluster

Procedure for upgrading the Oracle Grid Infrastructure of a Cloud VM Cluster.

> **Note:**
>
> - When planning to upgrade your Grid Infrastructure to 23ai, make sure that for each ASM diskgroup, `compatible.rdbms` has a value set to 19.0.0.0 and later.
> - Minimum requirements for upgrading Grid Infrastructure from 19c to 23ai:
>   - Exadata Guest VM running Exadata System Software 23.1.8
>   - Exadata Infrastructure running Exadata System Software 23.1.x

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters**.
4. In the list of cloud VM clusters, click the name of the cluster you want to patch to display the cluster details.
5. Under **Version**, click the **View Patches** link beside the **Updates Available** field.
6. Click **Updates** to view the list of available patches and upgrades.
7. Click the Actions icon (three dots) at the end of the row listing the Oracle Grid Infrastructure (GI) upgrade, then click **Upgrade Grid Infrastructure**.
8. In the **Upgrade Grid Infrastructure** dialog, confirm you want to upgrade the GI by clicking **Upgrade Grid Infrastructure**. If you haven't run a precheck, you have the option of clicking **Run Precheck** in this dialog to precheck your cloud VM cluster prior to the upgrade.

Using the API to Upgrade the Grid Infrastructure in a VM Cluster

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.

Use these API operations to upgrade the Oracle Grid Infrastructure in a cloud VM clusters and view the cluster's update history.

- ListCloudVmClusterUpdates
- ListCloudVmClusterUpdateHistoryEntries
- GetCloudVmClusterUpdate
- GetCloudVmClusterUpdateHistoryEntry
- UpdateVmCluster

For the complete list of APIs for the Database service, see Database Service API.

Upgrading Exadata Databases

This topic describes the procedures to upgrade an Exadata database instance to Oracle Database 19c (Long Term Release) by using the Console and the API. The upgrade is accomplished by moving the Exadata database to a Database Home that uses the target software version.

> **Note:**
>
> This topic applies only to Exadata Cloud Infrastructure instances using the new resource model. For information on converting an Exadata DB system to the new resource model, see Switching an Exadata DB System to the New Resource Model and APIs.

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

- Prerequisites to Upgrade Oracle Databases
  Review the list of prerequsites to upgrade an Exadata Cloud Infrastructure Oracle Database instance.
- About Upgrading a Database
- Using the Console to Upgrade a Database
  Procedures to precheck and upgrade a database, rollback a failed upgrade, and view the upgrade history.
- Using the API to upgrade Databases
  Use the following APIs to manage database upgrades:

Prerequisites to Upgrade Oracle Databases

Review the list of prerequsites to upgrade an Exadata Cloud Infrastructure Oracle Database instance.

- The Exadata Cloud Infrastructure system software must use Oracle Linux 7 (OL7). See *How to update the Exadata System Software (DomU) to 19 from 18 on the Exadata Cloud Service in OCI (My Oracle Support Doc ID 2521053.1)* for instructions on manually updating the operating system.
- The Oracle Grid Infrastructure must be version 19c. See *Upgrading Exadata Grid Infrastructure* for instructions on using the Oracle Cloud Infrastructure Console or API to upgrade Grid Infrastructure. If patches are available for your Grid Infrastructure, Oracle recommends applying them prior to performing a database upgrade.
- You must have an available Oracle Database Home that uses the four most recent versions of Oracle Database 19c available in Oracle Cloud Infrastructure. See *To Create a new Oracle Database Home in an existing Exadata Cloud Infrastructure Instance* for information on creating a Database Home. You can use Oracle-published software images or a *custom database software image* based on your patching requirements to create Database Homes.
- You must ensure that all pluggable databases in the container database that is being upgraded can be opened. Pluggable databases that cannot be opened by the system during the upgrade can cause an upgrade failure.

- If you are upgrading databases in a manually-created Data Guard association (an association not created using the Console or APIs), the following apply:
  - The databases must be registered with the Cloud tooling. See Updating Tooling on an Exadata Cloud Service Instance for more information.
  - Redo apply needs to be disabled during the upgrade of both the primary and standby. For Oracle 11.2 and 12.1 databases, the Data Guard configuration also has to be disabled.
  - If you have configured an observer, the observer needs to be disabled prior to upgrade.

Your Oracle database must be configured with the following settings in order to upgrade:

- The database must be in archive log mode.
- The database must have flashback enabled.

See the *Oracle Database documentation* for your database's release version to learn more about these settings.

Related Topics

- How to update the Exadata System Software (DomU) to 19 from 18 on the Exadata Cloud Service in OCI (Doc ID 2521053.1)
- Upgrading Exadata Grid Infrastructure
  This topic describes how to upgrade the Oracle Grid Infrastructure (GI) on an Exadata cloud VM cluster using the Oracle Cloud Infrastructure Console or API.
- To create a new Database Home in an existing Exadata Cloud Infrastructure instance
  To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.
- Oracle Database Software Images
- Oracle Database Documentation

About Upgrading a Database

For database software version upgrades, note the following:

- Database upgrades involve database downtime. Keep this in mind when scheduling your upgrade.
- Oracle recommends that you back up your database and test the new software version on a test system or a cloned version of your database before you upgrade a production database. See *to create an on-demand full backup of a database* for information on creating an on-demand manual backup.
- Oracle recommends running an upgrade precheck operation for your database prior to attempting an upgrade so that you can discover any issues that need mitigation prior to the time you plan to perform the upgrade. The precheck operation does not affect database availability and can be performed at any time that is convenient for you.
- If your databases uses Data Guard, you can upgrade either the primary or the standby first. To upgrade a primary, follow the steps in To upgrade or precheck an Exadata database. To upgrade a standby, follow the steps in To move a database to another Database Home
- If your databases uses Data Guard, upgrading a primary or standby will disable redo apply during the upgrade operation. After you upgrade both the primary and standby, redo apply

and open mode are re-enabled. Oracle recommends checking the redo apply and open mode configuration after upgrading.

- An upgrade operation cannot take place while an automatic backup operation is underway. Before upgrading, Oracle recommends disabling automatic backups and performing a manual backup. See *to configure automatic backups for a database* and *To create an on-demand full backup of a database* for more information.
- After upgrading, you cannot use automatic backups taken prior to the upgrade to restore the database to an earlier point in time.
- If you are upgrading an database that uses version 11.2 software, the resulting version 19c database will be a non-container database (non-CDB).
- How the Upgrade Operation Is Performed by the Database Service
  During the upgrade process, the Database service does the following:
- Rolling Back an Oracle Database Unsuccessful Upgrade
  If your upgrade does not complete successfully, then you have the option of performing a rollback.
- After Upgrading an Oracle Database
  After a successful upgrade, note the following:

Related Topics

- To create an on-demand backup of a database
- To configure automatic backups for a database

How the Upgrade Operation Is Performed by the Database Service

During the upgrade process, the Database service does the following:

- Executes an automatic precheck. This allows the system to identify issues needing mitigation and to stop the upgrade operation.
- Sets a guaranteed restore point, enabling it to perform a flashback in the event of an upgrade failure.
- Moves the database to a user-specified Oracle Database Home that uses the desired target software version.
- Runs the Database Upgrade Assistant (DBUA) software to perform the upgrade.
- For databases in Data Guard associations, redo apply is disabled until both the primary and standby databases are successfully upgraded, at which point redo apply is re-enabled by the system. The system then enables Open Mode after redo apply is enabled.

Rolling Back an Oracle Database Unsuccessful Upgrade

If your upgrade does not complete successfully, then you have the option of performing a rollback.

Details about the failure are displayed on the **Database Details** page in the Console, allowing you to analyze and resolve the issues causing the failure.

A rollback resets your database to the state prior to the upgrade. All changes to the database made during and after the upgrade will be lost. The rollback option is provided in a banner message displayed on the database details page of a database following an unsuccessful upgrade operation. See *Using the Console to Roll Back a Failed Database Upgrade* for more information.

For standby databases in Oracle Data Guard associations, rollback is accomplished by moving the standby back to the original Database Home. See To move a database to another Database Home for instructions.

Related Topics

- To roll back a failed database upgrade

After Upgrading an Oracle Database

After a successful upgrade, note the following:

- Check that automatic backups are enabled for the database if you disabled them prior to upgrading. See *Customizing the Automatic Backup Configuration* for more information.
- Edit the Oracle Database `COMPATIBLE` parameter to reflect the new Oracle Database software version. See *What Is Oracle Database Compatibility?* for more information.
- If your database uses a *`database_name`*`.env` file, ensure that the variables in the file have been updated to point to the 19c Database Home. These variables should be automatically updated during the upgrade process.
- If you are upgrading a non-container database to Oracle Database version 19c, you can convert the database to a pluggable database after converting. See *How to Convert Non-CDB to PDB (Doc ID 2288024.1)* for instructions on converting your database to a pluggable database.
- If your old Database Home is empty and will not be reused, you can remove it. See *Using the Console to Delete an Oracle Database Home* for more information.
- For databases in Data Guard associations, check the open mode and redo apply status after the upgrade is complete.

Related Topics

- Managing Exadata Database Backups by Using bkup_api
- What Is Oracle Database Compatibility?
- How to Convert Non-CDB to PDB - Step by Step Example (Doc ID 2288024.1)
- To delete a database home
  You cannot delete a Database Home that contains databases. You must first terminate the databases to empty the Database Home. See To terminate a database to learn how to terminate a database.

Using the Console to Upgrade a Database

Procedures to precheck and upgrade a database, rollback a failed upgrade, and view the upgrade history.

- To upgrade or precheck an Exadata database
  Procedure to upgrade or precheck an Exadata database.
- To roll back a failed database upgrade
- To view the the upgrade history of a database

To upgrade or precheck an Exadata database

Procedure to upgrade or precheck an Exadata database.

The following steps apply to databases for which either of the following apply:

- The database is the primary database in a Data Guard association
- The database is not part of a Data Guard association

To upgrade a standby database in a Data Guard configuration, move the standby to a Database Home using the Oracle Database version you are upgrading to. See To move a database to another Database Home for details.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Under**Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, click the name of the VM cluster that contains the database you want to upgrade.

> **Note:**
>
> If your database is in an Exadata Cloud Infrastructure instance that does not use the new Exadata resource model, you will need to switch the instance to the new model before you can upgrade your database.

4. In the list of databases on the details page of the VM cluster, click the name of the database you want to upgrade to view the Database Details page.
5. Click **More Actions**, then **Upgrade**.
6. In the **Upgrade Database** dialogue, select the following:
   - **Oracle Database version:** The drop-down selector lists only Oracle Database versions that are compatible with an upgrade from the current software version the database is using. The target software version must be higher than the database's current version.
   - **Target Database Home:** Select a Database Home for your database. The list of Database Homes is limited to those homes using the most recent versions of Oracle Database 19c software. Moving the database to the new Database Home results in the database being upgraded to the major release version and patching level of the new Database Home.
7. Click one of the following:
   - **Run Precheck:** This option starts an upgrade precheck to identify any issues with your database that need mitigation before you perform an upgrade.
   - **Upgrade Database:** This option starts upgrade operation. Oracle recommends performing an upgrade only after you have performed a successful precheck on the database.

To roll back a failed database upgrade

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure** .
2. Choose your **Compartment**.

   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, click the name of the VM cluster that contains the database with the failed upgrade.

4. Find the database that was unsuccessfully upgraded, and click its name to display details about it.
5. The database must display a banner at the top of the details page that includes a **Rollback** button and details about what issues caused the upgrade failure.
6. Click **Rollback**.
7. In the **Confirm rollback** dialog, confirm that you want to initiate a rollback to the previous Oracle Database version.

To view the the upgrade history of a database

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, click the name of the VM cluster that contains the database you want to upgrade.

> **Note:**
>
> If your database is in an Exadata Cloud Infrastructure instance that does not use the *new Exadata resource model* , you will need to *switch the instance to the new Exadata resource model* before you can upgrade your database.

4. In the list of databases on the details page of the VM cluster, click the name of the database for which you want to view the upgrade history.
5. On the Database Details page, under **Database Version**, click the **View** link that is displayed for databases that have been upgraded. This link does not appear for databases that have not been updated.
   The **Updates History** page is displayed. The table displayed on this page shows precheck and upgrade operations performed on the database.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.
- Switch an Exadata DB system to the new Exadata resource model
  Use the console to perform the switch to the new resource model.

Using the API to upgrade Databases

Use the following APIs to manage database upgrades:

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.

Use these API operations to manage database upgrades:

- ListDatabaseUpgradeHistoryEntries
- UpgradeDatabase

For the complete list of APIs for the Database service, see Database Service API.

> **Note:**
>
> When using the `UpgradeDatabase` API to upgrade an Exadata Cloud Infrastructure database, you must specify `DB_HOME` as the upgrade source.

Patching and Updating an Exadata Cloud Infrastructure System Manually

This topic describes the procedures for patching and updating various components in Exadata Cloud Service outside of the cloud automation.

For information related to patching and updating with dbaascli, refer to "*Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli*".

> **Note:**
>
> For more guidance on achieving continuous service during patching operations, see the *Application Checklist for Continuous Service for MAA Solutions* white paper.

- Patching Oracle Database and Oracle Grid Infrastructure Software Manually
  For daylight savings time, and some routine or one-off patches, it can be necessary for you to patch software manually.
- Updating the Exadata Cloud VM Cluster OS Manually
  You update the operating systems of Exadata compute nodes by using the patchmgr tool.
- Updating Tooling on an Exadata Cloud Infrastructure Instance
  Cloud-specific tooling is used on the Exadata Cloud Infrastructure Guest VMs for local operations, including dbaascli commands.

Related Topics

- Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli
  Learn to use the `dbaascli` utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata Cloud Infrastructure system.
- Application Checklist for Continuous Service for MAA Solutions

Patching Oracle Database and Oracle Grid Infrastructure Software Manually

For daylight savings time, and some routine or one-off patches, it can be necessary for you to patch software manually.

To perform routine patching of Oracle Database and Oracle Grid Infrastructure software, Oracle recommends that you use the facilities provided by Oracle Exadata Database Service on Dedicated Infrastructure. However, under some circumstances, it can be necessary for you to patch the Oracle Database or Oracle Grid Infrastructure software manually:

- **Daylight Savings Time (DST) Patching:** Because they cannot be applied in a rolling fashion, patches for the Oracle Database DST definitions are not included in the routine patch sets for Exadata Cloud Infrastructure. If you need to apply patches to the Oracle Database DST definitions, you must do so manually. See My Oracle Support Doc ID 412160.1.

- **Non-routine or One-off Patching:** If you encounter a problem that requires a patch which is not included in any routine patch set, then work with Oracle Support Services to identify and apply the appropriate patch.

For general information about patching Oracle Database, refer to information about patch set updates and requirements in *Oracle Database Upgrade Guide* for your release.

Related Topics

- https://support.oracle.com/epmos/faces/DocumentDisplay?cmd=show&type=NOT&id=1929745.1
- https://support.oracle.com/epmos/faces/DocumentDisplay?cmd=show&type=NOT&id=412160.1

Updating the Exadata Cloud VM Cluster OS Manually

You update the operating systems of Exadata compute nodes by using the patchmgr tool.

This utility manages the entire update of one or more compute nodes remotely, including running pre-reboot, reboot, and post-reboot steps. You can run the utility from either an Exadata compute node or a non-Exadata server running Oracle Linux. The server on which you run the utility is known as the "driving system." You cannot use the driving system to update itself. Therefore, if the driving system is one of the Exadata compute nodes on a system you are updating, you must run a separate operation on a different driving system to update that server.

The following two scenarios describe typical ways of performing the updates:

Scenario 1: Non-Exadata Driving System

The simplest way to run the update the Exadata system is to use a separate Oracle Linux server to update all Exadata compute nodes in the system.

Scenario 2: Exadata Node Driving System

You can use one Exadata compute node to drive the updates for the rest of the compute nodes in the system, and then use one of the updated nodes to drive the update on the original Exadata driver node.

For example: You are updating a half rack Exadata system, which has four compute nodes - node1, node2, node3, and node4. First, use node1 to drive the updates of node2, node3, and node4. Then, use node2 to drive the update of node1.

The driving system requires root user `SSH` access to each compute node the utility will update.

- Preparing for the OS Updates
  Determine the latest software version available, and connectivity to the proper `yum` repository
- To update the OS on all compute nodes of an Exadata Cloud Infrastructure instance
  Procedure to update all compute nodes using `patchmgr`.
- Installing Additional Operating System Packages
  Review these guidelines before you install additional operating system packages for Oracle Exadata Database Service on Dedicated Infrastructure.

Preparing for the OS Updates

Determine the latest software version available, and connectivity to the proper `yum` repository

> **Caution:**
>
> Do not install NetworkManager on the Exadata Cloud Infrastructure instance. Installing this package and rebooting the system results in severe loss of access to the system.

- Before you begin your updates, review *Exadata Cloud Service Software Versions* (Doc ID 2333222.1) to determine the latest software version and target version to use.
- Some steps in the update process require you to specify a YUM repository. The YUM repository URL is:

  ```
  http://yum-<region_identifier>.oracle.com/repo/EngineeredSystems/exadata/
  dbserver/<latest_version>/base/x86_64.
  ```

  Region identifiers are text strings used to identify Oracle Cloud Infrastructure regions (for example, `us-phoenix-1`). You can find a complete list of region identifiers in *Regions* .

  You can run the following `curl` command to determine the latest version of the YUM repository for your Exadata Cloud Service instance region:

  ```
  curl -s -X GET http://yum-<region_identifier>.oracle.com/repo/
  EngineeredSystems/exadata/dbserver/ |egrep "18.1."
  ```

  This example returns the most current version of the YUM repository for the US West (Phoenix) region:

  ```
  curl -s -X GET http://yum-us-phoenix-1.oracle.com/repo/EngineeredSystems/
  exadata/dbserver/ |egrep "18.1."
  <a href="18.1.4.0.0/">18.1.4.0.0/</a> 01-Mar-2018 03:36 -
  ```

- To apply OS updates, the system's **VCN** must be configured to allow access to the YUM repository. For more information, see *Option 2: Service Gateway to Both Object Storage and YUM repos.* .

Related Topics

- Regions
- Option 2: Service Gateway Access to Both Object Storage and YUM Repos
  You configure *both the client subnet and backup subnet* to use the service gateway for access to the Oracle Services Network, which includes both Object Storage and the Oracle YUM repos.

To update the OS on all compute nodes of an Exadata Cloud Infrastructure instance

Procedure to update all compute nodes using `patchmgr`.

This example procedure assumes the following:

- The system has two compute nodes, `node1` and `node2`.
- The target version is 18.1.4.0.0.180125.3.
- Each of the two nodes is used as the driving system for the update on the other one.

1. Gather the environment details.

   a. SSH to node1 as root and run the following command to determine the version of Exadata:

   ```
   [root@node1]# imageinfo -ver
   12.2.1.1.4.171128
   ```

   b. Switch to the grid user, and identify all computes in the cluster.

   ```
   [root@node1]# su - grid
   [grid@node1]$ olsnodes
   node1
   node1
   ```

2. Configure the driving system.

   a. Switch back to the root user on node1, check whether a root ssh key pair (id_rsaand id_rsa.pub) already exists. If not, then generate it.

   ```
   [root@node1 .ssh]#  ls /root/.ssh/id_rsa*
   ls: cannot access /root/.ssh/id_rsa*: No such file or directory
   [root@node1 .ssh]# ssh-keygen -t rsa
   Generating public/private rsa key pair.
   Enter file in which to save the key (/root/.ssh/id_rsa):
   Enter passphrase (empty for no passphrase):
   Enter same passphrase again:
   Your identification has been saved in /root/.ssh/id_rsa.
   Your public key has been saved in /root/.ssh/id_rsa.pub.
   The key fingerprint is:
   93:47:b0:83:75:f2:3e:e6:23:b3:0a:06:ed:00:20:a5
   root@node1.fraad1client.exadataclientne.oraclevcn.com
   The key's randomart image is:
   +--[ RSA 2048]----+
   |o..     + .      |
   |o.     o *       |
   |E     . o o      |
   | . .     =       |
   |  o .   S =      |
   |   +     = .     |
   |    +   o o      |
   |   . .   + .     |
   |      ...        |
   +-----------------+
   ```

   b. Distribute the public key to the target nodes, and verify this step. In this example, the only target node is node2.

   ```
   [root@node1 .ssh]# scp -i ~opc/.ssh/id_rsa ~root/.ssh/id_rsa.pub
   opc@node2:/tmp/id_rsa.node1.pub
   id_rsa.pub

   [root@node2 ~]# ls -al /tmp/id_rsa.node1.pub
   -rw-r--r-- 1 opc opc 442 Feb 28 03:33 /tmp/id_rsa.node1.pub
   [root@node2 ~]# date
   Wed Feb 28 03:33:45 UTC 2018
   ```

c. On the target node (`node2`, in this example), add the root public key of `node1` to the root `authorized_keys` file.

```
[root@node2 ~]# cat /tmp/id_rsa.node1.pub >> ~root/.ssh/authorized_keys
```

d. Download `dbserver.patch.zip` as `p21634633_12*_Linux-x86-64.zip` onto the driving system (`node1`, in this example), and unzip it. See *dbnodeupdate.sh and dbserver.patch.zip*: *Updating Exadata Database Server Software using the DBNodeUpdate Utility and patchmgr (Doc ID 1553103.1)* for information about the files in this .zip.

```
[root@node1 patch]# mkdir /root/patch
[root@node1 patch]# cd /root/patch
[root@node1 patch]# unzip p21634633_181400_Linux-x86-64.zip
Archive:  p21634633_181400_Linux-x86-64.zip   creating:
dbserver_patch_5.180228.2/
   creating: dbserver_patch_5.180228.2/ibdiagtools/
  inflating: dbserver_patch_5.180228.2/ibdiagtools/cable_check.pl
  inflating: dbserver_patch_5.180228.2/ibdiagtools/setup-ssh
  inflating: dbserver_patch_5.180228.2/ibdiagtools/VERSION_FILE
 extracting: dbserver_patch_5.180228.2/ibdiagtools/xmonib.sh
  inflating: dbserver_patch_5.180228.2/ibdiagtools/monitord
  inflating: dbserver_patch_5.180228.2/ibdiagtools/checkbadlinks.pl
   creating: dbserver_patch_5.180228.2/ibdiagtools/topologies/
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/
VerifyTopologyUtility.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/
verifylib.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Node.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Rack.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Group.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Switch.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topology-zfs
  inflating: dbserver_patch_5.180228.2/ibdiagtools/dcli
   creating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
remoteScriptGenerator.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
CommonUtils.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
SolarisAdapter.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
LinuxAdapter.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
remoteLauncher.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
remoteConfig.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/spawnProc.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
runDiagnostics.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/OSAdapter.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/SampleOutputs.txt
  inflating: dbserver_patch_5.180228.2/ibdiagtools/infinicheck
  inflating: dbserver_patch_5.180228.2/ibdiagtools/ibping_test
  inflating: dbserver_patch_5.180228.2/ibdiagtools/tar_ibdiagtools
```

```
 inflating: dbserver_patch_5.180228.2/ibdiagtools/verify-topology
 inflating: dbserver_patch_5.180228.2/installfw_exadata_ssh
  creating: dbserver_patch_5.180228.2/linux.db.rpms/
 inflating: dbserver_patch_5.180228.2/md5sum_files.lst
 inflating: dbserver_patch_5.180228.2/patchmgr
 inflating: dbserver_patch_5.180228.2/xcp
 inflating: dbserver_patch_5.180228.2/ExadataSendNotification.pm
 inflating: dbserver_patch_5.180228.2/ExadataImageNotification.pl
 inflating: dbserver_patch_5.180228.2/kernelupgrade_oldbios.sh
 inflating: dbserver_patch_5.180228.2/cellboot_usb_pci_path
 inflating: dbserver_patch_5.180228.2/exadata.img.env
 inflating: dbserver_patch_5.180228.2/README.txt
 inflating: dbserver_patch_5.180228.2/exadataLogger.pm
 inflating: dbserver_patch_5.180228.2/patch_bug_26678971
 inflating: dbserver_patch_5.180228.2/dcli
 inflating: dbserver_patch_5.180228.2/patchReport.py
extracting: dbserver_patch_5.180228.2/dbnodeupdate.zip
  creating: dbserver_patch_5.180228.2/plugins/
 inflating: dbserver_patch_5.180228.2/plugins/010-check_17854520.sh
 inflating: dbserver_patch_5.180228.2/plugins/020-check_22468216.sh
 inflating: dbserver_patch_5.180228.2/plugins/040-check_22896791.sh
 inflating: dbserver_patch_5.180228.2/plugins/000-check_dummy_bash
 inflating: dbserver_patch_5.180228.2/plugins/050-check_22651315.sh
 inflating: dbserver_patch_5.180228.2/plugins/005-check_22909764.sh
 inflating: dbserver_patch_5.180228.2/plugins/000-check_dummy_perl
 inflating: dbserver_patch_5.180228.2/plugins/030-check_24625612.sh
 inflating: dbserver_patch_5.180228.2/patchmgr_functions
 inflating: dbserver_patch_5.180228.2/exadata.img.hw
 inflating: dbserver_patch_5.180228.2/libxcp.so.1
 inflating: dbserver_patch_5.180228.2/imageLogger
 inflating: dbserver_patch_5.180228.2/ExaXMLNode.pm
 inflating: dbserver_patch_5.180228.2/fwverify
```

e. Create the `dbs_group` file that contains the list of compute nodes to update. Include the nodes listed after running the `olsnodes` command in step 1 except for the driving system node. In this example, `dbs_group` should include only `node2`.

```
[root@node1 patch]# cd /root/patch/dbserver_patch_5.180228
[root@node1 dbserver_patch_5.180228]# cat dbs_group
node2
```

3. Run a patching precheck operation.

> **Note:**
>
> You must run the precheck operation with the `-nomodify_at_prereq` option to prevent any changes to the system that could impact the backup you take in the next step. Otherwise, the backup might not be able to roll back the system to its original state, should that be necessary.

```
patchmgr -dbnodes dbs_group -precheck -yum_repo <yum_repository> -
target_version <target_version> -nomodify_at_prereq
```

The output should look like the following example:

```
[root@node1 dbserver_patch_5.180228]# ./patchmgr -dbnodes dbs_group -
precheck -yum_repo  http://yum-phx.oracle.com/repo/EngineeredSystems/
exadata/dbserver/18.1.4.0.0/base/x86_64 -target_version
18.1.4.0.0.180125.3  -nomodify_at_prereq

********************************************************************************
********************************
NOTE    patchmgr release: 5.180228 (always check MOS 1553103.1 for the
latest release of dbserver.patch.zip)
NOTE
WARNING Do not interrupt the patchmgr session.
WARNING Do not resize the screen. It may disturb the screen layout.
WARNING Do not reboot database nodes during update or rollback.
WARNING Do not open logfiles in write mode and do not try to alter them.
********************************************************************************
********************************
2018-02-28 21:22:45 +0000        :Working: DO: Initiate precheck on 1
node(s)
2018-02-28 21:24:57 +0000        :Working: DO: Check free space and verify
SSH equivalence for the root user to node2
2018-02-28 21:26:15 +0000        :SUCCESS: DONE: Check free space and
verify SSH equivalence for the root user to node2
2018-02-28 21:26:47 +0000        :Working: DO: dbnodeupdate.sh running a
precheck on node(s).
2018-02-28 21:28:23 +0000        :SUCCESS: DONE: Initiate precheck on
node(s).
```

4. Back up the current system.

> **Note:**
>
> This is the proper stage to take the backup, before any modifications are made to the system.

```
patchmgr -dbnodes dbs_group -backup -yum_repo <yum_repository> -
target_version <target_version>  -allow_active_network_mounts
```

The output should look like the following example:

```
[root@node1 dbserver_patch_5.180228]#  ./patchmgr -dbnodes dbs_group -
backup  -yum_repo  http://yum-phx.oracle.com/repo/EngineeredSystems/
exadata/dbserver/18.1.4.0.0/base/x86_64 -target_version
18.1.4.0.0.180125.3 -allow_active_network_mounts

********************************************************************************
********************************
NOTE    patchmgr release: 5.180228 (always check MOS 1553103.1 for the
latest release of dbserver.patch.zip)
NOTE
WARNING Do not interrupt the patchmgr session.
WARNING Do not resize the screen. It may disturb the screen layout.
```

```
WARNING Do not reboot database nodes during update or rollback.
WARNING Do not open logfiles in write mode and do not try to alter them.
********************************************************************************
********************************
2018-02-28 21:29:00 +0000        :Working: DO: Initiate backup on 1
node(s).
2018-02-28 21:29:00 +0000        :Working: DO: Initiate backup on node(s)
2018-02-28 21:29:01 +0000        :Working: DO: Check free space and verify
SSH equivalence for the root user to node2
2018-02-28 21:30:18 +0000        :SUCCESS: DONE: Check free space and
verify SSH equivalence for the root user to node2
2018-02-28 21:30:51 +0000        :Working: DO: dbnodeupdate.sh running a
backup on node(s).
2018-02-28 21:35:50 +0000        :SUCCESS: DONE: Initiate backup on
node(s).
2018-02-28 21:35:50 +0000        :SUCCESS: DONE: Initiate backup on 1
node(s).
```

5. Remove all custom RPMs from the target compute nodes that will be updated. Custom RPMs are reported in precheck results. They include RPMs that were manually installed after the system was provisioned.
   - If you are updating the system from version 12.1.2.3.4.170111, and the precheck results include `krb5-workstation-1.10.3-57.el6.x86_64`, remove it. (This item is considered a custom RPM for this version.)
   - Do **not** remove `exadata-sun-vm-computenode-exact` or `oracle-ofed-release-guest`. These two RPMs are handled automatically during the update process.
6. Run the `nohup` command to perform the update.

```
nohup patchmgr -dbnodes dbs_group -upgrade -nobackup -yum_repo
<yum_repository> -target_version <target_version> -
allow_active_network_mounts &
```

The output should look like the following example:

```
[root@node1 dbserver_patch_5.180228]# nohup ./patchmgr -dbnodes dbs_group -
upgrade -nobackup  -yum_repo  http://yum-phx.oracle.com/repo/
EngineeredSystems/exadata/dbserver/18.1.4.0.0/base/x86_64 -target_version
18.1.4.0.0.180125.3  -allow_active_network_mounts &

********************************************************************************
********************************
NOTE    patchmgr release: 5.180228 (always check MOS 1553103.1 for the
latest release of dbserver.patch.zip)
NOTE
NOTE    Database nodes will reboot during the update process.
NOTE
WARNING Do not interrupt the patchmgr session.
WARNING Do not resize the screen. It may disturb the screen layout.
WARNING Do not reboot database nodes during update or rollback.
WARNING Do not open logfiles in write mode and do not try to alter them.
********************************************************************************
******************************

2018-02-28 21:36:26 +0000        :Working: DO: Initiate prepare steps on
```

```
node(s).
2018-02-28 21:36:26 +0000        :Working: DO: Check free space and verify
SSH equivalence for the root user to node2
2018-02-28 21:37:44 +0000        :SUCCESS: DONE: Check free space and
verify SSH equivalence for the root user to node2
2018-02-28 21:38:43 +0000        :SUCCESS: DONE: Initiate prepare steps on
node(s).
2018-02-28 21:38:43 +0000        :Working: DO: Initiate update on 1
node(s).
2018-02-28 21:38:43 +0000        :Working: DO: Initiate update on node(s)
2018-02-28 21:38:49 +0000        :Working: DO: Get information about any
required OS upgrades from node(s).
2018-02-28 21:38:59 +0000        :SUCCESS: DONE: Get information about any
required OS upgrades from node(s).
2018-02-28 21:38:59 +0000        :Working: DO: dbnodeupdate.sh running an
update step on all nodes.
2018-02-28 21:48:41 +0000        :INFO    : node2 is ready to reboot.
2018-02-28 21:48:41 +0000        :SUCCESS: DONE: dbnodeupdate.sh running
an update step on all nodes.
2018-02-28 21:48:41 +0000        :Working: DO: Initiate reboot on node(s)
2018-02-28 21:48:57 +0000        :SUCCESS: DONE: Initiate reboot on node(s)
2018-02-28 21:48:57 +0000        :Working: DO: Waiting to ensure node2 is
down before reboot.
2018-02-28 21:56:18 +0000        :Working: DO: Initiate prepare steps on
node(s).
2018-02-28 21:56:19 +0000        :Working: DO: Check free space and verify
SSH equivalence for the root user to node2
2018-02-28 21:57:37 +0000        :SUCCESS: DONE: Check free space and
verify SSH equivalence for the root user to node2
2018-02-28 21:57:42 +0000        :SEEMS ALREADY UP TO DATE: node2
2018-02-28 21:57:43 +0000        :SUCCESS: DONE: Initiate update on node(s)
```

7. After the update operation completes, verify the version of the kernel on the compute node that was updated.

```
[root@node2 ~]# imageinfo -ver
18.1.4.0.0.180125.3
```

8. f the driving system is a compute node that needs to be updated (as in this example), repeat steps 2 through 7 of this procedure using an updated compute node as the driving system to update the remaining compute node. In this example update, you would use `node2` to update `node1`.
9. On each compute node, run the `uptrack-install` command as root to install the available ksplice updates.

```
uptrack-install --all -y
```

Installing Additional Operating System Packages

Review these guidelines before you install additional operating system packages for Oracle Exadata Database Service on Dedicated Infrastructure.

You are permitted to install and update operating system packages on Oracle Exadata Database Service on Dedicated Infrastructure as long as you do not modify the kernel or

InfiniBand-specific packages. However, Oracle technical support, including installation, testing, certification and error resolution, does not apply to any non-Oracle software that you install.

Also be aware that if you add or update packages separate from an Oracle Exadata software update, then these package additions or updates can introduce problems when you apply an Oracle Exadata software update. Problems can occur because additional software packages add new dependencies that can interrupt an Oracle Exadata update. For this reason, Oracle recommends that you minimize customization.

If you install additional packages, then Oracle recommends that you have scripts to automate the removal and reinstallation of those packages. After an Oracle Exadata update, if you install additional packages, then verify that the additional packages are still compatible, and that you still need these packages.

For more information, refer to *Oracle Exadata Database Machine Maintenance Guide*.

Related Topics

- Installing, Updating and Managing Non-Oracle Software

Updating Tooling on an Exadata Cloud Infrastructure Instance

Cloud-specific tooling is used on the Exadata Cloud Infrastructure Guest VMs for local operations, including dbaascli commands.

The cloud tooling is automatically updated by Oracle when new releases are made available. If needed, you can follow the steps in *Updating Cloud Tooling Using dbaascli* to ensure you have the latest version of the cloud tooling on all virtual machines in the VM cluster

Related Topics

- Updating Cloud Tooling Using dbaascli
  To update the cloud tooling release for Oracle Exadata Database Service on Dedicated Infrastructure, complete this procedure.

Use Oracle Data Guard with Exadata Cloud Infrastructure

Learn to configure and manage Data Guard associations in your VM cluster.

- About Using Oracle Data Guard with Exadata Cloud Infrastructure
  This topic explains how to use the Console or the API to manage Data Guard associations in your VM cluster.
- Prerequisites for Using Oracle Data Guard with Exadata Cloud Infrastructure
  An Exadata Cloud Infrastructure Oracle Data Guard implementation requires two existing Exadata VM Clusters: one containing an existing database that is to be duplicated by Data Guard, and one that will house the new standby database by Data Guard.
- Working with Data Guard
  Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data.
- Using the Console to Manage Oracle Data Guard Associations
  Learn how to enable a Data Guard association between databases, change the role of a database in a Data Guard association using either a switchover or a failover operation, and reinstate a failed database.
- Using the API to manage Data Guard associations
  Use these API operations to manage Data Guard associations on an Exadata Cloud Infrastructure instance:

About Using Oracle Data Guard with Exadata Cloud Infrastructure

This topic explains how to use the Console or the API to manage Data Guard associations in your VM cluster.

When you use the Console or the API to enable Data Guard for an Exadata database compute node database:

- The standby database is a physical standby.
- The versions of peer databases (primary and standby) are identical.
- You are limited to one standby database for each primary database.
- The standby database is deployed as an open, read-only database (Active Data Guard).

To configure a Data Guard system between on-premises and Exadata database compute nodes, or to configure your database with multiple standbys, you must access the database host directly and set up Data Guard manually.

For complete information on Oracle Data Guard, see the *Data Guard Concepts and Administration* documentation on the *Oracle Document Portal*.

Related Topics

- Data Guard Concepts and Administration
- Oracle Document Portal

Prerequisites for Using Oracle Data Guard with Exadata Cloud Infrastructure

An Exadata Cloud Infrastructure Oracle Data Guard implementation requires two existing Exadata VM Clusters: one containing an existing database that is to be duplicated by Data Guard, and one that will house the new standby database by Data Guard.

> **Note:**
>
> Oracle strongly recommends the primary and standby databases for any production workloads be on different Exadata Cloud Infrastructures for better fault isolation and disaster protection. If you are adding a new standby in the same region with multiple availability domains, Oracle recommends choosing a separate availability domain for complete availability domain or data center fault isolation. If you are adding a new standby across regions, the standby will have fault isolation for a regional failure as well.

When enabling Data Guard, you can create a new Database Home on the standby Exadata instance to house the new standby database during the enable Data Guard operation. Alternately, you can choose to provision the standby database in an existing Database Home on the standby instance. For information on creating the required resources for the standby system, see the following topics:

- To create a Cloud Exadata infrastructure resource
- To create a cloud VM cluster resource
- To create a new Database Home in an existing Exadata Cloud Infrastructure instance

You can use a custom database software image to that contains the necessary patches for your databases when creating a Database Home on either the primary or the standby Exadata instance. See Oracle Database Software Images for information on working with custom Oracle Database software images.

If you choose to provision a standby database in an existing Database Home, ensure that the target Database Home on the standby instance has all required patches that are in use for the primary database before you provision the standby database. See the following topics for more information on patching an existing Database Home:

- To patch the Oracle Database software in a Database Home (cloud VM cluster)
- To patch the Oracle Database software in a Database Home (DB system)

If you are creating a Data Guard Association and you are using customer managed keys to encrypt the database, you must have configured the Vault Service and created a master key. See *To administer Vault encryption keys* and *Key and Secret Management Concepts*.

- Network Requirements for Data Guard
  Describe the network requirements for using Exadata Cloud Infrastructure with Oracle Data Guard.
- Password Requirements
  For Data Guard operations to work, the `SYS` password and the TDE wallet password of the primary and standby databases must all be the same.
- Known Issues for Exadata Cloud Infrastructure and Data Guard
  Possible TDE key replication issue, and MRP and DG LCM operation failures.
- Adding a Node to a VM Cluster
- Removing a Node from a VM Cluster

Related Topics

- Customer-Managed Keys in Exadata Cloud Infrastructure
  Customer-managed keys for Exadata Cloud Infrastructure is a feature of Oracle Cloud Infrastructure (OCI) Vault service that enables you to encrypt your data using encryption keys that you control.
- To administer Vault encryption keys
  Use this procedure to rotate the Vault encryption key or or change the encryption management configuration.
- Key and Secret Management Concepts

Network Requirements for Data Guard

Describe the network requirements for using Exadata Cloud Infrastructure with Oracle Data Guard.

Ensure that your environment meets the following network requirements:

- The primary and standby databases can be part of VM clusters in different compartments.
- If you want to configure Oracle Data Guard across regions, then you must configure remote virtual cloud network (VCN) peering between the primary and standby databases. Networking is configured on the cloud VM cluster resource for systems using the *The new Exadata Resource Model*, and on the DB system resource for system using the old resource model. See *Remote VCN Peering using an RPC* .
  For Exadata Data Guard configurations, OCI supports the use of hub-and-spoke network topology for the VCNs within each region. This means that the primary and standby databases can each utilize a "spoke" VCN that passes network traffic to the "hub" VCN

that has a remote peering connection. See *Transit Routing inside a hub VCN* for information on setting up this network topology.

- To set up Oracle Data Guard within a single region, both Exadata Cloud Infrastructure instances must use the same VCN. When setting up Data Guard within the same region, Oracle recommends that the instance containing the standby database be in a different **availability domain** from the instance containing the primary database to improve availability and disaster recovery.
- Configure the ingress and egress security rules for the subnets of both Exadata Cloud Infrastructure instances in the Oracle Data Guard association to enable TCP traffic to move between the applicable ports. Ensure that the rules you create are stateful (the default).

  For example, if the subnet of the primary Exadata Cloud Infrastructure instance uses the source CIDR 10.0.0.0/24 and the subnet of the standby instance uses the source CIDR 10.0.1.0/24, then create rules as shown in the subsequent example.

> **Note:**
>
> The egress rules in the example show how to enable TCP traffic only for port 1521, which is a minimum requirement for Oracle Data Guard to work. If TCP traffic is already enabled for all destinations (0.0.0.0/0) on all of your outgoing ports, then you need not explicitly add these specific egress rules.

Security Rules for Subnet of Primary Exadata Cloud Infrastructure instance

Ingress Rules:

```
Stateless: No
Source: 10.0.1.0/24
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 1521
Allows: TCP traffic for ports: 1521
```

Egress Rules:

```
Stateless: No
Destination: 10.0.1.0/24
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 1521
Allows: TCP traffic for ports: 1521
```

Security Rules for Subnet of Standby Exadata Cloud Infrastructure instance

Ingress Rules:

```
Stateless: No
Source: 10.0.0.0/24
IP Protocol: TCP
Source Port Range: All
```

```
Destination Port Range: 1521
Allows: TCP traffic for ports: 1521

Egress Rules:

Stateless: No
Destination: 10.0.0.0/24
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 1521
Allows: TCP traffic for ports: 1521
```

For information about creating and editing rules, see *Security Lists* .

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.
- Remote VCN Peering using an RPC
- Transit Routing inside a hub VCN
- Security Lists

Password Requirements

For Data Guard operations to work, the `SYS` password and the TDE wallet password of the primary and standby databases must all be the same.

If you change any one of these passwords, you must update the rest of the passwords to match. See *Changing the Database Passwords* to learn how to change the SYS password or the TDE wallet password.

If you make any change to the TDE wallet (such as adding a master key for a new PDB or changing the wallet password), you must copy the wallet from the primary to the standby so that Data Guard can continue to operate. For Oracle Database versions earlier than 12.2, if you change the `SYS` password on one of the peers, you need to manually sync the password file between the DB systems.

Related Topics

- Changing the Database Passwords
  To change the SYS password, or to change the TDE wallet password, use this procedure.

Known Issues for Exadata Cloud Infrastructure and Data Guard

Possible TDE key replication issue, and MRP and DG LCM operation failures.

KMS RPM `libkmstdepkcs11_1.286-1.286-1-Linux.rpm` is the latest available which supports active replication of key between cross-region KMS vaults (source and target), and it is recommended to upgrade the RPM on clusters participating in Data Guard. OCI Vault cross-region Data Guard works with a lower version of RPM, but the older version does not guarantee active replication of keys. If the TDE keys have any replication issue between vaults, Data Guard replication might have an impact (MRP fails on standby cluster due to missing key on target vault) and MRP could resume only after the keys are replicated to the

target vault. To avoid MRP and DG LCM operation failures, upgrade the `libkms` RPM on both the clusters, and restart the databases (only databases using customer-managed keys).

Adding a Node to a VM Cluster

When adding a node to a VM cluster, an instance of the Data Guard database is automatically created on the new node. However, metadata updation on the remote database, that is, the primary database if addition is done on the standby database and vice versa, must be done manually.

This can be done by copying over the `addinstance` JSON file, `/var/opt/oracle/dbaas_acfs/<dbname>/addInstance.json` created at the end of instance addition and running the `/var/opt/oracle/ocde/rops update_instance <dbname> <path to addInstance JSON>` command on any node of the remote cluster.

Related Topics

- To add compute and storage resources to a flexible cloud Exadata infrastructure resource
  This task describes how to use the Oracle Cloud Infrastructure Console to scale a flexible cloud Exadata infrastructure resource.

Removing a Node from a VM Cluster

When removing a node from a VM cluster, the instance and it's metadata on the removing node is deleted automatically. However, deletion of the corresponding metadata on the remote database, that is, the primary database if removal is done on the standby database and vice versa, must be done manually.

This can be done by running the `/var/opt/oracle/ocde/rops remove_instance <dbname> <Instance Name>` command on any node of the remote cluster.

Related Topics

- To add compute and storage resources to a flexible cloud Exadata infrastructure resource
  This task describes how to use the Oracle Cloud Infrastructure Console to scale a flexible cloud Exadata infrastructure resource.

Working with Data Guard

Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data.

The Data Guard implementation requires two databases, one in a primary role and one in a standby role. The two databases compose a Data Guard association. Most of your applications access the primary database. The standby database is a transactionally consistent copy of the primary database.

Data Guard maintains the standby database by transmitting and applying redo data from the primary database. If the primary database becomes unavailable, you can use Data Guard to switch or fail over the standby database to the primary role.

- Switchover
  A switchover reverses the primary and standby database roles.
- Failover
  A failover transitions the standby database into the primary role after the existing primary database fails or becomes unreachable.

- Reinstate
  Reinstates a database into the standby role in a Data Guard association.

Switchover

A switchover reverses the primary and standby database roles.

Each database continues to participate in the Data Guard association in its new role. A switchover ensures no data loss. You can use a switchover before you perform planned maintenance on the primary database. Performing planned maintenance on a Exadata database virtual machine with a Data Guard association is typically done by switching the primary to the standby role, performing maintenance on the standby, and then switching it back to the primary role.

Failover

A failover transitions the standby database into the primary role after the existing primary database fails or becomes unreachable.

A failover might result in some data loss when you use **Maximum Performance** protection mode.

Reinstate

Reinstates a database into the standby role in a Data Guard association.

You can use the reinstate command to return a failed database into service after correcting the cause of failure.

> **Note:**
>
> You can't terminate a primary database that has a Data Guard association with a peer (standby) database. Delete the standby database first. Alternatively, you can switch over the primary database to the standby role, and then terminate the former primary.
>
> You can't terminate a VM cluster that includes Data Guard enabled databases. You must first remove the Data Guard association by terminating the standby database.

Using the Console to Manage Oracle Data Guard Associations

Learn how to enable a Data Guard association between databases, change the role of a database in a Data Guard association using either a switchover or a failover operation, and reinstate a failed database.

When you enable Data Guard, a separate Data Guard association is created for the primary and the standby database.

- Using the Console to Enable Data Guard on an Exadata Cloud Infrastructure System
  Learn to enable Data Guard association between databases.
- To view Data Guard associations of databases in a Cloud VM Cluster
  To view the role of each database in a Data Guard association in an Cloud VM Cluster, follow this procedure.

- To enable automatic backups on a standby database
  Learn to enable automatic backups on a standby database.
- To perform a database switchover
  You initiate a switchover operation by using the Data Guard association of the primary database.
- To edit the Oracle Data Guard association
- To perform a database failover
  You initiate a failover operation by using the Data Guard association of the standby database.
- To reinstate a database
  After you fail over a primary database to its standby, the standby assumes the primary role and the old primary is identified as a disabled standby. After you correct the cause of failure, you can reinstate the failed database as a functioning standby for the current primary by using its Data Guard association.
- To terminate a Data Guard association on an Exadata Cloud Infrastructure instance
  On an Exadata Cloud Infrastructure instance, you remove a Data Guard association by terminating the standby database.

Using the Console to Enable Data Guard on an Exadata Cloud Infrastructure System

Learn to enable Data Guard association between databases.

> **Note:**
>
> - When you enable Data Guard, replication of data happens only over the client network.
> - When you configure the Data Guard association for 23ai databases, the primary and standby databases must be on the same major release version while the standby database can be on a higher minor version.

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment** that contains the Exadata Cloud Infrastructure instance with the database for which you want to enable Oracle Data Guard..
3. Navigate to the cloud VM cluster or DB system that contains a database you want to assume the primary role:
   - *Cloud VM clusters (* new resource model*):* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
   - *DB systems:* Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. On the VM cluster or DB system details page, in the **Databases** section, click the name of the database you want to make primary.
5. On the Database Details page, under **Resources**, click **Data Guard Associations**.

6. In the **Data Guard Associations** section, click **Enable Data Guard**.
7. On the Enable Data Guard page, configure your Data Guard association.
   - In the **Select peer DB system** section, provide the following information for the standby database to obtain a list of available Exadata systems in which to locate the standby database:
     - **Region**: Select a region where you want to locate the standby database. The region where the primary database is located is selected, by default. You can choose to locate the standby database in a different region. The hint text associated with this field tells you in which region the primary database is located.
     - **Availability domain**: Select an availability domain for the standby database. The hint text associated with this field tells you in which availability domain the primary database is located.
     - **Shape**: Select the shape of the standby Exadata system.
     - **Data Guard peer resource type**: Select **DB System** or **VM Cluster**.
     - Select a DB system or cloud VM cluster from the drop-down list.
   - **Data Guard association details:**
     - **Data Guard Type:** Select Active Data Guard or Data Guard. Active Data Guard provides additional features including: Real-Time Query and DML Offload, Automatic Block Repair, Standby Block Change Tracking, Far Sync, Global Data Services, and Application Continuity. Note that Active Data Guard requires an Oracle Active Data Guard license. For more information on Active Data Guard, see Active Data Guard. For a complete overview of both Data Guard types, see Introduction to Oracle Data Guard
     - **Protection mode**: The protection mode can be **Maximum Performance** or **Maximum Availability**. See Oracle Data Duard Protection Modes for information on these options.
     - **Transport type:** The redo transport type used for this Data Guard association. See Redo Transport Servicesfor information on these options.
   - In the **Choose Database Home** section, choose one of the following:
     - **Select an existing Database Home:** If you use this option, select a home from the Database Home display name drop-down list.
     - **Create a new Database Home:** If you choose this option, enter a name for the new Database Home in the **Database Home display name** field. Click **Change Database Image** to select a database software image for the new Database Home. In the **Select a Database Software Image** panel, do the following:
       a. Select the compartment containing the database software image you want to use to create the new Database Home.
       b. Select the region containing the database software image you want to use to create the new Database Home. Region filter defaults to the currently connected region and lists all the software images created in that region. When you choose a different region, the software image list is refreshed to display the software images created in the selected region.
       c. Select the Oracle Database software version that the new Database Home will use, then choose an image from the list of available images for your selected software version.
       d. Click **Select**.

> **Note:**
>
> Oracle recommends applying the same list of patches to the Database Homes of the primary and standby databases.

- In the **Configure standby database:** section, provide standby database details.

  > **Note:**
  >
  > You cannot modify the `db_unique_name` and SID prefix after creating the database.

  - **Database unique name:** Optionally, specify a value for the `DB_UNIQUE_NAME` database parameter. This value must be unique across the primary and standby cloud VM clusters. The unique name must meet the requirements:
    * Maximum of 30 characters
    * Contain only alphanumeric or underscore (_) characters
    * Begin with an alphabetic character
    * Unique across the VM cluster. Recommended to be unique across the tenancy.

    If not specified, the system automatically generates a unique name value, as follows:

    <db_name>_<3_chars_unique_string>_<region-name>

  - **Database password:** Enter the database administrator password of the primary database. Use this same database administrator password for the standby database.

    > **Note:**
    >
    > The administrator password and the TDE wallet password must be identical. If the passwords are not identical, then follow the instructions in Changing the Database Passwords to ensure that they are.

8. Click **Show Advanced Options** to specify advanced options for the standby database:

   - **Management:**

     Oracle SID prefix: The Oracle Database instance number is automatically added to the SID prefix to create the `INSTANCE_NAME` database parameter. The `INSTANCE_NAME` parameter is also known as the SID. If not provided, then the SID prefix defaults to the first 12 characters of the `db_unique_name`.

     > **Note:**
     >
     > Entering an SID prefix is only available for Oracle 12.1 databases and above.

The SID prefix must meet the requirements:

- Maximum of 12 characters
- Contain only alphanumeric characters
- Begin with an alphabetic character
- Unique in the VM cluster and across primary and standby databases

9. Click **Enable Data Guard**. When you create the association, the details for a database and its peer display their respective roles as **Primary** or **Standby**.

A work request is issued to configure the Data Guard association. The progress of the request and the stages of provisoning can be viewed on the **Work Requests** page.

When the association is created, the details for a database and its peer display their respective roles as **Primary** or **Standby**.

- View Data Guard Provisioning Progress
  View the progress of Data Guard Provisioning tasks using the Work Requests page.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.
- Network Setup for Exadata Cloud Infrastructure Instances
  This topic describes the recommended configuration for the VCN and several related requirements for the Exadata Cloud Infrastructure instance.
- Changing the Database Passwords
  To change the SYS password, or to change the TDE wallet password, use this procedure.

View Data Guard Provisioning Progress

View the progress of Data Guard Provisioning tasks using the Work Requests page.

After you have completed the task To Enable Data Guard, multiple work requests are issued to complete the provisioning of the Data Guard association. To veiw the progress of these work requests:

1. Navigate to the **Work Requests Details** page. On the **Work Requests Details** page there is a bar in the Work Request Information tab that shows the overall progress of the Data Guard Provisioning
2. Under Resources, select **Log Messages**. The table shows a messsage for each task that is completed or in progress.

To view Data Guard associations of databases in a Cloud VM Cluster

To view the role of each database in a Data Guard association in an Cloud VM Cluster, follow this procedure.

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the cloud VM cluster that contains the databases you wish to view their roles in Data Guard associations.

4. In the **Databases** section under **Resources**, the role of each database in this VM Cluster is indicated in the **Data Guard role** column.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

To enable automatic backups on a standby database

Learn to enable automatic backups on a standby database.

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment that contains the Exadata Cloud Infrastructure instance with the database for which you want to enable automatic database.
3. Navigate to the cloud VM cluster or DB system that contains the primary database.
   - *Cloud VM clusters (* new resource model*):* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
   - *DB systems:* Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. On the VM cluster or DB system details page, in the **Databases** section, click the name of the primary database.
5. On the Database Details page, under **Resources**, click **Data Guard Associations**.
6. Click the name of the standby database for which you want to enable automatic backups.

   The system displays a banner if automatic backups are not enabled for this database.
7. Click **Enable automatic backups** on the banner.
8. On the resulting Configure Automatic Backups window, enter the following details:
   - **Enable automatic backup:** Check the check box to enable or disable automatic incremental backups for this database. If your database is in a security zone compartment, you must enable automatic backups.
   - **Backup Scheduling:**
     - **Full backup scheduling day:** Choose a day of the week for the initial and future L0 backups to start.
     - **Full backup scheduling time (UTC):** Specify the time window when the full backups start when the automatic backup capability is selected.
     - **Take the first backup immediately:** A full database backup includes all datafiles, control file, and parameter files associated with the target database. Archive backups are separate and decoupled and executed every 30 minutes. You can choose to execute the first full backup immediately or defer to the assigned full backup scheduling time. If you defer to the latter, the database will not be recoverable until the first backup completes.
   - **Backup Destination:** Object Storage is selected by default and you cannot change it.

> **Note:**
>
> - If automatic backup is enabled on the primary database and the backup destination is Autonomous Recovery Service, then you cannot enable backup on the standby database.
> - If automatic backup is enabled on the primary database and the backup destination is Object Storage, then you can enable backup on the standby database. Note that you can only select Object Storage as the backup destination.
> - If automatic backup is disabled on the primary database, you can still enable backup on the standby database by selecting Object Storage as the backup destination.

9. Click **Save Changes**.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

To perform a database switchover

You initiate a switchover operation by using the Data Guard association of the primary database.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose the **Compartment** that contains the Exadata Cloud Infrastructure instance with the database for which you want to enable Oracle Data Guard.
3. Navigate to the cloud VM cluster or DB system that contains the Data Guard association:

   Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems: Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. Under **Resources**, click **Data Guard Associations**.
5. For the Data Guard association on which you want to perform a switchover, click the Actions icon (three dots), and then click **Switchover**.
6. In the **Switchover Database** dialog box, enter the database admin password, and then click **OK**.

   This database should now assume the role of the standby, and the standby should assume the role of the primary in the Data Guard association.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

To edit the Oracle Data Guard association

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose the **Compartment** that contains the Exadata Cloud Service instance with the database for which you want to enable Oracle Data Guard.
3. Navigate to the cloud VM cluster or DB system that contains the Data Guard association:

   Cloud VM clusters (new resource model*):* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems: Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. Under **Resources**, click **Data Guard Associations**.
5. For the Data Guard association you want to manage, click the Actions icon (three dots), and then click **Edit Protection Mode**.
6. In the **Edit Data Guard Association** panel, configure the Data Guard association:
   - **Data Guard Type**: Select Active Data Guard or Data Guard. Active Data Guard provides additional features including: Real-Time Query and DML Offload, Automatic Block Repair, Standby Block Change Tracking, Far Sync, Global Data Services, and Application Continuity. Note that Active Data Guard requires an Oracle Active Data Guard license. For more information on Active Data Guard, see Active Data Guard. For a complete overview of both Data Guard types, see Introduction to Oracle Data Guard
   - **Protection mode**: The protection mode can be **Maximum Performance** or **Maximum Availability**. See *Oracle Data Guard Protection Modes* for information on these options.
   - **Transport type**: The redo transport type used for this Oracle Data Guard association.
   - **Database admin password**: Enter the ADMIN password for the database.
7. Click **Save**.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.
- Oracle Data Guard Protection Modes

To perform a database failover

You initiate a failover operation by using the Data Guard association of the standby database.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose the **Compartment** that contains the Exadata Cloud Infrastructure instance with the database for which you want to enable Oracle Data Guard.
3. Navigate to the cloud VM cluster or DB system that contains the Data Guard association:

Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

DB systems: Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

4. Under **Resources**, click **Data Guard Associations**.
5. For the Data Guard association on which you want to perform a failover, click **Failover**.
6. In the **Failover Database** dialog box, enter the database admin password, and then click **OK**.

   This database should now assume the role of the primary, and the old primary's role should display as **Disabled Standby**.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

To reinstate a database

After you fail over a primary database to its standby, the standby assumes the primary role and the old primary is identified as a disabled standby. After you correct the cause of failure, you can reinstate the failed database as a functioning standby for the current primary by using its Data Guard association.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose the **Compartment** that contains the Exadata Cloud Infrastructure instance with the database for which you want to enable Oracle Data Guard.
3. Navigate to the cloud VM cluster or DB system that contains the Data Guard association:

   *Cloud VM clusters (*new resource model*):* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems: Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. Under **Resources**, click **Data Guard Associations**.
5. For the Data Guard association on which you want to reinstate this database, click the Actions icon (three dots), and then click **Reinstate**.
6. In the **Reinstate Database** dialog box, enter the database admin password, and then click **OK**.

   This database should now be reinstated as the standby in the Data Guard association.

To terminate a Data Guard association on an Exadata Cloud Infrastructure instance

On an Exadata Cloud Infrastructure instance, you remove a Data Guard association by terminating the standby database.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Compartment** that contains the Exadata Cloud Infrastructure instance with the database for which you want to enable Oracle Data Guard.
3. Navigate to the cloud VM cluster or DB system that contains the standby database:

   Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   DB systems: Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
4. For the standby database you want to terminate, click the Actions icon (three dots), and then click **Terminate**.
5. In the **Terminate Database** dialog box, enter the name of the database, and then click **OK**.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

Using the API to manage Data Guard associations

Use these API operations to manage Data Guard associations on an Exadata Cloud Infrastructure instance:

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.

- CreateDataGuardAssociation
- ListDataGuardAssociations
- GetDataGuardAssociation
- UpdateDataGuardAssociation
- SwitchoverDataGuardAssociation
- FailoverDataGuardAssociation
- ReinstateDataGuardAssociation
- DeleteDatabase - To terminate an Exadata Cloud Infrastructure instance Data Guard association, you delete the standby database.

For the complete list of APIs for the Database service, see Database Service API.

Configure Oracle Database Features for Exadata Cloud Infrastructure

This topic describes how to configure Oracle Multitenant, tablespace encryption, and Huge Pages for use with your Exadata Cloud Infrastructure instance.

- Using Oracle Multitenant on an Exadata Cloud Infrastructure Instance

- Managing Tablespace Encryption
- Managing Huge Pages

Using Oracle Multitenant on an Exadata Cloud Infrastructure Instance

When you create an Exadata Cloud Infrastructure Instance that uses Oracle Database 12c or later, an Oracle Multitenant environment is created.

The multitenant architecture enables an Oracle database to function as a multitenant container database (CDB) that includes zero, one, or many pluggable databases (PDBs). A PDB is a portable collection of schemas, schema objects, and non-schema objects that appears to an Oracle Net Services client as a non-CDB. All Oracle databases using versions earlier than Oracle Database 12c are non-CDBs.

To use Oracle Transparent Data Encryption (TDE) in a pluggable database (PDB), you must create and activate a master encryption key for the PDB.

In a multitenant environment, each PDB has its own master encryption key which is stored in a single keystore used by all containers.

You must export and import the master encryption key for any encrypted PDBs you plug into your Exadata Cloud Infrastructure Instance CDB.

If your source PDB is encrypted, you must export the master encryption key and then import it.

You can export and import all of the TDE master encryption keys that belong to the PDB by exporting and importing the TDE master encryption keys from within a PDB. Export and import of TDE master encryption keys support the PDB unplug and plug operations. During a PDB unplug and plug, all of the TDE master encryption keys that belong to a PDB, as well as the metadata, are involved.

See "Exporting and Importing TDE Master Encryption Keys for a PDB" in *Oracle Database Advanced Security Guide* for Release 19, 18, 12.2 or 12.1.

See "ADMINISTER KEY MANAGEMENT" in *Oracle Database SQL Language Reference* for Release 19, 18, 12.2 or 12.1.

- To determine if you need to create and activate an encryption key for the PDB
- To create and activate the master encryption key in a PDB
- To export and import a master encryption key

To determine if you need to create and activate an encryption key for the PDB

1. Invoke SQL*Plus and log in to the database as the `SYS` user with `SYSDBA` privileges.
2. Set the container to the PDB:

   ```
   SQL> ALTER SESSION SET CONTAINER = pdb;
   ```

3. Query `V$ENCRYPTION_WALLET` as follows:

   ```
   SQL> SELECT wrl_parameter, status, wallet_type FROM v$encryption_wallet;
   ```

   If the `STATUS` column contains a value of `OPEN_NO_MASTER_KEY`, you need to create and activate the master encryption key.

To create and activate the master encryption key in a PDB

1. Set the container to the PDB:

```
SQL> ALTER SESSION SET CONTAINER = pdb;
```

2. Create and activate a master encryption key in the PDB by executing the following command:

```
SQL> ADMINISTER KEY MANAGEMENT SET KEY USING TAG 'tag' FORCE KEYSTORE
IDENTIFIED BY keystore-password WITH BACKUP USING 'backup_identifier';
```

In the previous command:

- `keystore-password` is the keystore password. By default, the keystore password is set to the value of the administration password that is specified when the database is created.
- The optional `USING TAG 'tag'` clause can be used to associate a tag with the new master encryption key.
- The `WITH BACKUP` clause, and the optional `USING 'backup_identifier'` clause, can be used to create a backup of the keystore before the new master encryption key is created.

See also `ADMINISTER KEY MANAGEMENT` in *Oracle Database SQL Language Reference for Release*19, 18 or 12.2.

> **Note:**
>
> To enable key management operations while the keystore is in use, Oracle Database 12c Release 2, and later, includes the `FORCE KEYSTORE` option to the `ADMINISTER KEY MANAGEMENT` command. This option is also available for Oracle Database 12c Release 1 with the October 2017, or later, bundle patch.
>
> If your Oracle Database 12c Release 1 database does not have the October 2017, or later, bundle patch installed, you can perform the following alternative steps:
>
> a. Close the keystore.
> b. Open the password-based keystore.
> c. Create and activate a master encryption key in the PDB by using `ADMINISTER KEY MANAGEMENT` without the `FORCE KEYSTORE` option.
> d. Update the auto-login keystore by using `ADMINISTER KEY MANAGEMENT` with the `CREATE AUTO_LOGIN KEYSTORE FROM KEYSTORE` option.

3. Query `V$ENCRYPTION_WALLET` again to verify that the `STATUS` column is set to `OPEN`:

```
SQL> SELECT wrl_parameter, status, wallet_type FROM v$encryption_wallet;
```

4. Query `V$INSTANCE` and take note of the value in the `HOST_NAME` column, which identifies the database server that contains the newly updated keystore files:

```
SQL> SELECT host_name FROM v$instance;
```

5. Copy the updated keystore files to all of the other database servers.

   To distribute the updated keystore, you must perform the following actions on each database server that does not contain the updated keystore files:

   a. Connect to the root container and query `V$ENCRYPTION_WALLET`. Take note of the keystore location contained in the `WRL_PARAMETER` column:

```
SQL> SELECT wrl_parameter, status FROM v$encryption_wallet;
```

   b. Copy the updated keystore files.

      You must copy all of the updated keystore files from a database server that is already updated. Use the keystore location observed in the `WRL_PARAMETER` column of `V$ENCRYPTION_WALLET`.

   Open the updated keystore:

```
SQL> ADMINISTER KEY MANAGEMENT SET KEYSTORE open FORCE KEYSTORE IDENTIFIED
BY keystore-password CONTAINER=all;
```

> **Note:**
>
> To enable key management operations while the keystore is in use, Oracle Database 12c Release 2, and later, includes the `FORCE KEYSTORE` option to the `ADMINISTER KEY MANAGEMENT` command. This option is also available for Oracle Database 12c Release 1 with the October 2017, or later, bundle patch.
>
> If your Oracle Database 12c Release 1 database does not have the October 2017, or later, bundle patch installed, you can perform the following alternative steps:
>
> a. Close the keystore before copying the updated keystore files.
>
> b. Copy the updated keystore files.
>
> c. Open the updated keystore by using `ADMINISTER KEY MANAGEMENT` without the `FORCE KEYSTORE` option.

6. Query `GV$ENCRYPTION_WALLET` to verify that the `STATUS` column is set to `OPEN` across all of the database instances:

```
SQL> SELECT wrl_parameter, status, wallet_type FROM gv$encryption_wallet;
```

To export and import a master encryption key

1. Export the master encryption key.

   a. Invoke SQL*Plus and log in to the PDB.

b. Execute the following command:

```
SQL> ADMINISTER KEY MANAGEMENT EXPORT ENCRYPTION KEYS WITH SECRET
"secret" TO 'filename' IDENTIFIED BY keystore-password;
```

2. Import the master encryption key.
   a. Invoke SQL*Plus and log in to the PDB.
   b. Execute the following command:

```
SQL> ADMINISTER KEY MANAGEMENT IMPORT ENCRYPTION KEYS WITH SECRET
"secret" FROM 'filename' IDENTIFIED BY keystore-password;
```

Managing Tablespace Encryption

By default, all new tablespaces that you create in an Exadata database are encrypted.

However, the tablespaces that are initially created when the database is created may not be encrypted by default.

- For databases that use Oracle Database 12c Release 2 or later, only the `USERS` tablespaces initially created when the database was created are encrypted. No other tablespaces are encrypted including the non-`USERS` tablespaces in:
  - The root container (`CDB$ROOT`).
  - The seed pluggable database (`PDB$SEED`).
  - The first PDB, which is created when the database is created.
- For databases that use Oracle Database 12c Release 1 or Oracle Database 11g, none of the tablespaces initially created when the database was created are encrypted.

For further information about the implementation of tablespace encryption in Exadata, along with how it impacts various deployment scenarios, see Oracle Database Tablespace Encryption Behavior in Oracle Cloud.

Creating Encrypted Tablespaces

User-created tablespaces are encrypted by default.

By default, any new tablespaces created by using the `SQL CREATE TABLESPACE` command are encrypted with the AES128 encryption algorithm. You do not need to include the `USING 'encrypt_algorithm'` clause to use the default encryption.

You can specify another supported algorithm by including the USING 'encrypt_algorithm' clause in the CREATE TABLESPACE command. Supported algorithms are AES256, AES192, AES128, and 3DES168.

Managing Tablespace Encryption

You can manage the software keystore (known as an Oracle wallet in Oracle Database 11g), the master encryption key, and control whether encryption is enabled by default.

Managing the Master Encryption Key

Tablespace encryption uses a two-tiered, key-based architecture to transparently encrypt (and decrypt) tablespaces. The master encryption key is stored in an external security module

(software keystore). This master encryption key is used to encrypt the tablespace encryption key, which in turn is used to encrypt and decrypt data in the tablespace.

When a database is created on an Exadata Cloud Service instance, a local software keystore is created. The keystore is local to the compute nodes and is protected by the administration password specified during the database creation process. The auto-login software keystore is automatically opened when the database is started.

You can change (rotate) the master encryption key by using the `ADMINISTER KEY MANAGEMENT SQL` statement. For example:

```
SQL> ADMINISTER KEY MANAGEMENT SET ENCRYPTION KEY USING TAG 'tag'
IDENTIFIED BY password WITH BACKUP USING 'backup';


keystore altered.
```

See "Managing the TDE Master Encryption Key" in *Oracle Database Advanced Security Guide* for Release 19, 18, 12.2 or 12.1 or "Setting and Resetting the Master Encryption Key" in *Oracle Database Advanced Security Administrator's Guide* for Release 11.2.

Controlling Default Tablespace Encryption

The `ENCRYPT_NEW_TABLESPACES` initialization parameter controls the default encryption of new tablespaces. In Exadata databases, this parameter is set to `CLOUD_ONLY` by default.

Values of this parameter are as follows.

Value	Description
`ALWAYS`	During creation, tablespaces are transparently encrypted with the AES128 algorithm unless a different algorithm is specified in the `ENCRYPTION` clause.
`CLOUD_ONLY`	Tablespaces created in an Exadata database are transparently encrypted with the AES128 algorithm unless a different algorithm is specified in the `ENCRYPTION` clause. For non-cloud databases, tablespaces are only encrypted if the `ENCRYPTION` clause is specified. `ENCRYPTION` is the default value.
`DDL`	During creation, tablespaces are not transparently encrypted by default, and are only encrypted if the `ENCRYPTION` clause is specified.

Note:

With Oracle Database 12c Release 2 (12.2), or later, you can no longer create an unencrypted tablespace in an Exadata database. An error message is returned if you set `ENCRYPT_NEW_TABLESPACES` to `DDL` and issue a `CREATE TABLESPACE` command without specifying an `ENCRYPTION` clause.

Managing Huge Pages

Huge Pages provide considerable performance benefits for Oracle Database on systems with large amounts of memory. Oracle Database on an Exadata Cloud Infrastructure instance provides configuration settings that make use of Huge Pages by default; however, you can make manual adjustments to optimize the configuration of Huge Pages.

Huge Pages is a feature integrated into the Linux kernel 2.6. Enabling Huge Pages makes it possible for the operating system to support large memory pages. Using Huge Pages can improve system performance by reducing the amount of system CPU and memory resources required to manage Linux page tables, which store the mapping between virtual and physical memory addresses. For Oracle Databases, using Huge Pages can drastically reduce the number of page table entries associated with the System Global Area (SGA).

On Exadata Cloud Infrastructure instances, a standard page is 4 KB, while a Huge Page is 2 MB by default. Therefore, an Oracle Database on an Exadata DB system with a 50 GB SGA requires 13,107,200 standard pages to house the SGA, compared with only 25,600 Huge Pages. The result is much smaller page tables, which require less memory to store and fewer CPU resources to access and manage.

Adjusting the Configuration of Huge Pages

The configuration of Huge Pages for Oracle Database is a two-step process:

- At the operating system level, the overall amount of memory allocated to Huge Pages is controlled by the vm.nr_hugepages entry in the /etc/sysctl.conf file. This setting is made on each compute node in the environment and it is strongly recommended that the setting is consistent across all of the compute nodes. To alter the Huge Page allocation, you can execute the following command on each compute node as the root user:

  ```
  # sysctl -w vm.nr_hugepages=value
  ```

  where `value` is the number of Huge Pages that you want to allocate.

  On Exadata Cloud Infrastructure instances, each Huge Page is 2 MB by default. Therefore, to allocate 50 GB of memory to Huge Pages you can execute the following command:

  ```
  # sysctl -w vm.nr_hugepages=25600
  ```

- At the Oracle Database level, the use of Huge Pages is controlled by the `USE_LARGE_PAGES` instance parameter setting. This setting applies to each database instance in a clustered database. Oracle strongly recommends a consistent setting across all of the database instances associated with a database. The following options are available:
  - `TRUE` — specifies that the database instance can use Huge Pages if they are available. For all versions of Oracle Database after 11.2.0.3, Oracle allocates as much of the SGA as it can, using Huge Pages. When the Huge Page allocation is exhausted, standard memory pages are used.
  - `FALSE` — specifies that the database instance does not use Huge Pages. This setting is generally not recommended if Huge Pages are available.
  - `ONLY` — specifies that the database instance must use Huge Pages. With this setting, the database instance fails to start if the entire SGA cannot be accommodated in Huge Pages.

If you make any adjustments at either the operating system or Oracle Database level, ensure that the overall configuration works.

For more information, see the *Oracle Database Administrator's Reference for Linux and UNIX-Based Operating Systems* for Release 19, 18, 12.1, or 11.2 for a general overview of Huge Pages and more information about configuring Huge Pages. Also, see `USE_LARGE_PAGES` in the *Oracle Database Reference* for Release 12.2, 12.1, or 11.2.

Managing Exadata Cloud Infrastructure I/O Resource Management (IORM)

This topic explains the I/O Resource Management (IORM) feature and how to enable it, modify the IORM settings, and disable it by using the Console or the API.

- About IORM
  The I/O Resource Management (IORM) feature allows you to manage how multiple databases share the I/O resources of an Oracle Exadata cloud VM cluster for systems using the new resource model or DB system
- Using the Console to Manage IORM
- Using the API to manage the I/O resources of an Exadata cloud VM cluster

About IORM

The I/O Resource Management (IORM) feature allows you to manage how multiple databases share the I/O resources of an Oracle Exadata cloud VM cluster for systems using the new resource model or DB system

On an Exadata VM cluster or DB system, all databases share dedicated storage servers which include flash storage. By default, the databases are given equal priority with respect to these resources. The Exadata storage management software uses a first come, first served approach for query processing. If a database executes a major query that overloads I/O resources, overall system performance can be slowed down.

IORM allows you to assign priorities to your databases to ensure critical queries are processed first when workloads exceed their resource allocations. You assign priorities by creating directives that specify the number of shares for each database. The number of shares corresponds to a percentage of resources given to that database when I/O resources are stressed.

Directives work together with an overall optimization objective you set for managing the resources. The following objectives are available:

- **Auto** - (Default objective). IORM determines the optimization objective and continuously and dynamically determines the optimal settings, based on the workloads observed, and resource plans enabled.
- **Balanced** - For critical OLTP and DSS workloads. This setting balances low disk latency and high throughput. This setting limits disk utilization of large I/Os to a lesser extent than low latency to achieve a balance between good latency and good throughput.
- **High throughput** - For critical DSS workloads that require high throughput.
- **Low latency** - For critical OLTP workloads. This setting provides the lowest possible latency by significantly limiting disk utilization.

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.
- Managing I/O Resources

Using the Console to Manage IORM

- To enable IORM on your Exadata cloud VM cluster
  This topic only applies to Exadata Cloud Infrastructure systems using the new resource model.
- To modify the IORM configuration on your cloud VM cluster
  This topic only applies to Exadata Cloud Infrastructure systems using the new resource model.
- To enable IORM on your Exadata DB system
- To modify the IORM configuration on your Exadata DB system
  Use this procedure to change your IORM settings or to disable IORM.

To enable IORM on your Exadata cloud VM cluster

This topic only applies to Exadata Cloud Infrastructure systems using the new resource model.

> **Note:**
>
> If you are enabling IORM for an Exadata DB system, see To enable IORM on your Exadata DB system.

Enabling IORM includes specifying an optimization objective and configuring your resource plan directives.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters** under the **Oracle Exadata Database Service on Dedicated Infrastructure**.
4. In the list of VM clusters, find the VM cluster for which you want to enable IORM, and click its highlighted name. The cluster's details are displayed, showing the IORM status as "Disabled."
5. Click **Enable IORM**.

   It might take a minute for the Enable I/O Resource Management dialog to retrieve the VM cluster information.
6. Select the objective to apply to the resource plan:
   - **Auto** - (Default objective) Dynamically changes the objective based on the resource plan and observed workloads.
   - **Balanced** - Weighs high throughput and low latency evenly.
   - **High throughput** - Provides the best throughput for DSS workloads.

- **Low latency** - Provides the best latency for critical OLTP workloads.

7. Configure the resource plan default directive by setting the number of shares. This number of shares is assigned to each database not associated with a specific directive.
8. In the Resource Plan Directives section, add a directive for each database you want to assign a greater or lesser number of shares than the default directive.

   To add a directive, click **+ Additional Directive**, then specify the database and the number of shares for that database.
9. When you are done adding directives, click **Enable**.

   While the IORM configuration settings are being applied, the VM cluster details page shows the IORM status as "Updating." The update might take several minutes to complete but should have no impact on your ability to perform normal operations on your VM cluster. After a successful update, the IORM status shows as "Enabled."

Related Topics

- The New Exadata Cloud Infrastructure Resource Model
  Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.
- To enable IORM on your Exadata DB system

To modify the IORM configuration on your cloud VM cluster

This topic only applies to Exadata Cloud Infrastructure systems using the new resource model.

> **Note:**
>
> If you are updating an Exadata DB system, see To modify the IORM configuration on your Exadata DB system

Use this procedure to change your IORM settings or to disable IORM.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
4. In the list of VM clusters, find the VM cluster for which you want to update IORM, and click its highlighted name. The cluster's details are displayed, showing the IORM status as "Enabled."
5. Click **Update IORM**.
6. In the Update I/O Resource Management dialog, take one of the following actions:
   - Change your settings - Specify a new objective and adjust your directives, as applicable, and then click **Update**.
   - Disable IORM - Click **Disable IORM**. Disabling IORM removes all your resource plan directives and restores the Auto (default) objective for I/O resource management.

   While the new IORM configuration settings are being applied, the system details page shows the IORM status as "Updating." The update might take several minutes to complete

but should have no impact on your ability to perform normal operations on your DB system. After a successful update, the IORM status shows as "Enabled" or "Disabled," depending on the action you took.

Related Topics

- To modify the IORM configuration on your Exadata DB system
  Use this procedure to change your IORM settings or to disable IORM.

To enable IORM on your Exadata DB system

> **Note:**
>
> This topic only applies to Exadata Cloud Infrastructure instances using the DB system resource model.

Enabling IORM includes specifying an optimization objective and configuring your resource plan directives.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. In the list of DB systems, find the Exadata DB system for which you want to enable IORM, and click its highlighted name.

   The system details are displayed, showing the IORM status as "Disabled."
4. Click **Enable IORM**.

   It might take a minute for the Enable I/O Resource Management dialog to retrieve the DB system information.
5. Select the objective to apply to the resource plan:
   - **Auto** - (Recommended) Dynamically changes the objective based on the resource plan and observed workloads.
   - **Balanced** - Weighs high throughput and low latency evenly.
   - **High throughput** - Provides the best throughput for DSS workloads.
   - **Low latency** - Provides the best latency for critical OLTP workloads.
6. Configure the resource plan default directive by setting the number of shares. This number of shares is assigned to each database not associated with a specific directive.
7. In the Resource Plan Directives section, add a directive for each database you want to assign a greater or lesser number of shares than the default directive.

   To add a directive, click **+ Additional Directive**, then specify the database and the number of shares for that database.
8. When you are done adding directives, click **Enable**.

   While the IORM configuration settings are being applied, the system details page shows the IORM status as "Updating." The update might take several minutes to complete but should have no impact on your ability to perform normal operations on your DB system. After a successful update, the IORM status shows as "Enabled."

To modify the IORM configuration on your Exadata DB system

Use this procedure to change your IORM settings or to disable IORM.

> **Note:**
>
> This topic only applies to Exadata Cloud Infrastructure instances using the DB system resource model.

1. Open the navigation menu. Click **Oracle Database**, then click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
3. In the list of DB systems, find the Exadata DB system for which you want to modify the IORM configuration, and click its highlighted name.

   The system details are displayed, showing the IORM status as "Enabled."
4. Click **Update IORM**.
5. In the Update I/O Resource Management dialog, take one of the following actions:
   - Change your settings - Specify a new objective and adjust your directives, as applicable, and then click **Update**.
   - Disable IORM - Click **Disable IORM**. Disabling IORM removes all your resource plan directives and restores the Auto (default) objective for I/O resource management.

   While the new IORM configuration settings are being applied, the system details page shows the IORM status as "Updating." The update might take several minutes to complete but should have no impact on your ability to perform normal operations on your DB system. After a successful update, the IORM status shows as "Enabled" or "Disabled," depending on the action you took.

Using the API to manage the I/O resources of an Exadata cloud VM cluster

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.

Use these API operations to manage the I/O resources of an Exadata cloud VM cluster. (see The New Exadata Cloud Service Resource Model for more information on this resource type).

- ListCloudVmClusters
- GetCloudVmCluster
- GetCloudVmClusterIormConfig
- UpdateCloudVmClusterIormConfig

Use these API operations to manage the I/O resources of an Exadata DB system.

- ListDbSystems
- GetDbSystem
- GetExadataIormConfig
- UpdateExadataIormConfig

Migrate to Exadata Cloud Infrastructure

For general guidance on methods and tools to migrate databases to Oracle Cloud Infrastructure database services, including Exadata Cloud Infrastructure see "Migrating Databases to the Cloud".

A recommended approach for migrating to Exadata Cloud Infrastructure is using Zero Downtime Migration

- Moving to Oracle Cloud Using Zero Downtime Migration

Related Topics

- Migrating Databases to the Cloud

Moving to Oracle Cloud Using Zero Downtime Migration

Oracle now offers the Zero Downtime Migration service, a quick and easy way to move on-premises databases to Oracle Cloud Infrastructure.

Zero Downtime Migration leverages Oracle Active Data Guard to create a standby instance of your database in an Oracle Cloud Infrastructure system. You switch over only when you are ready, and your source database remains available as a standby. Use the Zero Downtime Migration service to migrate databases individually or at the fleet level. See *Move to Oracle Cloud Using Zero Downtime Migration* for more information.

Related Topics

- Move to Oracle Cloud Using Zero Downtime Migration

Switch an Exadata DB System to the New Resource Model and APIs

If you have existing Exadata DB systems in Oracle Cloud Infrastructure, you can switch them to the new resource model and APIs.

This does not change to the underlying hardware or shape family of your Exadata Cloud Infrastructure instance. The existing DB system APIs will be deprecated for Exadata by Oracle Cloud Infrastructure for all users following written notification and a transition period allowing you to switch to the new API and Console interfaces. Note that this change will not affect bare metal and virtual machine DB systems. Switching to the new resource model does not impact the DB system's existing Exadata databases or client connections. If you have created automation that uses the existing DB system API, your applications may need to be updated to use the new API.

> **Note:**
>
> No new systems can be provisioned with the old DB system resource model/APIs after May 15th, 2021. Support for the old DB system resource model/APIs on existing systems will end on November 15th, 2021. Oracle recommends that you migrate your Exadata Cloud Infrastructure instances to the new resource model APIs as soon as possible.

Converting to the new resource model does not involve any system downtime. After converting your DB system, you will have two new resources in place of the DB system resource: a cloud Exadata infrastructure resource, and a cloud VM cluster resource.

What to Expect After Switching

- Your new cloud Exadata infrastructure resource and cloud VM cluster are created in the same compartment as the DB system they replace.
- Your new cloud Exadata infrastructure resource and cloud VM cluster use the same networking configuration as the DB system they replace.
- After the switch, you cannot perform operations on the old Exadata DB system resource.
- Switching is permanent, and the change cannot be undone
- X6, X7, X8 and Exadata base systems retain their fixed shapes after the switch, and cannot be expanded.
- Switch an Exadata DB system to the new Exadata resource model
  Use the console to perform the switch to the new resource model.

Related Topics

- Switch an Exadata DB system to the new Exadata resource model
  Use the console to perform the switch to the new resource model.

Switch an Exadata DB system to the new Exadata resource model

Use the console to perform the switch to the new resource model.

You have a Exadata DB System that needs to be switched to the new resource Model

1.
2. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
3. Choose your **Compartment**.
4. In the list of DB systems, find the Exadata DB system you want to switch to the new resource model, and click its highlighted name to view the system details.
5. Click **More Actions**, then **Switch to New API**.
6. In the displayed confirmation page, read the **What to expect after switching** section. When you are ready to switch to the new resource model and APIs, click **Start**.

> **Note:**
>
> Switching an Exadata DB system to the new resource model and APIs cannot be reversed. If you have automation for your system that utilizes the DB system APIs, you may need to update your applications prior to switching.

Connect Identity and Access Management (IAM) Users to Oracle Exadata Database Service on Dedicated Infrastructure

You can configure Oracle Exadata Database Service on Dedicated Infrastructure to use Oracle Cloud Infrastructure Identity and Access Management (IAM) authentication and authorization to allow IAM users to access an Oracle Database with IAM credentials.

- Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database
  Learn to enable an Oracle Database instance on Oracle Exadata Database Service on Dedicated Infrastructure to allow user access with an Oracle Cloud Infrastructure IAM database password (using a password verifier), or SSO tokens.
- Prerequisites for Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database
  Review the prerequisites for Identity and Access Management (IAM) authentication on an Oracle Database.
- Enable, Disable, and Re-enable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database
  Learn to enable, disable, and re-enable Identity and Access Management (IAM) Authentication on Oracle Database.
- Manage Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Groups and Policies, Users, Roles, and Database Passwords
- Configuring Client Connection
  Configure various clients to use IAM authentication.
- Configuring Authorization for IAM Users and Oracle Cloud Infrastructure Applications
  An Oracle DBaaS database administrator can map IAM users and Oracle Cloud Infrastructure (OCI) applications to the Oracle Database global schemas and global roles.

Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database

Learn to enable an Oracle Database instance on Oracle Exadata Database Service on Dedicated Infrastructure to allow user access with an Oracle Cloud Infrastructure IAM database password (using a password verifier), or SSO tokens.

- About Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database
  IAM users can connect to the database instance by using either an IAM database password verifier or an IAM token.
- Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password Verifier Authentication
  You can enable an Oracle Database instance to allow user access with an Oracle Cloud Infrastructure IAM database password (using a password verifier).
- Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) SSO Token Based Authentication
  For IAM token access to the database, the client application or tool requests a database token from IAM for the IAM user.

About Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database

IAM users can connect to the database instance by using either an IAM database password verifier or an IAM token.

Using the IAM database password verifier is similar to the database password authentication process. However, instead of the password verifier (encrypted hash of the password) being stored in the database, the verifier is instead stored as part of the OCI IAM user profile.

The second connection method, the use of an IAM token for the database, is more modern. The use of token-based access is a better fit for Cloud resources such as Oracle Databases in the Exadata Cloud Infrastructure. The token is based on the strength that the IAM endpoint can enforce. This can be multi-factor authentication, which is stronger than the use of passwords alone. Another benefit of using tokens is that the password verifier (which is considered sensitive) is never stored or available in memory.

> **Note:**
>
> Oracle Database supports the Oracle DBaaS integration for Oracle Cloud Infrastructure (OCI) IAM with identity domains as well as the legacy IAM, which does not include identity domains. Both default and non-default domain users and groups are supported when using IAM with Identity Domains.
>
> Support for non-default custom domains are only available with Oracle Database Release 19c, Version 19.21 and higher (but not Oracle Database Release 21c).

Oracle Cloud Infrastructure IAM integration with Oracle Exadata Database Service on Dedicated Infrastructure supports the following:

- *Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password Verifier Authentication*
- *Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) SSO Token Based Authentication*

See *Authenticating and Authorizing IAM Users for Oracle DBaaS Databases* for complete details about the architecture for using IAM users on Oracle Exadata Database Service on Dedicated Infrastructure.

Related Topics

- Authenticating and Authorizing IAM Users for Oracle DBaaS Databases
- Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password Verifier Authentication
  You can enable an Oracle Database instance to allow user access with an Oracle Cloud Infrastructure IAM database password (using a password verifier).
- Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) SSO Token Based Authentication
  For IAM token access to the database, the client application or tool requests a database token from IAM for the IAM user.

Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password Verifier Authentication

You can enable an Oracle Database instance to allow user access with an Oracle Cloud Infrastructure IAM database password (using a password verifier).

> **Note:**
>
> Any supported 12c and above database client can be used for IAM database password access to Oracle Database.

An Oracle Cloud Infrastructure IAM database password allows an IAM user to log in to an Oracle Database instance as Oracle Database users typically log in with a username and password. The user enters their IAM username and IAM database password. An IAM database password is a different password than the Oracle Cloud Infrastructure Console password. Using an IAM user with a password verifier, you can log in to Oracle Database with any supported database client.

For password verifier database access, you create the mappings for IAM users and OCI applications to the Oracle Database instance. The IAM user accounts themselves are managed in IAM. The user accounts and user groups can be in either the default domain or in a custom, non-default domain.

For more information about managing IAM database password, see *Managing User Credentials*.

Related Topics

- Managing User Credentials

Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) SSO Token Based Authentication

For IAM token access to the database, the client application or tool requests a database token from IAM for the IAM user.

The client application will pass the database token directly to the database client through the database client API.

If the application or tool has not been updated to request an IAM token, then the IAM user can use OCI CLI to request and store the database token. You can request a database access token (`db-token`) using the following credentials:

- Security tokens (with IAM authentication), delegation tokens (in the OCI cloud shell) and `API-keys`, which are credentials that represent the IAM user to enable the authentication
- Instance principal tokens, which enable instances to be authorized actors (or principals) to perform actions on OCI resources after authentication
- Resource principal token, which is a credential that enables the application to authenticate itself to other OCI services
- Using an IAM user name and IAM database password (can only be requested by database client)

When the IAM users logs into the client with a slash `/` login and the `OCI_IAM` parameter is configured (`sqlnet.ora`, `tnsnames.ora`, or as part of a connect string), then the database client retrieves the database token from a file. If the IAM user submits a user name and password, the connection will use the IAM database verifier access described for client connections that use IAM database password verifiers. If the parameter `PASSWORD_AUTH=OCI_TOKEN`, then the database driver will instead use the username and password to connect directly to IAM and request a database token. The instructions in this guide show how to use the OCI CLI as a helper for the database token. If the application or tool has been updated to work with IAM, then follow the instructions for the application or tool. Some common use cases include the following: SQL*Plus on-premises, SQLcl on-premises, SQL*Plus in Cloud Shell, or applications that use SEP wallets.

There are several ways a database client can obtain an IAM database token:

- A client application or tool can request the database token from IAM for the user and can pass the database token through the client API. Using the API to send the token overrides other settings in the database client. Using IAM tokens requires the latest Oracle Database client 19c (at least 19.16). Some earlier clients (19c and 21c) provide a limited set of capabilities for token access. Oracle Database client 21c does not fully support the IAM token access feature:
  - JDBC-thin on all platforms
    - See *Support for IAM Token-Based Authentication* and *JDBC and UCP Downloads* for more information.
  - SQL*Plus and Oracle Instant Client OCI-C on Linux:
    See *Identity and Access Management (IAM) Token -Based Authentication* for more information
  - Oracle Data Provider for .NET (ODP.NET) Core: .NET clients (latest version of Linux or Windows). .NET software components are available as a free download from the following sites:
    - *Oracle Data access Components - .NET Downloads*
    - *NuGet Gallery*
    - *Visual Studio Code Market Place*
- If the application or tool does not support requesting an IAM database token through the client API, the IAM user can first use the Oracle Cloud Infrastructure command line interface (CLI) to retrieve the IAM database token and save it in a file location. For example, to use SQL*Plus and other applications and tools using this connection method, you first obtain the database token using the Oracle Cloud Infrastructure (OCI) Command Line Interface (CLI). For more information, see db-token get. If the database client is configured for IAM database tokens, when a user logs in with the slash login form, the database driver uses the IAM database token that has been saved in default or specified file location.
- A client application or tool can use an Oracle Cloud Infrastructure IAM instance principal or resource principal to get an IAM database token and use the IAM database token to authenticate itself to an Oracle Database instance. For more information, see *Mapping Instance and Resource Principals*.
- IAM users and OCI applications can request a database token from IAM with several methods, including using an API key. See *Configuring a Client Connection for SQL*Plus That Uses an IAM Token* for an example. See *Authenticating and Authorizing IAM Users for Oracle DBaaS Databases* for a description of other methods such as using a delegation token within an OCI cloud shell.

> **Note:**
>
> If your database is in Restricted Mode, only DBAs with the `RESTRICTED SESSION` privilege can connect to the database.

If a user enters a username/password to login, then the database driver uses the password verifier method to access the database. If the parameter `PASSWORD_AUTH=OCI_TOKEN`, then the database driver will instead user the username and password to connect directly to IAM and request a database token.

Related Topics

- Support for IAM Token-Based Authentication
- JDBC and UCP Downloads
- Identity and Access Management (IAM) Token-Based Authentication
- db-token get
- Oracle Data access Components - .NET Downloads
- NuGet Gallery
- Visual Studio Code Marketplace
- Mapping Instance and Resource Principals
- Configuring a Client Connection for SQL*Plus That Uses an IAM Token
- Authenticating and Authorizing IAM Users for Oracle DBaaS Databases

Prerequisites for Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database

Review the prerequisites for Identity and Access Management (IAM) authentication on an Oracle Database.

- Prerequisites for IAM Authentication on Oracle Database
  Before using IAM authentication on databases in the Exadata Cloud Infrastructure, you must use the Networking service to add a service gateway, a route rule, and an egress security rule to the Virtual Cloud Network (VCN) and subnets where your database resources reside.
- Disable External Authentication Scheme
  Review the prerequisites for enabling IAM user access to Oracle Database.
- Configure TLS to Use IAM Tokens
  When sending IAM tokens from the database client to the database server, a TLS connection must be established. The TLS wallet with the database certificate for the ExaDB-D service instance must be stored under the `WALLET_ROOT` location. Create a tls directory so it looks like: `WALLET_ROOT/<PDB GUID>/tls`.

Prerequisites for IAM Authentication on Oracle Database

Before using IAM authentication on databases in the Exadata Cloud Infrastructure, you must use the Networking service to add a service gateway, a route rule, and an egress security rule to the Virtual Cloud Network (VCN) and subnets where your database resources reside.

1. Create a service gateway in the VCN where your database resources reside by following the instructions in *Task 1: Create the service gateway* in OCI documentation.
2. After creating the service gateway, add a route rule and an egress security rule to each subnet (in the VCN) where the database resources reside so that these resources can use the gateway to use IAM authentication:
   a. Go to the **Subnet Details** page for the subnet.
   b. In the **Subnet Information** tab, click the name of the subnet's Route Table to display its **Route Table Details** page.
   c. In the table of existing Route Rules, check whether there is already a rule with the following characteristics:
      - **Destination**: All IAD Services In Oracle Services Network
      - **Target Type**: Service Gateway
      - **Target**: The name of the service gateway you just created in the VCN

      If such a rule does not exist, click **Add Route Rules** and add a route rule with these characteristics.
   d. Return to the Subnet Details page for the subnet.
   e. In the subnet's Security Lists table, click the name of the subnet's security list to display its Security List Details page.
   f. In the side menu, under **Resources**, click **Egress Rules**.
   g. In the table of existing Egress Rules, check whether there is already a rule with the following characteristics:
      - **Stateless**: No
      - **Destination**: All IAD Services In Oracle Services Network
      - **IP Protocol**: TCP
      - **Source Port Range**: All
      - **Destination Port Range**: 443
   h. If such a rule does not exist, click **Add Egress Rules** and add an egress rule with these characteristics.

Related Topics

- Task 1: Create the service gateway

Disable External Authentication Scheme

Review the prerequisites for enabling IAM user access to Oracle Database.

If the database is enabled for another external authentication scheme, verify that you want to use IAM on the Oracle Database instance. There can only be one external authentication scheme enabled at any given time.

If you want to use IAM and another external authentication scheme is enabled, you must first disable the other external authentication scheme.

Configure TLS to Use IAM Tokens

When sending IAM tokens from the database client to the database server, a TLS connection must be established. The TLS wallet with the database certificate for the ExaDB-D service

instance must be stored under the `WALLET_ROOT` location. Create a tls directory so it looks like: `WALLET_ROOT/<PDB GUID>/tls`.

When configuring TLS between the database client and server there are several options to consider.

- Using a self-signed database server certificate vs a database server certificate signed by a commonly known certificate authority
- One-way TLS (TLS) vs Mutual or two-way TLS (mTLS)
- Client with or without a wallet

Self-Signed Certificate

Using a self-signed certificate is a common practice for internally facing IT resources since you can create these yourself and it's free. The resource (in our case, the database server) will have a self-signed certificate to authenticate itself to the database client. The self-signed certificate and root certificate will be stored in the database server wallet. For the database client to be able to recognize the database server certificate, a copy of the root certificate will also be needed on the client. This self-created root certificate can be stored in a client-side wallet or installed in the client system default certificate store (Windows and Linux only). When the session is established, the database client will check to see that the certificate sent over by the database server has been signed by the same root certificate.

A Well-Known Certificate Authority

Using a commonly known root certificate authority has some advantages in that the root certificate is most likely already stored in the client system default certificate store. There is no extra step for the client to store the root certificate if it is a common root certificate. The disadvantage is that this normally has a cost associated with it.

One-Way TLS

In the standard TLS session, only the server provides a certificate to the client to authenticate itself. The client doesn't need to have a separate client certificate to authenticate itself to the server (similar to how HTTPS sessions are established). While the database requires a wallet to store the server certificate, the only thing the client needs to have is the root certificate used to sign the server certificate.

Two-Way TLS (also called Mutual TLS, mTLS)

In mTLS, both the client and server have identity certificates that are presented to each other. In most cases, the same root certificate will have signed both of these certificates so the same root certificate can be used with the database server and client to authenticate the other certificate. mTLS is sometimes used to authenticate the user since the user identity is authenticated by the database server through the certificate. This is not necessary for passing IAM tokens but can be used when passing IAM tokens.

Client with a Wallet

A client wallet is mandatory when using mTLS to store the client certificate. However, the root certificate can be stored either in the same wallet or in the system default certificate store.

A Client without a Wallet

Clients can be configured without a wallet when using TLS under these conditions: 1) One-way TLS is being configured where the client does not have its own certificate and 2) the root certificate that signed the database server certificate is stored in the system default certificate store. The root certificate would most likely already be there if the server certificate is signed by a common certificate authority. If it's a self-signed certificate, then the root certificate would need to be installed in the system default certificate store to avoid using a client wallet.

For details on how to configure TLS between the database client and database server including the options described above, see *Configuring Transport Layer Security Authentication* in the *Oracle Database Security Guide*.

If you choose to use self-signed certificates and for additional wallet related tasks, see *Managing Public Key Infrastructure (PKI) Elements* in the *Oracle Database Security Guide*.

Related Topics

- Configuring Transport Layer Security Authentication
- Managing Public Key Infrastructure (PKI) Elements

Enable, Disable, and Re-enable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database

Learn to enable, disable, and re-enable Identity and Access Management (IAM) Authentication on Oracle Database.

- Enable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database
  Review the steps to enable or re-enable IAM user access to Oracle Database.
- Disable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database
  Describes the steps to disable IAM external authentication user access for Oracle Database.
- Using Oracle Database Tools with Identity and Access Management (IAM) Authentication
  Review the notes for using Oracle Database tools with IAM authentication enabled.

Enable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database

Review the steps to enable or re-enable IAM user access to Oracle Database.

> **Note:**
>
> Oracle Database supports the Oracle DBaaS integration for Oracle Cloud Infrastructure (OCI) IAM with identity domains as well as the legacy IAM, which does not include identity domains. Both default and non-default domain users and groups are supported when using IAM with Identity Domains.

1. Perform the prerequisites for IAM authorization and authentication on Oracle Database.See *Prerequisites for Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database* for more information.
2. Enable Oracle Cloud Infrastructure (IAM) Authentication and Authorization using the `ALTER SYSTEM` command.

```
ALTER SYSTEM SET IDENTITY_PROVIDER_TYPE=OCI_IAM SCOPE=BOTH;
```

3. Verify the value of `IDENTITY_PROVIDER_TYPE` system parameter.

```
SELECT NAME, VALUE FROM V$PARAMETER WHERE NAME='identity_provider_type';


NAME                      VALUE
------------------------  -------
identity_provider_type    OCI_IAM
```

Related Topics

- Prerequisites for Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database
  Review the prerequisites for Identity and Access Management (IAM) authentication on an Oracle Database.
- Disable External Authentication Scheme
  Review the prerequisites for enabling IAM user access to Oracle Database.

Disable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database

Describes the steps to disable IAM external authentication user access for Oracle Database.

To disable IAM user access on your Oracle Database instance:

1. Disable IAM integration using the `ALTER SYSTEM` command.

```
ALTER SYSTEM RESET IDENTITY_PROVIDER_TYPE SCOPE=BOTH;
```

2. If you also want to remove the IAM policy to allow database access, you may need to review and either modify or remove the IAM groups and the policies you set up to allow access to the database by IAM users.

Using Oracle Database Tools with Identity and Access Management (IAM) Authentication

Review the notes for using Oracle Database tools with IAM authentication enabled.

- Oracle APEX is not supported for IAM users with Oracle Database.
- Database Actions is not supported for IAM users with Oracle Database. See *Provide Database Actions Access to Database Users* for information on using regular database users with Oracle Database.
- Oracle Machine Learning Notebooks and other components are not supported for IAM Authorized users with Oracle Database. See *Add Existing Database User Account to Oracle Machine Learning Components* for information on using regular database users with Oracle Database.

Related Topics

- Provide Database Actions Access to Database Users
- Add Existing Database User Account to Oracle Machine Learning Components

Manage Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Groups and Policies, Users, Roles, and Database Passwords

- Create Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Groups and Policies for IAM Users
  Review the steps to write policy statements for an IAM group to enable IAM user access to Oracle Cloud Infrastructure resources, specifically Oracle Database instances using IAM database tokens.
- Authorize Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Users on Oracle Database
  Review the steps to authorize IAM users on an Oracle Database instance.
- To Exclusively Map a Local IAM User to an Oracle Database Global User
  You can map a local IAM user exclusively to an Oracle Database global user.
- Add Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Roles on Oracle Database
  Optionally, create global roles to provide additional database roles and privileges to IAM users when multiple IAM users are mapped to the same shared global user.
- Create Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password for IAM Users
  To add an IAM user and allow the IAM user to login to Oracle Database by supplying a username and password, you must create an IAM database password.

Create Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Groups and Policies for IAM Users

Review the steps to write policy statements for an IAM group to enable IAM user access to Oracle Cloud Infrastructure resources, specifically Oracle Database instances using IAM database tokens.

A policy is a group of statements that specifies who can access particular resources, and how. Access can be granted for the entire tenancy, databases in a compartment, or individual databases. This means you write a policy statement that gives a specific group a specific type of access to a specific type of resource within a specific compartment.

Note: Defining a policy is required to use IAM tokens to access Oracle Database. A policy is not required when using IAM database password verifiers to access Oracle Database.

1. Create an IAM group for IAM users that will access the database. Review OCI IAM documentation for creating groups and adding IAM users to a group.
   For example, create the group *DBUsers*. For more information, see *Managing Groups*.
2. Write policy statements to enable access to Oracle Cloud Infrastructure resources.
   a. In the Oracle Cloud Infrastructure console, click **Identity and Security**, and then click **Policies**.
   b. To write a policy, click **Create Policy**, and then enter a **Name** and a **Description.**
   c. Use the **Policy Builder** to create a policy. For example, to create a policy to allow users in IAM group DBUsers to access any Oracle Database in their tenancy:

```
Allow group DBUsers to use database-connections in tenancy
```

Where, `database-connections` is the OCI resource name to connect to the database. `Use` is the minimum verb to allow access to the database. Both `use` and `manage` can be used.

For example to create a policy that limits members of *DBUsers* group to access Oracle Databases in the compartment *testing_compartment* only:

```
allow group DBUsers to use database-connections in compartment
testing_compartment
```

For example, to create a policy that limits group access to a single database in a compartment:

```
allow group DBUsers to use database-connections in compartment
testing_compartment where target.database.id =
'ocid1.database.oc1.iad.aaaabbbbcccc'
```

d. Click **Create**.
For more information about policies, see Managing Policies.

Notes for creating policies for use with IAM users on Oracle Database:

- Policies can allow IAM users to access Oracle Database instances across the entire tenancy, in a compartment, or can limit access to a single Oracle Database instance.
- You must use dynamic groups for Instance Principals and Resource Principals. You can create Dynamic Groups and reference dynamic groups in the policies you create to access Oracle Cloud Infrastructure. See *Accessing Cloud Resources by Configuring Policies and Roles and Managing Dynamic Groups* for details.

Related Topics

- Managing Groups
- Accessing Cloud Resources by Configuring Policies and Roles
- Managing Dynamic Groups

Authorize Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Users on Oracle Database

Review the steps to authorize IAM users on an Oracle Database instance.

To authorize IAM users to allow access to Oracle Database, map database global users to IAM groups or directly to IAM users with `CREATE USER` or `ALTER USER` statements with `IDENTIFIED GLOBALLY AS` clause.

The authorization of IAM users to an Oracle Database instance works by mapping IAM global users (schemas) to IAM users (exclusive mapping) or IAM groups (shared schema mapping).

To authorize IAM users on a database instance:

1. Log in as a user with DBA privileges to the database that is enabled to use IAM. A user with the DBA role will need the required `CREATE USER` and `ALTER USER` system privileges for these steps.

2. Create a mapping between the Oracle Database user (schema) with `CREATE USER` or `ALTER USER` statements and include the `IDENTIFIED GLOBALLY AS` clause, specifying the IAM group name. Use the following syntax to map a global user to an IAM group:

```
CREATE USER global_user IDENTIFIED GLOBALLY AS
'IAM_GROUP_NAME=IAM_GROUP_NAME';
```

For example, to map an IAM group named db_sales_group to a shared database global user named sales_group:

```
CREATE USER sales_group IDENTIFIED GLOBALLY AS
'IAM_GROUP_NAME=db_sales_group';
```

This creates a shared global user mapping. The mapping, with the global user `sales_group` is effective for all users in the IAM group. Thus, anyone in the `db_sales_group` can log in to the database using their IAM credentials through the shared mapping of the `sales_group` global user.

If you want to create additional global user mappings for other IAM groups or users, follow these steps for each IAM group or user.

> **Note:**
>
> Database users that are not `IDENTIFIED GLOBALLY` can continue to login as before, even when the Oracle Database is enabled for IAM authentication.

To Exclusively Map a Local IAM User to an Oracle Database Global User

You can map a local IAM user exclusively to an Oracle Database global user.

1. Log in as an user with DBA privileges to the database that is enabled to use IAM. A user with the DBA role has will need the required `CREATE USER` and `ALTER USER` system privileges that you need for these steps.
2. Create a mapping between the Oracle Database user (schema) with `CREATE USER` or `ALTER USER` statements and include the `IDENTIFIED GLOBALLY AS` clause, specifying the IAM local IAM user name. For example, to create a new database global user named `peter_fitch` and map this user to an existing local IAM user named `peterfitch`:

```
CREATE USER peter_fitch IDENTIFIED GLOBALLY AS
'IAM_PRINCIPAL_NAME=peterfitch'
```

You can use either instance principal or resource principal to retrieve database tokens to establish a connection from your application to an Oracle Database instance.

If you are using an instance principal or resource principal, you must map a dynamic group. Thus, you cannot exclusively map instance and resource principals. You only can map them through a shared mapping and putting the instance or resource instance in an IAM dynamic group

Add Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Roles on Oracle Database

Optionally, create global roles to provide additional database roles and privileges to IAM users when multiple IAM users are mapped to the same shared global user.

Creating global roles is optional, but useful when assigning users to a shared schema.

Use a global role to optionally differentiate users who use the same shared schema. For example, a set of users can all have the same shared schema and the shared schema could have the `CREATE SESSION` privilege. Then global roles can be used to provide differentiated privileges and roles assigned to different groups of users who all use the same shared schema.

Granting additional roles to IAM users in Oracle Database works by mapping Oracle Database global roles to IAM groups.

1. Log in as a user with DBA privileges to the database that is enabled to use IAM. A user with the DBA privileges `CREATE ROLE` and `ALTER ROLE` system privileges is needed for these steps.
2. Set database authorization for Oracle Database roles with `CREATE ROLE` or `ALTER ROLE` statements and include the `IDENTIFIED GLOBALLY AS` clause, specifying the IAM group name. Use the following syntax to map a global role to an IAM group:

   ```
   CREATE ROLE global_role IDENTIFIED GLOBALLY AS
   'IAM_GROUP_NAME=IAM_GROUP_of_WHICH_the_IAM_USER_IS_a_MEMBER';
   ```

   For example, to map an IAM group named ExporterGroup to a shared database global role named export_role:

   ```
   CREATE ROLE export_role IDENTIFIED GLOBALLY AS
   'IAM_GROUP_NAME=ExporterGroup';
   ```

3. Use the `GRANT` statements to grant the required privileges or other roles to the global role.

   ```
   GRANT CREATE SESSION TO export_role;
   GRANT DWROLE TO export_role;
   ```

4. If you want an existing database role to be associated with an IAM group, then use the `ALTER ROLE` statement to alter the existing database role to map the role to an IAM group. Use the following syntax to alter an existing database role to map it to an IAM group:

   ```
   ALTER ROLE existing_database_role IDENTIFIED GLOBALLY AS
   'IAM_GROUP_NAME=IAM_Group_Name';
   ```

Follow these steps for each IAM group to add additional global role mappings for other IAM groups.

Create Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password for IAM Users

To add an IAM user and allow the IAM user to login to Oracle Database by supplying a username and password, you must create an IAM database password.

For more information, see *Working with IAM Database Passwords*.

Related Topics

- Working with IAM Database Passwords

Configuring Client Connection

Configure various clients to use IAM authentication.

- Configure a Client Connection for SQL*Plus that Uses an IAM Database Password Verifier
  You can configure SQL*Plus to use an IAM database password verifier.
- Configure Client Connection for SQL*Plus that Uses an IAM Token
  You can configure a client connection for SQL*Plus that uses an IAM token.
- Client Connections That Use a Token Requested by an IAM User Name and Database Password
  You can create a client connection that uses a token requested by an IAM user name and database password.
- Use Instance Principal to Access Database with IAM Authentication
  After the ADMIN user enables OCI IAM on the database, an application can access the database through an OCI IAM database token using an instance principal.
- Configure Proxy Authentication
  Proxy authentication allows an IAM user to proxy to a database schema for tasks such as application maintenance.
- Use Database Link with IAM Authenticated Users
  You can use a database link to connect from one database instance to another as an OCI IAM user.

Configure a Client Connection for SQL*Plus that Uses an IAM Database Password Verifier

You can configure SQL*Plus to use an IAM database password verifier.

As the IAM user, log in to the database by using the following syntax:

```
CONNECT user_name@db_connect_string
Enter password: password
```

In this specification, `user_name` is the IAM user name. There is a limit of 128 bytes for the combined `domain_name/user_name`.

The following example shows how IAM user `peter_fitch` can log in to a database instance.

```
sqlplus /nolog
connect peter_fitch@db_connect_string
Enter password: password
```

Some special characters will require double quotation marks around `user_name` and . For example:

```
"peter_fitch@example.com"@db_connect_string

"IAM database password"
```

Configure Client Connection for SQL*Plus that Uses an IAM Token

You can configure a client connection for SQL*Plus that uses an IAM token.

1. Ensure you have an IAM user account.
2. Check with an IAM administrator and the database administrator to ensure you have a policy allowing you to access the database in the compartment or your tenancy and that you are mapped to a global schema in the database.
3. If your application or tool does not support direct IAM integration, then download, install, and configure the OCI CLI. (See OCI Command Line Interface Quickstart.) Set up an API key as part of the OCI CLI configuration and select default values.
   a. Set up the API key access for the IAM user.
   b. Retrieve the `db-token`. For example:
      - Retrieve a `db-token` with an `API-key` using the OCI CLI:

        ```
        oci iam db-token get
        ```

      - Retrieve `db-token` with a security (or session) token:

        ```
        oci iam db-token get --auth security_token
        ```

      - Retrieve `db-token` with a delegation token: When you log in to the cloud shell, the delegation token is automatically generated and placed in the `/etc` directory. To get this token, execute the following command in the OCI CLI:

        ```
        oci iam db-token get
        ```

      - Using an instance principal to retrieve a `db-token` using OCI CLI:

        ```
        oci iam db-token get --auth instance_principal
        ```

      If the security token has expired, a window will appear so the user can log in to OCI again. This generates the security token for the user. OCI CLI will use this refreshed token to get the `db-token`.

      See Required Keys and OCIDs for more information.
4. Ensure that you are using the latest release updates for the Oracle Database client releases 19c.
   This configuration only works with the Oracle Database client release 19c.
5. Follow the existing process to download the wallet from the database and then follow the directions for configuring it for use with SQL*Plus.
   a. Confirm that DN matching is enabled by looking for `SSL_SERVER_DN_MATCH=ON` in `sqlnet.ora`.

b. Configure the database client to use the IAM token by adding `TOKEN_AUTH=OCI_TOKEN` to the `sqlnet.ora` file. Because you will be using the default locations for the database token file, you do not need to include the token location.

The `TOKEN_AUTH` and `TOKEN_LOCATION` values in the `tnsnames.ora` connect strings take precedence over the `sqlnet.ora` settings for that connection. For example, for the connect string, assuming that the token is in the default location (`~/.oci/db-token` for Linux):

```
(description=
  (retry_count=20)(retry_delay=3)
  (address=(protocol=tcps)(port=1522)
  (host=example.us-phoenix-1.oraclecloud.com))

(connect_data=(service_name=aaabbbccc_exampledb_high.example.oraclecloud.co
m))
  (security=(ssl_server_cert_dn="CN=example.uscom-east-1.oraclecloud.com,
     OU=Oracle BMCS US, O=Example Corporation,
     L=Redwood City, ST=California, C=US")
  (TOKEN_AUTH=OCI_TOKEN)))
```

After the connect string is updated with the `TOKEN_AUTH` parameter, the IAM user can log in to the database instance by running the following command to start SQL*Plus. You can include the connect descriptor itself or use the name of the descriptor from the `tnsnames.ora` file.

```
connect /@exampledb_high
```

Or:

```
connect /@(description=
  (retry_count=20)(retry_delay=3)
  (address=(protocol=tcps)(port=1522)
  (host=example.us-phoenix-1.oraclecloud.com))

(connect_data=(service_name=aaabbbccc_exampledb_high.example.oraclecloud.com))
  (security=(ssl_server_cert_dn="CN=example.uscom-east-1.oraclecloud.com,
     OU=Oracle BMCS US, O=Example Corporation,
     L=Redwood City, ST=California, C=US")
  (TOKEN_AUTH=OCI_TOKEN)))
```

The database client is already configured to get a `db-token` because `TOKEN_AUTH` has already been set, either through the `sqlnet.ora` file or in a connect string. The database client gets the `db-token` and signs it using the private key and then sends the token to the database. If an IAM user name and IAM database password are specified instead of slash `/`, then the database client will connect using the password instead of using the `db-token`.

Client Connections That Use a Token Requested by an IAM User Name and Database Password

You can create a client connection that uses a token requested by an IAM user name and database password.

- IAM users can connect to the Oracle DBaaS instance by using an IAM token that was retrieved using an IAM user name and IAM database password.

For more information, see *About Client Connections That Use a Token Requested by an IAM User Name and Database Password*

- To set these parameters, you modify either the `sqlnet.ora` file or the `tnsnames.ora` file.
  For more information, see *Parameters to Set for Client Connections That Use a Token Requested by an IAM User Name and Database Password*
- You can configure the database client to retrieve the IAM database token using the provided IAM user name and IAM database password.
  For more information, see *Configuring the Database Client to Retrieve a Token Using an IAM User Name and Database Password*
- You can enable an IAM user name and a secure external password store (SEPS) to request the IAM database token.
  For more information, see *Configuring a Secure External Password Store Wallet to Retrieve an IAM Token*

Related Topics

- Client Connections That Use a Token Requested by an IAM User Name and Database Password
- About Client Connections That Use a Token Requested by an IAM User Name and Database Password
- Parameters to Set for Client Connections That Use a Token Requested by an IAM User Name and Database Password
- Configuring the Database Client to Retrieve a Token Using an IAM User Name and Database Password
- Configuring a Secure External Password Store Wallet to Retrieve an IAM Token

Use Instance Principal to Access Database with IAM Authentication

After the ADMIN user enables OCI IAM on the database, an application can access the database through an OCI IAM database token using an instance principal.

For more information, see *Accessing the Oracle Cloud Infrastructure API Using Instance Principals*.

For more Information, see *Accessing the Database Using an Instance Principal or a Resource Principal*.

Related Topics

- Accessing the Oracle Cloud Infrastructure API Using Instance Principals
- Accessing the Database Using an Instance Principal or a Resource Principal

Configure Proxy Authentication

Proxy authentication allows an IAM user to proxy to a database schema for tasks such as application maintenance.

Proxy authentication is typically used to authenticate the real user and then authorize them to use a database schema with the schema privileges and roles in order to manage an application. Alternatives such as sharing the application schema password are considered insecure and unable to audit which actual user performed an action.

A use case can be in an environment in which a named IAM user who is an application database administrator can authenticate by using their credentials and then proxy to a

database schema user (for example, `hrapp`). This authentication enables the IAM administrator to use the `hrapp` privileges and roles as user `hrapp` in order to perform application maintenance, yet still use their IAM credentials for authentication. An application database administrator can sign in to the database and then proxy to an application schema to manage this schema.

You can configure proxy authentication for both the password authentication and token authentication methods.

Configuring Proxy Authentication for the IAM User

To configure proxy authentication for an IAM user, the IAM user must already have a mapping to a global schema (exclusive or shared mapping). A separate database schema for the IAM user to proxy to must also be available.

After you ensure that you have this type of user, alter the database user to allow the IAM user to proxy to it.

1. Log in to the database instance as a user who has the `ALTER USER` system privileges.
2. Grant permission for the IAM user to proxy to the local database user account. An IAM user cannot be referenced in the command so the proxy must be created between the database global user (mapped to the IAM user) and the target database user.In the following example, `hrapp` is the database schema to proxy to, and `peterfitch_schema` is the database global user exclusively mapped to user `peterfitch`.

   ```
   ALTER USER hrapp GRANT CONNECT THROUGH peterfitch_schema;
   ```

At this stage, the IAM user can log in to the database instance using the proxy. For example:

- To connect using a password verifier:

  ```
  CONNECT peterfitch[hrapp]@connect_string
  Enter password: password
  ```

- To connect using a token:

  ```
  CONNECT [hrapp]/@connect_string
  ```

Validating the IAM User Proxy Authentication

You can validate the IAM user proxy configuration for both password and token authentication methods.

1. Connect as the IAM user and proxied to the database user. Run the `SHOW USER` and `SELECT SYS_CONTEXT` commands.
   For example, suppose you want to check the proxy authentication of the IAM user *peterfitch* when they proxy to database user *hrapp*. You will need to connect to the database using the different types of authentication methods shown here, but the output of the commands that you execute will be the same for all types.
   - For password authentication:

     ```
     CONNECT peterfitch[hrapp]/password\!@connect_string SHOW USER;


     --The output should be USER is "HRAPP"
     SELECT SYS_CONTEXT('USERENV','AUTHENTICATION_METHOD') FROM DUAL;
     --The output should be "PASSWORD_GLOBAL"
     ```

```
SELECT SYS_CONTEXT('USERENV','PROXY_USER') FROM DUAL;
--The output should be "PETERFITCH_SCHEMA"
SELECT SYS_CONTEXT('USERENV','CURRENT_USER') FROM DUAL;
--The output should be "HRAPP"
```

- For token authentication:

```
CONNECT [hrapp]/@connect_string
SHOW USER;


--The output should be USER is "HRAPP "
SELECT SYS_CONTEXT('USERENV','AUTHENTICATION_METHOD') FROM DUAL;
--The output should be "TOKEN_GLOBAL"
SELECT SYS_CONTEXT('USERENV','PROXY_USER') FROM DUAL;
--The output should be "PETERFITCH_SCHEMA"
SELECT SYS_CONTEXT('USERENV','CURRENT_USER') FROM DUAL;
--The output should be "HRAPP"
```

Use Database Link with IAM Authenticated Users

You can use a database link to connect from one database instance to another as an OCI IAM user.

You can use either connected user or fixed user database link to connect to a database as an OCI IAM user.

> **Note:**
>
> Current user database link is not supported for connecting to a database in Exadata Cloud Infrastructure as an OCI IAM user.

- **Connected User Database Link**: For a connected user database link, an IAM user must be mapped to a schema in both the source and target databases connected by a database link. You can use a database password verifier or an IAM database token to use a connected user database link.
- **Fixed User Database Link**: A fixed user database link can be created using a database user or an IAM user. When using an IAM user as a fixed user database link, the IAM user must have a schema mapping in the target database. The IAM user for a database link can be configured with a password verifier only.

Configuring Authorization for IAM Users and Oracle Cloud Infrastructure Applications

An Oracle DBaaS database administrator can map IAM users and Oracle Cloud Infrastructure (OCI) applications to the Oracle Database global schemas and global roles.

- About Configuring Authorization for IAM Users and Oracle Cloud Infrastructure Applications
  You create the mappings for IAM users and Oracle Cloud Infrastructure (OCI) applications to database users (schemas) in the Oracle DBaaS.

- Mapping an IAM Group to a Shared Oracle Database Global User
  Oracle Database global users that are mapped to IAM groups and IAM dynamic groups give IAM users and OCI applications a schema when they log in along with the privileges and roles granted to that schema.
- Mapping an IAM Group to an Oracle Database Global Role
  Oracle Database global roles that are mapped to IAM groups and dynamic groups give member users and applications additional privileges and roles above what they have been granted through their login schemas.
- Exclusively Mapping an IAM User to an Oracle Database Global User
  You can map an IAM user exclusively to an Oracle Database global user.
- Altering or Migrating an IAM User Mapping Definition
  You can update an IAM user to a database global user mapping by using the `ALTER USER` statement.
- Mapping Instance and Resource Principals
  Instance principals and resource principals can be used by applications to retrieve database tokens to establish a connection to an Oracle DBaaS instance.
- Verifying the IAM User Logon Information
  After you configure and authorize an IAM user for the Oracle DBaaS instance, you can verify the user logon information by executing a set of SQL queries on the Oracle database side.

About Configuring Authorization for IAM Users and Oracle Cloud Infrastructure Applications

You create the mappings for IAM users and Oracle Cloud Infrastructure (OCI) applications to database users (schemas) in the Oracle DBaaS.

There is a difference with authorization between IAM database password authentication and using IAM token based authentication. IAM database password verifier authorization is only based on mappings of database schemas and global roles to IAM users and group. With IAM token based authentication, IAM policies are an additional authorization for IAM users to access their tenancy databases. An IAM user must be authorized through an IAM policy **and** be authorized through a mapping to a database global schema (exclusive or shared).

For both token and password verifier database access, you create the mappings for IAM users and OCI applications to the Oracle DBaaS instance. The IAM user accounts themselves are managed in IAM. The user accounts and user groups can be in either the default domain or in a custom, non-default domain.

When the IAM user accesses the Oracle DBaaS instance with a token, the database will perform an authorization check against IAM policies to ensure the user is allowed to access the database. If the IAM user is allowed to access the database by IAM policy, then the database will query IAM for the user groups. When using password verifier authentication, the database will query IAM for user groups once the IAM user successfully completes authentication. The database queries the IAM endpoint to find the groups of which the user is a member. If your deployment is using shared schemas, then one of the IAM groups will map to a shared database schema and the IAM user will be assigned to that database schema. The IAM user will have the roles and privileges that are granted to the database schema. Because multiple IAM users can be assigned to the same shared database schema, only the minimal set of roles and privileges should be granted to the shared schema. In some cases, no privileges and roles should be granted to the shared schema. Users will be assigned the appropriate set of roles and schemas through database global roles. Global roles are mapped to IAM groups. This way, different users can have different roles and privileges even if they are mapped to the same database shared schema. A newly hired user will be assigned to an IAM

group mapped to a shared schema and then to one or more additional groups mapped to global roles to gain the additional roles and privileges required to complete their tasks. The combination of shared schemas and global roles allows for centralized authorization management with minimal changes to the database operationally. The database must be initially provisioned with the set of shared schemas and global roles mapped to the appropriate IAM groups, but then user authorization management can happen within IAM.

Ensure that the IAM user is only mapped to one schema, either through exclusive mapping to a database schema or as a member of one IAM group that is mapped to a shared database schema. If more than one schema is mapped for an IAM user, then the database will take exclusive mapping as precedence over any group mapping to a shared schema. If more than one group is mapped for a user, then the database will select the oldest mapping.

When using global roles to grant privileges and roles to the user, remember that the maximum number of enabled roles in a session is 150.

If you drop and recreate IAM users and groups using the same names, then the mappings from the database to IAM using the same names will continue to work. However, recreating an IAM user will require the IAM user to do one or more of the following: create the IAM database password, re-upload the API public key, update the OCI configuration file, and then re-examine the IAM policy for database authentication and authorization with IAM. If the IAM policy specifies a group that can use or manage the `database-connections` and `autonomous-database-family` resource types, then the user will need to be added to that group to allow IAM authentication and authorization.

Accessing the database with tokens requires the user to be authorized by IAM policy and by database mapping. Accessing the database with the IAM database password verifier requires authorization through database mapping. If no database schema mapping exists for the IAM user, the IAM user is prevented from accessing the database even if they have a valid token or password.

IAM users get their authorizations to perform various tasks based on the roles that they have been granted. The following scenarios are possible:

- **IAM group mapped to a shared Oracle Database global user:** With the shared database global user account, an IAM user is assigned to a shared database schema (user) through the mapping of an IAM group to the shared schema. The IAM users that are members of the group can connect to the database through this shared schema. Use of shared schemas allows for centralized management of user authorization in IAM.
- **IAM group mapped to an Oracle Database global role:** The privileges that have been granted to the shared Oracle Database global role become available to the users who have added to the IAM group.
- **Local IAM user exclusively mapped to an Oracle Database global user:** With an exclusive global user mapping, a dedicated database user is exclusively mapped to a local IAM user. Not as common as the shared database schema, this user is created for when the user requires their own schema objects. Oracle recommends that you grant database privileges to these users through global roles, which facilitates authorization management. These users can also have direct privilege and role grants to their exclusive schema.
  In IAM with Identity Domains, users and groups are supported in the default domain as well as custom non-default domains. When you specify users and groups in the default domain, then no domain prefix is required. When you specify users and groups in a non-default domain, then the domain must be prefixed.

Mapping an IAM Group to a Shared Oracle Database Global User

Oracle Database global users that are mapped to IAM groups and IAM dynamic groups give IAM users and OCI applications a schema when they log in along with the privileges and roles granted to that schema.

1. Log in to the Oracle DBaaS instance as a user who has the `CREATE USER` or `ALTER USER` system privilege.
2. Run the `CREATE USER` or `ALTER USER` statement with the `IDENTIFIED GLOBALLY AS` clause specifying the IAM group name (which can be a dynamic group).

   For example, to create a new database global user account named `shared_sales_schema` and map it to an existing IAM group named `WidgetSalesGroup`:

   ```
   CREATE USER shared_sales_schema IDENTIFIED GLOBALLY AS
   'IAM_GROUP_NAME=WidgetSalesGroup';
   ```

   The following example shows how to accomplish this for a non-default domain:

   ```
   CREATE USER shared_sales_schema IDENTIFIED GLOBALLY AS
   'IAM_GROUP_NAME=sales_domain/WidgetSalesGroup';
   ```

Mapping an IAM Group to an Oracle Database Global Role

Oracle Database global roles that are mapped to IAM groups and dynamic groups give member users and applications additional privileges and roles above what they have been granted through their login schemas.

Global roles cannot be granted to a database schema (user), they can only be mapped to a group and be assigned to an IAM user when accessing the database.

1. Log in to the Oracle DBaaS instance as a user who has been granted the `CREATE ROLE` or `ALTER ROLE` system privilege
2. Run the `CREATE ROLE` or `ALTER ROLE` statement with the `IDENTIFIED GLOBALLY AS` clause specifying the name of the IAM group (which can be a dynamic group).

   For example, to create a new database global role named `widget_mgr_role` and map it to an existing IAM group named `WidgetManagerGroup`, using the default domain:

   ```
   CREATE ROLE widget_mgr_role IDENTIFIED GLOBALLY AS
   'IAM_GROUP_NAME=WidgetManagerGroup';
   ```

   The following example shows how to create the role by specifying a non-default domain, `sales_domain`:

   ```
   CREATE ROLE widget_sales_role IDENTIFIED GLOBALLY AS
   'IAM_GROUP_NAME=sales_domain/WidgetManagerGroup';
   ```

   All members of the `WidgetManagerGroup` in the `sales_domain` domain will be authorized with the database global role `widget_sales_role` when they log in to the database.

Exclusively Mapping an IAM User to an Oracle Database Global User

You can map an IAM user exclusively to an Oracle Database global user.

1. Log in to the Oracle DBaaS instance as a user who has been granted the `CREATE USER` or `ALTER USER` system privilege.
2. Run the `CREATE USER` or `ALTER USER` statement with the `IDENTIFIED GLOBALLY AS` clause specifying the IAM database user name.

   By default, the IAM database user name is the same as the IAM user name, including the domain name. You can also create a unique IAM database user name for ease of authentication to the database. In your OCI IAM user profile, you can create a unique IAM database user name for ease of authentication to the database. This can be set when you create and manage your IAM database password in your IAM profile. Adding or changing the IAM database user name will invalidate the IAM user to schema mapping, so the database schema will need to be remapped to the new IAM database user name.

   For example, to create a new database global user named `peter_fitch` and map this user to an existing IAM user named with an IAM database user name of `peterfitch`, using the default domain:

   ```
   CREATE USER peter_fitch IDENTIFIED GLOBALLY AS
   'IAM_PRINCIPAL_NAME=peterfitch';
   ```

   The following example shows how to create the user by specifying a non-default domain, `sales_domain`:

   ```
   CREATE USER peter_fitch2 IDENTIFIED GLOBALLY AS
   'IAM_PRINCIPAL_NAME=sales_domain/peterfitch';
   ```

Altering or Migrating an IAM User Mapping Definition

You can update an IAM user to a database global user mapping by using the `ALTER USER` statement.

You can update database schemas that were mapped to an IAM user, and whose accounts were created using any of the `CREATE USER` statement clauses: `IDENTIFIED BY` password, `IDENTIFIED EXTERNALLY`, or `IDENTIFIED GLOBALLY`. This is useful when migrating existing schemas to using IAM. If you delete and recreate an IAM user or an IAM group using the exact same name as the previous IAM user or group, then the existing mapping from the database that uses that IAM user or IAM group name will continue to work.

1. Log in to the Oracle DBaas instance as a user who has been granted the `ALTER USER` system privilege.
2. Run the `ALTER USER` statement with the `IDENTIFIED GLOBALLY AS` clause.

   For example, suppose you want to change the existing schema `shared_sales_schema` to a different IAM group:

   ```
   ALTER USER shared_sales_schema IDENTIFIED GLOBALLY AS
   'IAM_GROUP_NAME=BiggerWidgetSalesGroup';
   ```

The following example shows how to modify the schema by specifying a non-default domain, `sales_domain`:

```
ALTER USER shared_sales_schema IDENTIFIED GLOBALLY AS
'IAM_GROUP_NAME=sales_domain/BiggerWidgetSalesGroup';
```

Mapping Instance and Resource Principals

Instance principals and resource principals can be used by applications to retrieve database tokens to establish a connection to an Oracle DBaaS instance.

Only dynamic groups can be mapped when you use instance and resource principals. You cannot exclusively map instance and resource principals; you only can map them through a shared mapping and putting the instance or resource instance in an IAM dynamic group.

Related Topics

- Managing Dynamic Groups
- Calling Services from an Instance
- Accessing Other Oracle Cloud Infrastructure Resources from Running Functions

Verifying the IAM User Logon Information

After you configure and authorize an IAM user for the Oracle DBaaS instance, you can verify the user logon information by executing a set of SQL queries on the Oracle database side.

1. Log in to the Oracle DBaaS instance as an IAM user that you have just configured and authorized.

   For example, to log in to the database instance `inst1` as the database global user `peterfitch`, who is using the default domain in IAM:

   ```
   sqlplus /nolog
   CONNECT "peterfitch"@inst1
   Enter password: password
   ```

   This example shows how to log in if user `peterfitch` is in a non-default domain, `sales_domain`:

   ```
   sqlplus /nolog
   CONNECT "sales_domain/peterfitch"@inst1
   Enter password: password
   ```

2. Verify the mapped global user.

   The mapped global user is the database user account that has the IAM user authorization. User `PETER_FITCH_SCHEMA` is considered a global user with exclusive mapping for the IAM user `peterfitch`, while user `WIDGET_SALES` is considered a global user with shared mapping for IAM group `widget_sales_group` of which `peterfitch` is a member.

   ```
   SHOW USER;
   ```

   Output similar to the following appears, depending on if it is an exclusive mapping or a shared mapping:

   ```
   USER is "PETER_FITCH_SCHEMA"
   ```

   Or

   ```
   USER is "WIDGET_SALES"
   ```

3. Find the roles that have been granted to the centrally managed user.

```
SELECT ROLE FROM SESSION_ROLES ORDER BY ROLE;
```

Output similar to the following appears:

```
ROLE
----------------------------------------------------------------
WIDGET_SALES_ROLE
...
```

4. Run the following queries to check the `SYS_CONTEXT` namespace values for the current schema being used in this database session, current user name, session user name, authentication method, authenticated identity, enterprise identity, identification type, and server type.

- Verify the current schema that is being used in this database session. A database schema is an object container that identifies the objects it contains. The current schema is the default container for objects name resolution in this database session.

```
SELECT SYS_CONTEXT('USERENV', 'CURRENT_SCHEMA') FROM DUAL;
```

Output similar to the following appears, depending on if it is an exclusive mapping or a shared mapping:

```
SYS_CONTEXT('USERENV','CURRENT_SCHEMA')
--------------------------------------------------------------------------
PETER_FITCH_SCHEMA
```

Or

```
SYS_CONTEXT('USERENV','CURRENT_SCHEMA')
--------------------------------------------------------------------------
WIDGET_SALES
```

- Verify the current user. In this case, the current user is the same as the current schema.

```
SELECT SYS_CONTEXT('USERENV', 'CURRENT_USER') FROM DUAL;
```

Output similar to the following appears, depending on if it is an exclusive mapping or a shared mapping:

```
SYS_CONTEXT('USERENV','CURRENT_USER')
--------------------------------------------------------------------------
PETER_FITCH_SCHEMA
```

Or

```
SYS_CONTEXT('USERENV','CURRENT_USER')
--------------------------------------------------------------------------
WIDGET_SALES
```

- Verify the session user.

```
SELECT SYS_CONTEXT('USERENV', 'SESSION_USER') FROM DUAL;
```

Output similar to the following appears, depending on if it is an exclusive mapping or a shared mapping:

```
SYS_CONTEXT('USERENV','SESSION_USER')
--------------------------------------------------------------------------------
PETER_FITCH_SCHEMA
```

Or

```
SYS_CONTEXT('USERENV','SESSION_USER')
--------------------------------------------------------------------------------
WIDGET_SALES
```

- Verify the authentication method.

```
SELECT SYS_CONTEXT('USERENV', 'AUTHENTICATION_METHOD') FROM DUAL;
```

Output similar to the following appears:

```
SYS_CONTEXT('USERENV','AUTHENTICATION_METHOD')
--------------------------------------------------------------------------------
PASSWORD_GLOBAL
```

If the user is authenticating with a token, then the output is `TOKEN_GLOBAL`.

- Verify the authenticated identity for the enterprise user. The IAM authenticated user identity is captured and audited when this user logs on to the database.

```
SELECT SYS_CONTEXT('USERENV', 'AUTHENTICATED_IDENTITY') FROM DUAL;
```

Output similar to the following appears:

```
SYS_CONTEXT('USERENV','AUTHENTICATED_IDENTITY')
--------------------------------------------------------------------------------
sales_domain/peterfitch
```

- If a user nickname has been set for the enterprise user, then verify this nickname.

```
SELECT SYS_CONTEXT('USERENV', 'USER_NICKNAME') FROM DUAL;
```

Output similar to the following appears:

```
SYS_CONTEXT('USERENV','USER_NICKNAME')
--------------------------------------------------------------------------------
pfitch
```

- Verify the centrally managed user's enterprise identity.

```
SELECT SYS_CONTEXT('USERENV', 'ENTERPRISE_IDENTITY') FROM DUAL;
```

Enterprise Identity will show the OCI Identity (OCID) of the IAM user or OCI application. Output similar to the following appears:

```
SYS_CONTEXT('USERENV','ENTERPRISE_IDENTITY')
--------------------------------------------------------------------------------
ocid1.user.region1..aaaaaaaaj7ot4g2sagkjtw3enbg4ied3x554zwyywurgrm2232j4crm5zha
```

- Verify the identification type.

```
SELECT SYS_CONTEXT('USERENV', 'IDENTIFICATION_TYPE') FROM DUAL
```

Output similar to the following appears, depending on if it is an exclusive mapping or a shared mapping:

```
SYS_CONTEXT('USERENV','IDENTIFICATION_TYPE')
--------------------------------------------------------------------------------
GLOBAL EXCLUSIVE
```

Or

```
SYS_CONTEXT('USERENV','IDENTIFICATION_TYPE')
------------------------------------------------------------------------
GLOBAL SHARED
```

- Verify the server type.

```
SELECT SYS_CONTEXT('USERENV', 'LDAP_SERVER_TYPE') FROM DUAL;
```

Output similar to the following appears. In this case, the LDAP server type is IAM.

```
SYS_CONTEXT('USERENV','LDAP_SERVER_TYPE')
------------------------------------------------------------------------
OCI_IAM
```

Authenticating and Authorizing Microsoft Azure Active Directory Users for Oracle Databases

An Oracle Database instance can be configured for Microsoft Azure AD users to connect using Azure `OAuth2` access tokens.

- Introduction to Authorizing Microsoft Azure AD Users for an Oracle Database
  Before you begin authenticating and authorizing Microsoft Azure AD users for an Oracle Database, you should understand the overall process.
- Configuring the Oracle Database for Microsoft Azure AD Integration
  The Microsoft Azure AD integration with the Oracle Database instance requires the database to be registered with Azure AD so that the database can request the Azure AD public key.
- Mapping Oracle Database Schemas and Roles
  Azure AD users will be mapped to one database schema and optionally to one or more database roles.
- Configuring Azure AD Client Connections to the Oracle Database
  You can configure client connections to connect with the Azure AD registered database
- Trace Files for Troubleshooting Oracle Database Client Connections with Azure AD
  You can use trace files to troubleshoot Oracle Database client connections with Azure AD connections.

Introduction to Authorizing Microsoft Azure AD Users for an Oracle Database

Before you begin authenticating and authorizing Microsoft Azure AD users for an Oracle Database, you should understand the overall process.

- About Authorizing Microsoft Azure AD Users for an Oracle Exadata Database Service on Dedicated Infrastructure
  Users for Oracle Exadata Database Service on Dedicated Infrastructure can be centrally managed in a Microsoft Azure Active Directory (Azure AD) service.

- Azure AD Users Mapping to an Oracle Database Schema and Roles
  Microsoft Azure users must be mapped to an Oracle Database schema and have the necessary privileges (through roles) before being able to authenticate to the Oracle Database instance.
- Use Cases for Connecting to an Oracle Database Using Azure AD
  Oracle Database supports four types of use cases for connecting to an Oracle Database instance using Microsoft Azure Active Directory.
- General Process of Integrating Microsoft Azure AD with Oracle Exadata Database Service on Dedicated Infrastructure
  Both the Oracle and the Microsoft Azure administrators play roles in configuring the connection between Oracle Exadata Database Service on Dedicated Infrastructure and Microsoft Azure AD.

About Authorizing Microsoft Azure AD Users for an Oracle Exadata Database Service on Dedicated Infrastructure

Users for Oracle Exadata Database Service on Dedicated Infrastructure can be centrally managed in a Microsoft Azure Active Directory (Azure AD) service.

You can perform this integration in the following Oracle Database environments:

- On-premises Oracle Database release 19.16 and later, but not for Oracle Database 21c
- Oracle Exadata Database Service on Dedicated Infrastructure on Database versions 19.17 and later. This feature is not supported on Oracle Database release 21c.
- Oracle Base Database Service

The instructions for configuring Azure AD use the term "Oracle Database" to encompass these environments.

This type of integration enables the Azure AD user to access an Oracle Exadata Database Service on Dedicated Infrastructure instance. Azure AD users and applications can log in with Azure AD Single Sign On (SSO) credentials to get an Azure AD `OAuth2` access token to send to the database.

The administrator creates and configures the application registration (app registration) of the Oracle Exadata Database Service on Dedicated Infrastructure instance with Azure AD. The database administrator also creates application (app) roles for the database app registration in Azure AD, and assigns these roles to Azure AD users, groups, and applications. These app roles will be mapped to the database global schemas and global roles. An Azure AD principal that is assigned to an app role will be mapped to either a database global schema or database global role. An Oracle global schema can also be mapped exclusively to an Azure AD user. When the principal is a guest user or service principal, they can only be mapped to the database schema through an Azure app role. An Oracle global role can only be mapped to an Azure app role.

Tools and applications that are updated to support Azure AD tokens can authenticate users directly with Azure AD and pass the database access token to the Oracle Exadata Database Service on Dedicated Infrastructure instance. You can configure existing database tools such as SQL*Plus to use an Azure AD token from a file location. In these cases, Azure AD tokens can be retrieved using tools like Microsoft PowerShell or Azure CLI and put into a file location. An Azure AD `OAuth2` database access token is a bearer token with an expiration time. The Oracle Database client driver will ensure that the token is in a valid format and that it has not expired before passing it to the database. The token is scoped for the database. Assigned app roles for the Azure AD principal are included as part of the access token. The directory location for the Azure AD token should only have enough permission for the user to write the token file

to the location and the database client to retrieve these files (for example, just read and write by the process user). Because the token allows access to the database, it should be protected within the file system.

Azure AD users can request a token as a client registered with Azure AD app registration by using methods such as the following:

- Passing the Azure AD user name and password through a command line, script, file, or any other supported method
- Entering the Azure AD credentials into an Azure AD authentication screen with or without multi-factor authentication

Oracle Exadata Database Service on Dedicated Infrastructure supports the following Azure AD authentication flows:

- Resource owner password credential (ROPC), which is used in non-graphic user interface environments when a pop-up window cannot be used to authenticate a user.
- Authorization code, which is used when a browser can be used to enter credentials for the user
- Client credentials, which are for applications that connect as themselves (and not the end-user)
- On-Behalf-Of (OBO), where an application requests an access token on behalf of a logged-in user to send to the database

Oracle Exadata Database Service on Dedicated Infrastructure accepts tokens representing the following Azure AD principals:

- Azure AD user, who is registered user in the Azure AD tenancy
- Guest user, who is registered as a guest user in the Azure AD tenancy
- Service, which is the registered application connecting to the database as itself with the client credential flow (connection pool use case)

Azure AD Users Mapping to an Oracle Database Schema and Roles

Microsoft Azure users must be mapped to an Oracle Database schema and have the necessary privileges (through roles) before being able to authenticate to the Oracle Database instance.

In Microsoft Azure, an Azure AD administrator can assign users, groups, and applications to the database app roles.

Exclusively mapping an Azure AD schema to a database schema requires the database administrator to create a database schema when the Azure AD user joins the organization or is authorized to the database. The database administrator must also modify the privileges and roles that are granted to the database schema to align them with the tasks the Azure AD user is assigned to. When the Azure AD user leaves the organization, the database administrator must drop the database schema so that an unused account is not left on the database. Using the database app roles enables the Azure AD administrator to control access and roles by assigning users to app roles that are mapped to global schemas and global roles. This way, user access to the database is managed by Azure AD administrators and database administrators do not need to create, manage, and drop schemas for every user.

An Azure AD user can be mapped to a database schema (user) either exclusively or through an app role.

- **Creating an exclusive mapping between an Azure AD user and an Oracle Database schema.** In this type of mapping, the database schema must be created for the Azure AD user. Database privileges and roles that are needed by the Azure AD user must be granted to the database schema. The database schema not only must be created when the Azure AD user is authorized to the database, but the granted privileges and roles must be modified as the Azure AD roles and tasks change. Finally, the database schema must be dropped when the Azure AD user leaves the organization.
- **Creating a shared mapping between an Azure AD app role and an Oracle Database schema.** This type of mapping, which is more common than exclusive mappings, is for Azure AD users who have been assigned directly to the app role or is a member of an Azure AD group that is assigned to the app role. The app role is mapped to an Oracle Database schema (shared schema mapping). Shared schema mapping allows multiple Azure AD users to share the same Oracle Database schema so a new database schema is not required to be created every time a new user joins the organization. This operational efficiency allows database administrators to focus on database application maintenance, performance, and tuning tasks instead of configuring new users, updating privileges and roles, and removing accounts.

In addition to database roles and privileges being granted directly to the mapped global schema, additional roles and privileges can be granted through mapped global roles. Different Azure AD users mapped to the same shared global schema may need different privileges and roles. Azure app roles can be mapped to Oracle Database global roles. Azure AD users who are assigned to the app role or are a member of an Azure AD group that is assigned to the app role will be granted the Oracle Database global role when they access the database.

The following diagram illustrates the different types of assignments and mappings that are available.

Figure 5-1 Assignments and Mappings Between Azure AD and Oracle Database

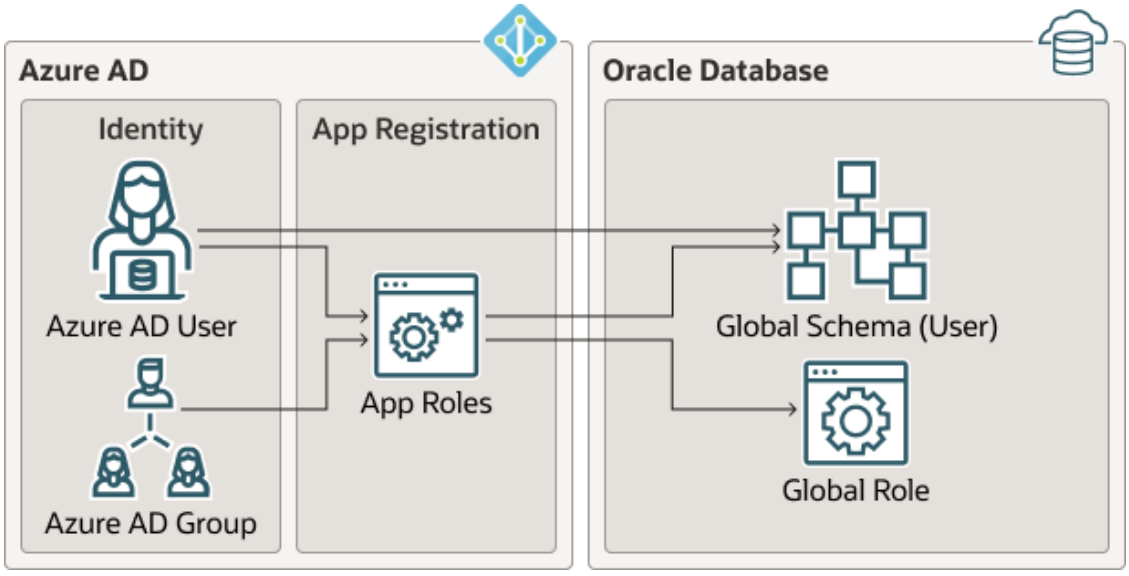

These mappings are as follows:

- An Azure AD user can be mapped directly to an Oracle Database global schema (user).
- An Azure AD user, Azure AD group, or application is assigned to an app role, which is then mapped to either an Oracle Database global schema (user) or a global role.

Use Cases for Connecting to an Oracle Database Using Azure AD

Oracle Database supports four types of use cases for connecting to an Oracle Database instance using Microsoft Azure Active Directory.

- **Connection using OAuth 2.0 authorization flow:** The client directs the resource owner to an authorization server, which in turn directs the resource owner back to the client with the authorization code. See the Microsoft Azure article Microsoft identity platform and OAuth 2.0 authorization code flow.
- **Connection using the resource owner password credentials:** The resource owner password credentials (that is, the user name and password) can be used directly to obtain an access token. Azure AD requires an additional client Id and a secret for this flow. (The secret is not required for public client.) See the Microsoft Azure article Microsoft identity platform and OAuth 2.0 Resource Owner Password Credentials.
- **Connection using the client credentials:** The client acts on its own behalf (the client is also the resource owner) or requests access to protected resources based on an authorization arranged with the authorization server. This flow is used to get the Azure `OAuth2` access token for the service principal. An application can also request an Azure AD `OAuth2` access token directly from Azure AD and pass it through a database client API. See the Microsoft Azure article Get Azure AD tokens by using a service principal.
- **Connection using on-behalf-of (OBO) token:** An Azure application requests an OBO token for a logged in user. The OBO token will also be an access token for the database with the Azure AD user identity and assigned app roles for the database. This enables the Azure AD user to log in to the database as the user and not the application. Only an application can request an OBO token for its Azure AD user and pass it to the database client through the API.

General Process of Integrating Microsoft Azure AD with Oracle Exadata Database Service on Dedicated Infrastructure

Both the Oracle and the Microsoft Azure administrators play roles in configuring the connection between Oracle Exadata Database Service on Dedicated Infrastructure and Microsoft Azure AD.

The general process is as follows:

1. The Oracle administrator ensures that the Oracle Database environment meets the requirements for the Microsoft Azure AD integration. See Oracle Database Requirements for the Microsoft Azure AD Integration.
2. The Oracle administrator registers the database instance with the Microsoft Azure AD tenancy and then enables the connection between the Oracle Exadata Database Service on Dedicated Infrastructure and the Azure AD endpoint.
   As part of the registration process, the Oracle administrator or the Azure administrator creates or designates Azure app roles to be used for the mappings between the Oracle database and the Microsoft Azure endpoint.
3. The Oracle administrator creates and maps global schemas to either an Azure AD user (exclusive schema mapping) or to an Azure app role (shared schema mapping). The Azure AD user or application must be mapped to one schema.
4. Optionally, the Oracle administrator creates and maps global Oracle Database roles to Azure app roles.

5. The Azure AD end user who wants to connect with the Oracle Exadata Database Service on Dedicated Infrastructure instance registers the client application as an Azure AD client (similar to how the Oracle database is registered).
   The Azure AD client will have a client identification and a client secret, unless the application client is public. If the application client is public, then only the application client identification is necessary.
6. The Azure AD end user (who can be a database administrator) connects using an utility such as PowerShell or the Azure command-line interface to retrieve the token and store it in a local file directory. An application can also request an Azure AD `OAuth2` access token directly from Azure AD and pass it through a database client API. Refer to the following Oracle Database client documentation for information about passing Azure AD `OAuth2` tokens:
   - JDBC-thin clients: *Oracle Database JDBC Developer's Guide*
   - Oracle Call Interface (OCI): *Oracle Call Interface Developer's Guide*
   - Oracle Data Provider for .NET (ODP): *Oracle Data Provider for .NET Developer's Guide.*Connecting to Oracle Database
7. Once connected to the Oracle Exadata Database Service on Dedicated Infrastructure instance, the Azure AD end user performs tasks as needed.

Configuring the Oracle Database for Microsoft Azure AD Integration

The Microsoft Azure AD integration with the Oracle Database instance requires the database to be registered with Azure AD so that the database can request the Azure AD public key.

- Prerequisites for Azure AD Authentication
  Before using Azure AD authentication on databases in the Exadata Cloud Infrastructure, you must use the Networking service to add a service gateway, a route rule, and an egress security rule to the Virtual Cloud Network (VCN) and subnets where your database resources reside.
- Configure TLS to Use Azure AD tokens
  When sending Azure AD tokens from the database client to the database server, a TLS connection must be established. The TLS wallet with the database certificate for the ExaDB-D service instance must be stored under the `WALLET_ROOT` location. Create a `tls` directory so it looks like: `WALLET_ROOT/<PDB GUID>/tls`.
- Oracle Database Requirements for the Microsoft Azure AD Integration
  Before you can configure an Oracle Database instance with Microsoft Azure AD, you must ensure that your environment meets special requirements.
- Registering the Oracle Database Instance with a Microsoft Azure AD Tenancy
  A user with administrator privileges uses Microsoft Azure AD to register the Oracle Database instance with the Microsoft Azure AD tenancy.
- Enabling Microsoft Entra ID v2 Access Tokens
  To enable the Microsoft Entra ID v2 access token, you must configure it to use the `upn` attribute from the Azure portal.
- Testing the Accessibility of the Azure Endpoint
  You must ensure that your Oracle Database instance can access the Azure AD endpoint.
- Managing App Roles in Microsoft Entra ID
  In Entra ID, you can create and manage app roles that will be assigned to Azure users and groups and also be mapped to Oracle Database global schemas and roles.

- Enabling Azure AD External Authentication for Oracle Database
  You can enable a Microsoft Azure AD external authentication with Oracle Database.
- Disabling Azure AD External Authentication for Oracle Database
  To disable Azure AD External authentication for an Oracle Database instance, you must set parameters with the `ALTER SYSTEM` statement.

Prerequisites for Azure AD Authentication

Before using Azure AD authentication on databases in the Exadata Cloud Infrastructure, you must use the Networking service to add a service gateway, a route rule, and an egress security rule to the Virtual Cloud Network (VCN) and subnets where your database resources reside.

1. Create a service gateway in the VCN where your database resources reside by following the instructions in *Task 1: Create the service gateway* in OCI documentation.
2. After creating the service gateway, add a route rule and an egress security rule to each subnet (in the VCN) where the database resources reside so that these resources can use the gateway to use Azure AD authentication:
   a. Go to the **Subnet Details** page for the subnet.
   b. In the **Subnet Information** tab, click the name of the subnet's Route Table to display its **Route Table Details** page.
   c. In the table of existing Route Rules, check whether there is already a rule with the following characteristics:
      - **Destination**: 0.0.0.0/0
      - **Target Type**: NAT Gateway
      - **Target**: The name of the NAT gateway you just created in the VCN

      If such a rule does not exist, click **Add Route Rules** and add a route rule with these characteristics.
   d. Return to the **Subnet Details** page for the subnet.
   e. In the subnet's Security Lists table, click the name of the subnet's security list to display its **Security List Details** page.
   f. In the side menu, under **Resources**, click **Egress Rules**.
   g. In the table of existing Egress Rules, check whether there is already a rule with the following characteristics:
      - **Destination Type**: CIDR
      - **Destination**: 0.0.0.0/0
      - **IP Protocol**: TCP
      - **Source Port Range**: 443
      - **Destination Port Range**: All
   h. If such a rule does not exist, click **Add Egress Rules** and add an egress rule with these characteristics.

Related Topics

- Task 1: Create the service gateway

Configure TLS to Use Azure AD tokens

When sending Azure AD tokens from the database client to the database server, a TLS connection must be established. The TLS wallet with the database certificate for the ExaDB-D service instance must be stored under the `WALLET_ROOT` location. Create a `tls` directory so it looks like: `WALLET_ROOT/<PDB GUID>/tls`.

When configuring TLS between the database client and server there are several options to consider.

- Using a self-signed database server certificate vs a database server certificate signed by a commonly known certificate authority
- One-way TLS (TLS) vs Mutual or two-way TLS (mTLS)
- Client with or without a wallet

Self-Signed Certificate

Using a self-signed certificate is a common practice for internally facing IT resources since you can create these yourself and it's free. The resource (in our case, the database server) will have a self-signed certificate to authenticate itself to the database client. The self-signed certificate and root certificate will be stored in the database server wallet. For the database client to be able to recognize the database server certificate, a copy of the root certificate will also be needed on the client. This self-created root certificate can be stored in a client-side wallet or installed in the client system default certificate store (Windows and Linux only). When the session is established, the database client will check to see that the certificate sent over by the database server has been signed by the same root certificate.

A Well-Known Certificate Authority

Using a commonly known root certificate authority has some advantages in that the root certificate is most likely already stored in the client system default certificate store. There is no extra step for the client to store the root certificate if it is a common root certificate. The disadvantage is that this normally has a cost associated with it.

One-Way TLS

In the standard TLS session, only the server provides a certificate to the client to authenticate itself. The client doesn't need to have a separate client certificate to authenticate itself to the server (similar to how HTTPS sessions are established). While the database requires a wallet to store the server certificate, the only thing the client needs to have is the root certificate used to sign the server certificate.

Two-Way TLS (also called Mutual TLS, mTLS)

In mTLS, both the client and server have identity certificates that are presented to each other. In most cases, the same root certificate will have signed both of these certificates so the same root certificate can be used with the database server and client to authenticate the other certificate. mTLS is sometimes used to authenticate the user since the user identity is authenticated by the database server through the certificate. This is not necessary for passing IAM tokens but can be used when passing IAM tokens.

Client with a Wallet

A client wallet is mandatory when using mTLS to store the client certificate. However, the root certificate can be stored either in the same wallet or in the system default certificate store.

A Client without a Wallet

Clients can be configured without a wallet when using TLS under these conditions: 1) One-way TLS is being configured where the client does not have its own certificate and 2) the root certificate that signed the database server certificate is stored in the system default certificate store. The root certificate would most likely already be there if the server certificate is signed by a common certificate authority. If it's a self-signed certificate, then the root certificate would need to be installed in the system default certificate store to avoid using a client wallet.

For details on how to configure TLS between the database client and database server including the options described above, see *Configuring Transport Layer Security Authentication* in the *Oracle Database Security Guide*.

If you choose to use self-signed certificates and for additional wallet related tasks, see *Managing Public Key Infrastructure (PKI) Elements* in the *Oracle Database Security Guide*.

Related Topics

- Configuring Transport Layer Security Authentication
- Managing Public Key Infrastructure (PKI) Elements

Oracle Database Requirements for the Microsoft Azure AD Integration

Before you can configure an Oracle Database instance with Microsoft Azure AD, you must ensure that your environment meets special requirements.

The Microsoft Azure AD integration with the Oracle Exadata Database on Dedicated Infrastructure requires:

- The ExaCS Database to be version 19.17 or higher.
- Connectivity to the database using TLS. Non TLS connections are not supported.
- Outbound network connectivity to Azure AD so that the database can request the Azure AD public key.
- The ExaDB-D Database to be registered with Azure AD.

Note the following:

- The Oracle Database server must be able to request the Azure AD public key. Depending on the enterprise network connectivity setup, you may need to configure a proxy setting.
- Users and applications that need to request an Azure AD token must also be able to have network connectivity to Azure AD. You may need to configure a proxy setting for the connection.
- You must configure Transport Layer Security (TLS) between the Oracle Database client and the Oracle Database server so that the token can be transported securely. This TLS connection can be either one-way or mutual.
- You can create the TLS server certificate to be self-signed or be signed by a well known certificate authority. The advantage of using a certificate that is signed by a well known Certificate Authority (CA) is that the database client can use the system default certificate store to validate the Oracle Database server certificate instead of having to create and maintain a local wallet with the root certificate. Note that this applies to Linux and Windows clients only.

Related Topics

-

Registering the Oracle Database Instance with a Microsoft Azure AD Tenancy

A user with administrator privileges uses Microsoft Azure AD to register the Oracle Database instance with the Microsoft Azure AD tenancy.

1. Log in to the Azure portal as an administrator who has Microsoft Azure AD privileges to register applications.
2. In the Azure Active directory admin center page, from the left navigation bar, select **Azure Active Directory**.
3. In the MS - App registrations page, select **App registrations** from the left navigation bar.
4. Select **New registration**.

   The Register an application window appears.

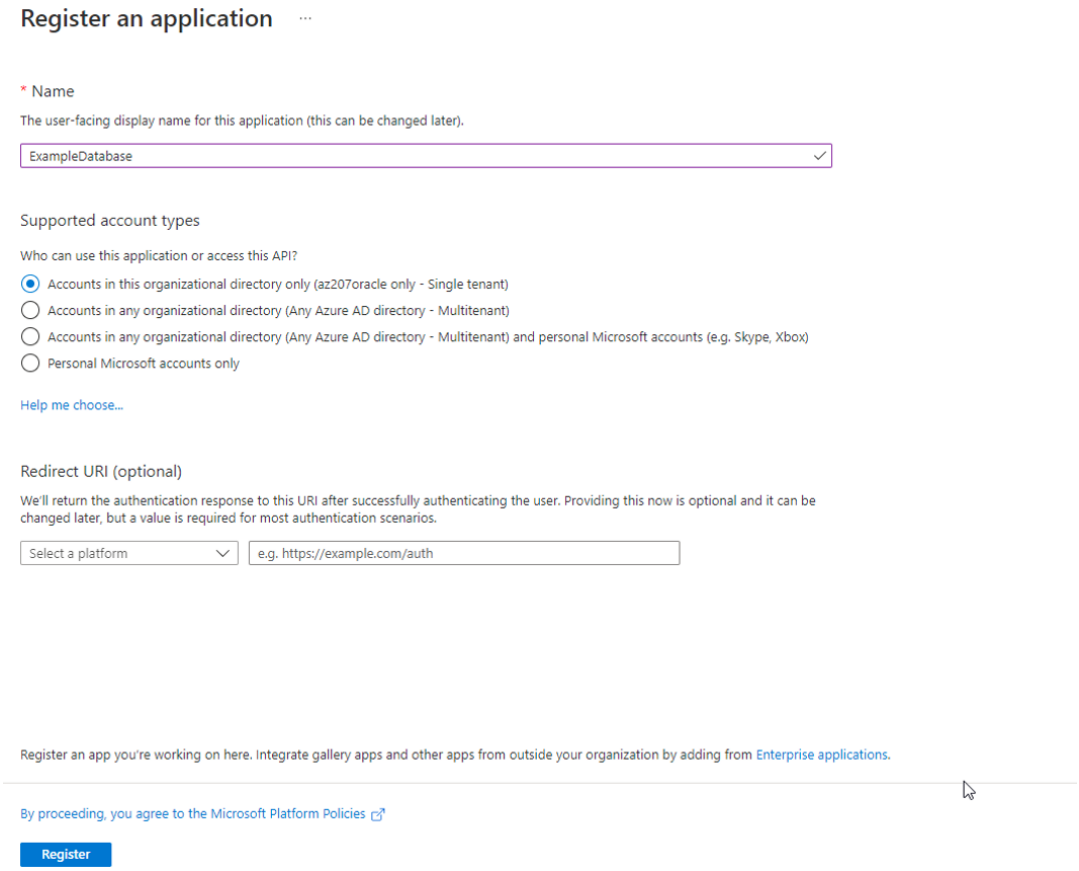

5. In the Register an application page, enter the following Oracle Database instance registration information:
   - In the **Name** field, enter a name for the Oracle Database instance connection (for example, `Example Database`).
   - Under Supported account types, select the account type that matches your use case.
     - **Accounts in this organizational directory only (*tenant_name* only - Single tenant)**
     - **Accounts in any organizational directory (Any Azure AD directory - Multitenant)**

- **Accounts in any organizational directory (Any Azure AD directory - Multitenant) and personal Microsoft accounts (e.g. Skype, Xbox)**
- **Personal Microsoft accounts only**

6. Bypass the Redirect URI (Optional) settings. You do not need to create a redirect URI.
7. Click **Register**.

   After you click **Register**, Azure AD displays the app registration's Overview pane, which will show the Application (client) ID under Essentials. This value is a unique identifier for the application in the Microsoft identity platform.
8. Register a scope, which will set the permission for the registered app.
   a. In the left navigation bar, select **Expose an API**.
   b. Under Set the App ID URI, in the **Application ID URI** field, enter the app ID URI for the database connection using the following format, and then click **Save**:

      your_tenancy_url/application_(client)_id

      In this specification:

      - *your_tenancy_url* must include `https` as the prefix and the fully qualified domain name of your Azure AD tenancy.
      - *application_(client)_id* is the ID that was generated when you registered the Oracle Database instance with Azure AD. It is displayed in the Overview pane of the app registration.

      For example:

      ```
      https://sales_west.example.com/1aa11111-1a1z-1a11-1a1a-11aa11a1aa1a
      ```

   c. Select **Add a scope** and then enter the following settings:

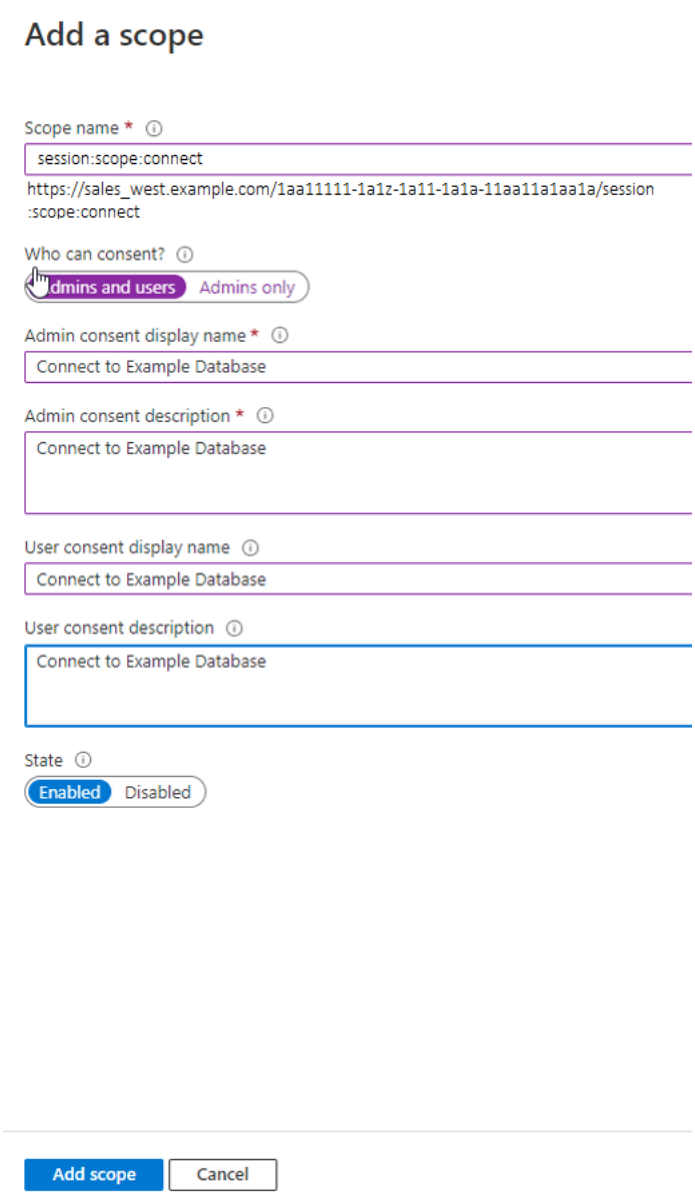

- **Scope name** specifies a name for the scope. Enter the following name:

  ```
  session:scope:connect
  ```

  This name can be any text. However, a scope name must be provided. You will need to use this scope name later when you give consent to the database client application to access the database.
- **Who can consent** specifies the necessary permissions. Select **Admins and users**, or for higher restrictions, **Admins only**.
- **Admin consent** display name describes the scope's purpose (for example, `Connect to Oracle`), which only administrators can see.
- **Admin consent display name** describes the scope's purpose (for example, `Connect to Example Database`), which only administrators can see.

- **User consent display name** is a short description of the purpose of the scope (for example, `Connect to Example Database`), which users can see if you specify **Admins and users** in **Who can consent**.
- **User consent description** is a more detailed description of the purpose of the scope (for example, `Connect to Example Database`), which users can see if you specify **Admins and users** in **Who can consent**.
- **State** enables or disables the connection. Select **Enabled**.

After you complete these steps, you are ready to add one or more Azure app roles, and then perform the mappings of Oracle schemas and roles.

Related Topics

- Quickstart: Register an application with the Microsoft identity platform

Enabling Microsoft Entra ID v2 Access Tokens

To enable the Microsoft Entra ID v2 access token, you must configure it to use the `upn` attribute from the Azure portal.

Oracle Database supports the Entra ID v2 token as well as the default v1 token. However, to use the Entra ID v2 token, you must perform some additional steps to ensure it works with the Oracle Database.

1. Check the version of the Entra ID access token that you are using.
2. Log in to the Microsoft Entra ID portal.
3. Search for and select **Entra ID**.
4. Under **Manage**, select **App registrations**.
5. Choose the application for which you want to configure optional claims based on your scenario and desired outcome.
6. Under **Manage**, select **Token configuration**.
7. Click **Add optional claim** and select **upn**.

- Checking the Entra ID Access Token Version
  You can check the version of the Entra ID access token that your site uses by using the JSON Web Tokens web site.

Related Topics

- Checking the Entra ID Access Token Version
  You can check the version of the Entra ID access token that your site uses by using the JSON Web Tokens web site.

Checking the Entra ID Access Token Version

You can check the version of the Entra ID access token that your site uses by using the JSON Web Tokens web site.

By default, Entra ID v1 access token, but your site may have chosen to use v2. Oracle Database supports v1 tokens and Autonomous Database Serverless supports v2 tokens, as well. If you want to use the v2 access tokens, then you can enable their use for the Oracle database. To find the version of the Entra ID access token that you are using, you can either check with your Entra ID administrator, or confirm the version from the JSON Web Tokens website, as follows.

1. Go to the JSON Web Tokens website.

   ```
   https://jwt.io/
   ```

2. Copy and paste the token string into the **Encoded** field.
3. Check the **Decoded** field, which displays information about the token string.

   Near or at the bottom of the field, you will see a claim entitled `ver`, which indicates either of the following versions:

   - `"ver": "1.0"`
   - `"ver": "2.0"`

Related Topics

- Enabling Microsoft Entra ID v2 Access Tokens
  To enable the Microsoft Entra ID v2 access token, you must configure it to use the `upn` attribute from the Azure portal.

Testing the Accessibility of the Azure Endpoint

You must ensure that your Oracle Database instance can access the Azure AD endpoint.

For an Oracle database to accept Azure AD `OAuth2` tokens, the database must request the public key from the Azure AD endpoint.

- Run the following test to determine if the database can connect with the Azure AD endpoint:

```
SET SERVEROUTPUT ON SIZE 40000
DECLARE
  req UTL_HTTP.REQ;
  resp UTL_HTTP.RESP;
BEGIN
  UTL_HTTP.SET_WALLET(path => 'system:');
  req := UTL_HTTP.BEGIN_REQUEST('https://login.windows.net/common/
discovery/keys');
  resp := UTL_HTTP.GET_RESPONSE(req);
  DBMS_OUTPUT.PUT_LINE('HTTP response status code: ' || resp.status_code);
  UTL_HTTP.END_RESPONSE(resp);
END;
/
```

If this test is successful, then a `PL/SQL procedure successfully completed` message appears.

If the following messages appear, then it means that a database network access control list (ACL) policy blocked your test and you will need to temporarily set an access control list policy to allow you to test this:

```
ORA-29273: HTTP request failed
ORA-24247: network access denied by access control list (ACL)
```

1. Set the ACL as follows:

```
BEGIN
DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
  host => '*',
  ace  =>  xs$ace_type(privilege_list => xs$name_list('connect'),
                       principal_name => 'username_placeholder',
                       principal_type => xs_acl.ptype_db));
END;
/
```

   Replace `username_placeholder` with the user name of the database user who is running the test. For example:

```
BEGIN
DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
  host => '*',
  ace  =>  xs$ace_type(privilege_list => xs$name_list('connect'),
                       principal_name => 'DBA_DEBRA',
                       principal_type => xs_acl.ptype_db));
END;
/
```

2. Try running the test again.
3. Remove the ACL, because you now no longer need it. For example, assuming your user name is `dba_debra`:

```
BEGIN
DBMS_NETWORK_ACL_ADMIN.REMOVE_HOST_ACE(
  host => '*',
  ace  =>  xs$ace_type(privilege_list => xs$name_list('connect'),
                       principal_name => 'DBA_DEBRA',
                       principal_type => xs_acl.ptype_db));
END;
/
```

If the database cannot connect with the Azure AD endpoint, even after you set the ACL policy, you will most likely need to set the `HTTP_PROXY` package for your database. Review the topics listed in Related Topics, depending if you are using a default Oracle Database environment or an Oracle Real Application Clusters RAC environment. Your network administrator should be able to tell you what the correct `HTTP_PROXY` setting should be.

Related Topics

- Creating the Network Proxy for the Default Oracle Database Environment
- Creating the Network Proxy for an Oracle Real Application Clusters Environment

Managing App Roles in Microsoft Entra ID

In Entra ID, you can create and manage app roles that will be assigned to Azure users and groups and also be mapped to Oracle Database global schemas and roles.

- Creating a Microsoft Azure AD App Role
  Azure AD users, groups, and applications will be assigned to the app roles.

- Assigning Users and Groups to the Microsoft Azure AD App Role
  Before Microsoft Azure AD users can have access to the Oracle Database instance, they must first be assigned to the app roles that will be mapped to Oracle Database schema users or roles.
- Assigning an Application to an App Role
  You can assign an Azure AD client application to a app role.

Creating a Microsoft Azure AD App Role

Azure AD users, groups, and applications will be assigned to the app roles.

See the Microsoft Azure article Create and assign a custom role in Azure Active Directory for detailed steps on how to create an app role. The following steps describe how to create the app role for use with an Oracle Database integration.

1. Log in to Azure AD as an administrator who has privileges for creating app roles.
2. Access the Oracle Database app registration that you created.
   a. Use the **Directory + subscription** filter to locate the Azure Active Directory tenant that contains the Oracle Database app registration.
   b. Select **Azure Active Directory**.
   c. Under **Manage**, select **App registrations**, and then select the Oracle Database instance that you registered earlier.
3. Under **Manage**, select **App roles**.
4. In the App roles page, select **Create app role**.
5. In the Create app role page, enter the following information:
   - **Display name** is the displayed name of the role (for example, `HR App Schema`). You can include spaces in this name.
   - **Value** is the actual name of the role (for example, `HR_APP`). Ensure that this setting matches exactly the string that is referenced in the application's code. Do not include spaces in this name.
   - **Description** provides a description of the purpose of this role.
   - **Do you want to enable this app role?** enables you to activate the role.
6. Click **Apply**.

   The app role appears in the App roles pane.

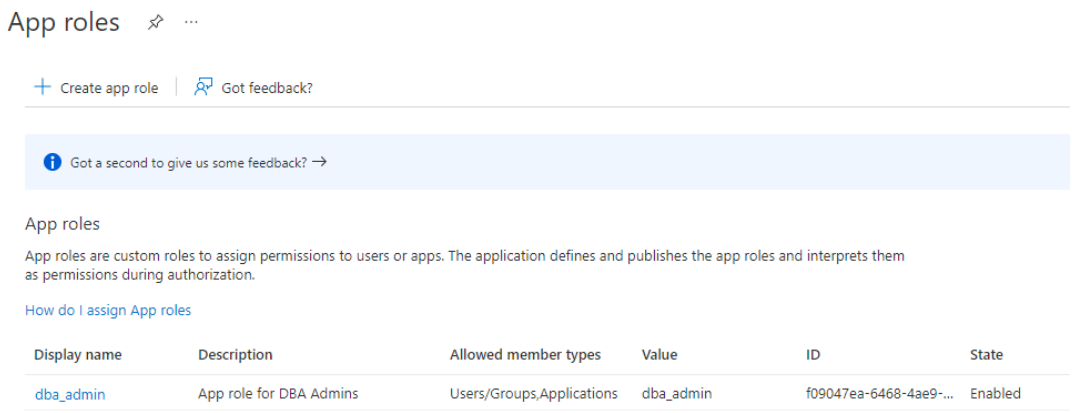

Assigning Users and Groups to the Microsoft Azure AD App Role

Before Microsoft Azure AD users can have access to the Oracle Database instance, they must first be assigned to the app roles that will be mapped to Oracle Database schema users or roles.

See the Microsoft Azure article Add app roles to your application and receive them in the token for detailed steps assigning users and groups to an app role. The following steps explain how to do this for an Oracle Database integration.

1. Log in to Azure AD as an administrator who has privileges for assigning Azure AD users and groups to app roles.
2. In Enterprise applications, access the Oracle Database application that you registered.
    a. Use the **Directory + subscription** filter to locate the Azure Active Directory tenant that contains the Oracle connection.
    b. Select **Azure Active Directory**.
    c. Under **Manage**, select **Enterprise applications**, and then select the Oracle Database application name that you registered earlier.
3. Under Getting Started, select **Assign users and groups**.
4. Select **Add user/group**.
5. In the Add assignment window, select **Users and groups** to display a list of users and security groups.
6. From this list, select the users and groups that you want to add to the app role, and then click **Select**.
7. In the Add assignment window, select **Select a role** to display a list of the app roles that you have created.
8. Select the app role and then select **Select**.
9. Click **Assign**.

Assigning an Application to an App Role

You can assign an Azure AD client application to a app role.

1. Log in to Azure AD as an administrator who has privileges for assigning Azure AD users and groups to app roles.
2. Access the app registration for the application.
3. Under Manage, select **API permissions**.
4. In the Configured permissions area, select **+ Add a permission**.
5. In the Request API permission pane, select the **My APIs** tab.
6. Select the Oracle Database app that you want to give permission for this application to access. Then select the **Application permissions** option.
7. Select the database app roles to assign to the application and then click the **Add Permission** box at the bottom of the screen to assign the app roles and close the dialog box. Ensure that the app roles that you just assigned appear under Configured permissions.

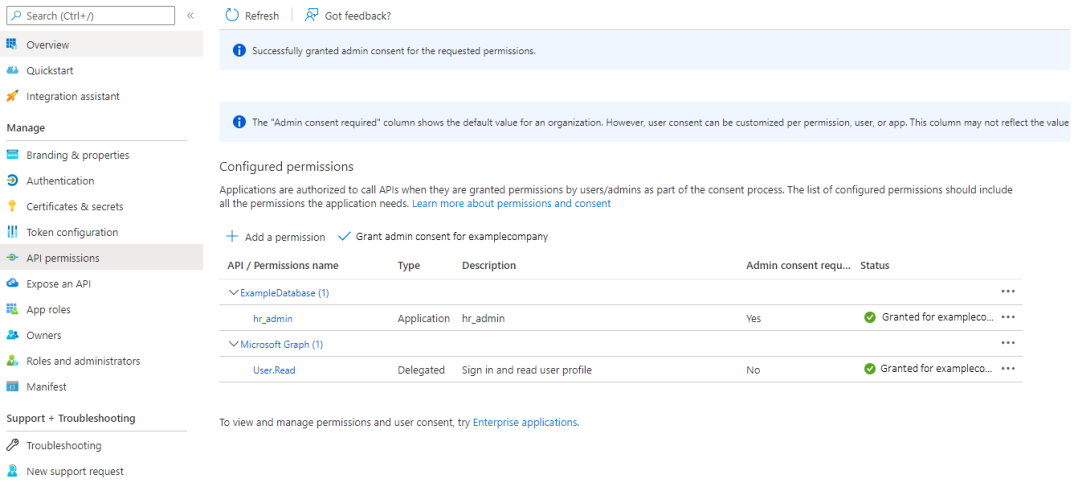

8. Select **Grant admin consent for *tenancy*** to grant consent for the tenancy users, then select **Yes** in the confirmation dialog box.

Related Topics

- Configure the admin consent workflow

Enabling Azure AD External Authentication for Oracle Database

You can enable a Microsoft Azure AD external authentication with Oracle Database.

1. Log in to the Oracle Database instance as a user who has been granted the `ALTER SYSTEM` system privilege.
2. Set the `IDENTITY_PROVIDER_TYPE` parameter as follows:

```
ALTER SYSTEM SET IDENTITY_PROVIDER_TYPE=AZURE_AD SCOPE=BOTH;
```

3. Ensure that you set the `IDENTITY_PROVIDER_TYPE` parameter correctly.

```
SELECT NAME, VALUE FROM V$PARAMETER WHERE NAME='identity_provider_type';
```

The following output should appear:

```
NAME                      VALUE
------------------------  -------
identity_provider_type    AZURE_AD
```

4. Set the `IDENTITY_PROVIDER_CONFIG` parameter by using the following syntax:

```
ALTER SYSTEM SET IDENTITY_PROVIDER_CONFIG =
{
   application_id_uri : string , // from registered app, to be mapped in
jwt "aud" claim;
                                 // Domain qualified to support cross
tenancy resource access
   tenant_id : string,           // from tenant config
   app_id: string                // from registered resource app
}SCOPE=BOTH;
```

For example:

```
ALTER SYSTEM SET IDENTITY_PROVIDER_CONFIG =
{
  "application_id_uri" : "https://www.example.com/11aa1a11-
aaaa-1111-1111-1111aa11111",
  "tenant_id" : 111a1111-a11a-111a-1a1a-1111111111a,
  "app_id" : 11aa1a11-aaaa-1111-1111-1111aa11111
}SCOPE=BOTH;
```

Disabling Azure AD External Authentication for Oracle Database

To disable Azure AD External authentication for an Oracle Database instance, you must set parameters with the `ALTER SYSTEM` statement.

1. Log in to the Oracle Database instance as a user who has been granted the `ALTER SYSTEM` system privilege.
2. Set the identity provider parameters as follows:

```
ALTER SYSTEM RESET IDENTITY_PROVIDER_CONFIG SCOPE=BOTH;
ALTER SYSTEM RESET IDENTITY_PROVIDER_TYPE SCOPE=BOTH;
```

Mapping Oracle Database Schemas and Roles

Azure AD users will be mapped to one database schema and optionally to one or more database roles.

- Exclusively Mapping an Oracle Database Schema to a Microsoft Azure AD User
  You can exclusively map an Oracle Database schema to a Microsoft Azure AD user.
- Mapping a Shared Oracle Schema to an App Role
  In this mapping, an Oracle schema is mapped to an app role.
- Mapping an Oracle Database Global Role to an App Role
  Oracle Database global roles that are mapped to Azure app roles give Azure users and applications additional privileges and roles above those that they have been granted through their login schemas.

Exclusively Mapping an Oracle Database Schema to a Microsoft Azure AD User

You can exclusively map an Oracle Database schema to a Microsoft Azure AD user.

1. Log in to the Oracle Database instance as a user who has been granted the `CREATE USER` or `ALTER USER` system privilege.
2. Run the `CREATE USER` or `ALTER USER` statement with the `IDENTIFIED GLOBALLY AS` clause specifying the Azure AD user name.

   For example, to create a new database schema user named `peter_fitch` and map this user to an existing Azure AD user named `peter.fitch@example.com`:

```
CREATE USER peter_fitch IDENTIFIED GLOBALLY AS
'AZURE_USER=peter.fitch@example.com';
```

3. Grant the `CREATE SESSION` privilege to the user.

```
GRANT CREATE SESSION TO peter_fitch;
```

Mapping a Shared Oracle Schema to an App Role

In this mapping, an Oracle schema is mapped to an app role.

1. Log in to the Oracle Database instance as a user who has the `CREATE USER` or `ALTER USER` system privilege.
2. Run the `CREATE USER` or `ALTER USER` statement with the `IDENTIFIED GLOBALLY AS` clause specifying the Azure application role name.

   For example, to create a new database global user account (schema) named `dba_azure` and map it to an existing Azure AD application role named `AZURE_DBA`:

```
CREATE USER dba_azure IDENTIFIED GLOBALLY AS 'AZURE_ROLE=AZURE_DBA';
```

Mapping an Oracle Database Global Role to an App Role

Oracle Database global roles that are mapped to Azure app roles give Azure users and applications additional privileges and roles above those that they have been granted through their login schemas.

1. Log in to the Oracle Database instance as a user who has been granted the `CREATE ROLE` or `ALTER ROLE` system privilege
2. Run the `CREATE ROLE` or `ALTER ROLE` statement with the `IDENTIFIED GLOBALLY AS` clause specifying the name of the Azure AD application role.

   For example, to create a new database global role named `widget_sales_role` and map it to an existing Azure AD application role named `WidgetManagerGroup`:

```
CREATE ROLE widget_sales_role IDENTIFIED GLOBALLY AS
'AZURE_ROLE=WidgetManagerGroup';
```

Configuring Azure AD Client Connections to the Oracle Database

You can configure client connections to connect with the Azure AD registered database

- About Configuring Client Connections to Azure ADs
  There are numerous ways that you can configure a client to connect with an Oracle Database instance using Azure AD tokens.
- Supported Client Drivers for Azure AD Connections
  Oracle Database supports several types of client drivers for Azure AD connections.
- Operational Flow for SQL*Plus Client Connection in PowerShell to Oracle Database
  The connection between the Azure user, Azure AD, and the Oracle Database instance relies on the passing of the `OAuth2` token throughout these components.
- Registering a Client with Azure AD Application Registration
  This type of registration is similar to registering Oracle Database with Azure AD app registration.

- Examples of Retrieving Azure AD OAuth2 Tokens
  These examples show different ways that you can retrieve Azure AD `OAuth2` tokens.
- Configuring SQL*Plus for Azure AD Access Tokens
  You must configure SQL*Plus to retrieve the Azure AD database access token from a location and use it when the `/` slash login is used.

About Configuring Client Connections to Azure ADs

There are numerous ways that you can configure a client to connect with an Oracle Database instance using Azure AD tokens.

You should choose the client connection method that works best with your environment. This guide provides examples of connecting SQL*Plus with different methods of getting an Azure AD OAuth2 access token. All Oracle Database clients version 19.16 and above can accept a token that is passed as a file. The JDBC-thin, Instant Client, and ODP.net drivers also accept the token through the database client API from an application. Oracle Database tools such as SQL*Plus cannot retrieve the tokens directly, so tools such as PowerShell or Azure CLI must be used to retrieve the Azure AD OAuth2 access token. To retrieve an Azure AD token, the client must be registered through the Azure AD app registration process. Registering the client is similar to registering the Oracle Database server with Azure AD using app registration. Both the database and client must be registered with Azure AD.

The database must be registered so the client can get permission to get an access token for the database. The client must be registered so that Azure AD can recognize a trusted client is asking for an access token.

See the following Microsoft Azure articles for more information about connecting clients to Azure AD:

- *Quickstart: Configure a client application to access a web API*
- *Choose the right Azure command-line tool*
- *Get Azure AD tokens by using the Microsoft Authentication Library*
- *Install the Azure CLI on Linux*

Related Topics

- Quickstart: Configure a client application to access a web API
- Choose the right Azure command-line tool
- Get Azure AD tokens by using the Microsoft Authentication Library
- Install the Azure CLI on Linux

Supported Client Drivers for Azure AD Connections

Oracle Database supports several types of client drivers for Azure AD connections.

- JDBC-thin: Oracle Database 19.16 (July 2022), Oracle Database 21.8 (October 2022)
- OCI (C driver): Oracle Database 19.16 (July 2022)
- Oracle Instant Client based on OCI
- Oracle Data Provider (core): Oracle Database 19.16, Oracle Database 21.7
- Oracle Data Provider (unmanaged): based on OCI
- Oracle Data Provider (managed): Oracle Database 19.16, Oracle Database 21.7

- All other drivers built on OCI adopts the OCI compatibility

Operational Flow for SQL*Plus Client Connection in PowerShell to Oracle Database

The connection between the Azure user, Azure AD, and the Oracle Database instance relies on the passing of the `OAuth2` token throughout these components.

This example shows the use of the Resource Owner Password Credential (ROPC) flow with a public client. See the Microsoft Azure article Microsoft identity platform and OAuth 2.0 Resource Owner Password Credentials for detailed information about ROPC.

Figure 5-2 ROPC Operational Flow with a Public Client

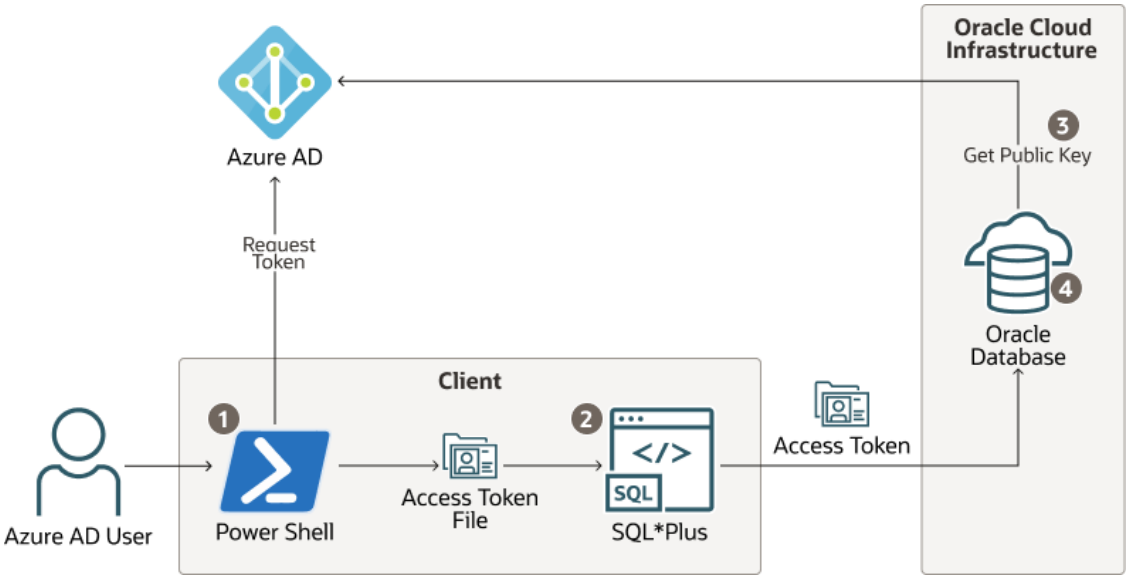

1. The Azure user requests an Azure AD access token for the database in PowerShell and the returned token is written into a file called `token` at a file location.
2. The Azure user connects to the database using `/` slash login. Either the `sqlnet.ora` or `tnsnames.ora` connection string tells the instant client that an Azure AD `OAuth2` token is needed and to retrieve it from a specified file location. The access token is sent to the database.
3. The database verifies that the access token came from Azure AD (using the Azure AD public key) and then checks the token for additional claims.
4. The database finds the schema mapping (exclusive or shared) and creates the session. The database will also grant any global roles that the Azure user is also assigned to through an app role.

Registering a Client with Azure AD Application Registration

This type of registration is similar to registering Oracle Database with Azure AD app registration.

- Confidential and Public Client Registration
  You can register the database client with Azure as either confidential or public depending on your use case.

- Registering a Database Client App with Azure AD
  Creating the client app registration is similar to creating the Oracle Database instance with the Microsoft Azure AD tenancy.

Confidential and Public Client Registration

You can register the database client with Azure as either confidential or public depending on your use case.

See the Microsoft Azure article Authentication flows and application scenarios for detailed information about authentication flows and application scenarios.

Registering a confidential client app requires that the client have a secret, in addition to the client ID. The confidential client app uses both the client ID and the secret when it makes Azure AD requests. However, in an enterprise, it is not practical for every SQL*Plus and SQLcl user to create a separate app registration with its own secret. In addition, a secret is no longer a secret when you start to share it within an organization. It is far better to just create a public client app. A public client app does not have a secret; it only has a client ID. All database tool users can use the public client ID when they connect to Azure AD to get an access token. The Azure AD user still needs to authenticate to Azure AD with their own user credential.

Registering a Database Client App with Azure AD

Creating the client app registration is similar to creating the Oracle Database instance with the Microsoft Azure AD tenancy.

1. Log in to the Azure portal as an administrator who has Microsoft Azure AD privileges to register applications.
2. In the Azure Active directory admin center page, from the left navigation bar, select **Azure Active Directory**.
3. In the MS - App registrations page, select **App registrations** from the left navigation bar.
4. Select **New registration**.
5. In the Register an application page, enter the following Oracle Database instance registration information:
   - In the **Name** field, enter a name for the client app (for example, `DatabaseClientApplication`)..
   - Under Supported account types, select the account type that matches your use case.
     - **Accounts in this organizational directory only (*tenant_name* only - Single tenant)**
     - **Accounts in any organizational directory (Any Azure AD directory - Multitenant)**
     - **Accounts in any organizational directory (Any Azure AD directory - Multitenant) and personal Microsoft accounts (e.g. Skype, Xbox)**
     - **Personal Microsoft accounts only**
6. Under Redirect URI (optional), configure the redirect URI for the client app.

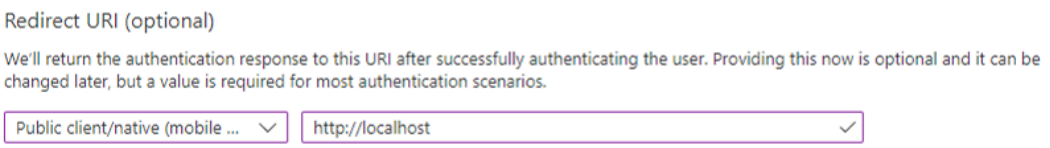

- Select **Public client/native (mobile & desktop)**, **Web**, or **Single-page application (SPA)**. Choose **Public client** if this client app will be used by multiple users such as database administrators who need to use SQL*Plus to access the Oracle Database instance.
- Add a redirect URI of `http://localhost`, unless you have another address to use. This redirect URI is needed for the authorization flow.

7. Click **Register**.

   At this stage, the database client has been registered with Azure AD. Next, you must add the new client to the list of authorized client apps for the Oracle Database instance.

8. To add the new client to this list of client apps, do the following:

   a. Make a note of the new client's Application (client) ID. This ID is in the Overview page for the app.

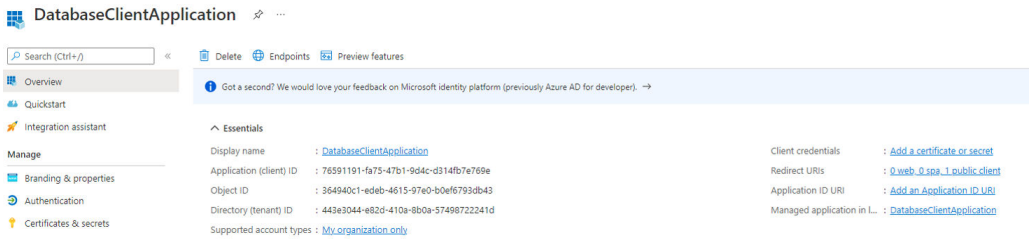

   b. On the App registrations page, open the app registration page for the database server by selecting it from the menu.

   c. On the left side, select **Expose an API**.

   d. Scroll down on the main page until you see **Authorized client applications**.

   e. Select **+** to add a client application.

   f. Copy the new client's Application (client) ID to the **Client Id** field.

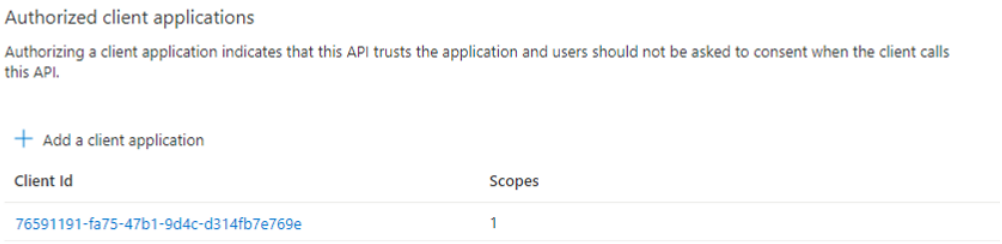

   g. Click **Add application**.

Related Topics

- Quickstart: Register an application with the Microsoft identity platform

Examples of Retrieving Azure AD OAuth2 Tokens

These examples show different ways that you can retrieve Azure AD `OAuth2` tokens.

- Example: Using PowerShell to Get a Token Using Resource Owner Password Credentials
  This example shows how to use PowerShell to get an Azure AD access token by using Resource Owner Password Credentials (ROPC).

- Example: Using Python with Microsoft Authentication Library Using an Authorization Flow
  Because this example with the Microsoft Authentication Library (MSAL) is in Python, it can be run on a variety of platforms such as PowerShell and Linux.
- Example: Using Curl with a Resource Owner Password Credential Flow
  This example shows how to use the `curl` command against the Azure AD API uses a Resource Owner Password Credential (ROPC) flow with a public Azure AD client.
- Example: Azure CLI Using Authorization Flow
  This example shows how to use the Azure CLI to retrieve an access token and then write the token to a file.

Example: Using PowerShell to Get a Token Using Resource Owner Password Credentials

This example shows how to use PowerShell to get an Azure AD access token by using Resource Owner Password Credentials (ROPC).

You can retrieve the `OAuth2` access token by making a REST call from PowerShell. This configuration requires several values that were generated or that you specified when you registered the Oracle Database instance with Azure AD.

1. If necessary install the Azure Active Directory PowerShell module.

   Follow the instructions in the Microsoft article Install the Azure Az PowerShell module to download and install Azure PowerShell. It takes about 20 minutes or longer to perform the installation. You may want to set debug options for Azure PowerShell so that you can see how the installation is progressing.

2. After the Azure PowerShell installation is complete, log in to PowerShell and then set the following variables in the order shown.

   a. `$TenantDomain = "`*`user_tenancy_domain_name`*`"`

      This value is the tenancy domain name. For example:

      ```
      $TenantDomain = "example.com"
      ```

   b. `$AppClientId ="`*`application_client_id`*`"`

      This value sets the application client ID for the database client, not the database server. This is the **Application (client) ID** value in the app registration's Overview pane. For example:

      ```
      $AppClientId ="111a1a1a-aa1a-1a1a-11aa-1a11111111aa"
      ```

   c. `$Username = "`*`user_name`*`"`, which

      This value is the name of the Azure user who wants to access the Oracle Database instance. For example:

      ```
      $Username = "peter.fitch@example.com"
      ```

   d. Entering the user password in PowerShell scripts depends on your corporate or personal security standards. Use your own method to capture the password securely or use this example that hides the password from the command history and command line window. You should delete the password variables after they have been used. Enter these in the order shown:

      i. `$securePassword = Read-Host " Enter Password" -AsSecureString`

ii. `$Password = [System.Runtime.InteropServices.Marshal]::PtrToStringAuto([System.Runtime.InteropServices.Marshal]::SecureStringToBSTR($securePassword))`

e. `$Scope = "database_app_id_uri/scope"`

This value sets the app ID URI for the database and the scope (permission) for the database, separated by a / slash. These values can be found in the database app registration Expose an API page. In the following example, `https://example.com/111aa1aa-1111-1111-a1a1-1a11a111111a` is the app ID URI and `session:scope:connect` is the scope.

```
$Scope = "https://example.com/111aa1aa-1111-1111-a1a1-1a11a111111a/session:scope:connect"
```

f. `$requestBody = @{client_id=$AppClientId;grant_type="password";username=$Username;password=$Password;scope=$Scope;}`

This is the request body of the upcoming REST call.

g. `$OAuthResponse = Invoke-RestMethod -Method Post -Uri https://login.microsoftonline.com/$TenantDomain/oauth2/v2.0/token -Body $requestBody`

This gets the `OAuth2` access token for the user.

h. You can remove the password from the variables if you do not need them.

- `$securePassword = $null`, for secure password strings
- `$Password = $null`, for clear text password strings

i. `$AccessToken = $OAuthResponse.access_token | Out-File -FilePath .\token -Encoding ASCII`, which writes the `OAuth2` token to the current file location using ASCII encoding

3. Optionally, because the Azure AD `OAuth2` access token is a JSON Web Token (JWT) formatted token, you can view the cleartext of the encoded content by copying and pasting the token content into the website:

https://jwt.io/

Note the following:

- The default PowerShell UTF16 file encoding cannot be used for the token. Use ASCII encoding as an alternative.
- Tokens may not work cross-platform (for example, Windows to Linux or Linux to Windows), depending on encoding changes to the file when it is moved.

At this stage, the `OAuth2` access token has been retrieved and stored as a file. The next step is to enable the SQL*Plus client to use the store access token and send it to the database.

Example: Using Python with Microsoft Authentication Library Using an Authorization Flow

Because this example with the Microsoft Authentication Library (MSAL) is in Python, it can be run on a variety of platforms such as PowerShell and Linux.

When multi-factor authentication is enabled for the user, an `OAuth2` authorization flow is necessary for a user to add the second authentication. Because the authorization flow requires two round trips to Azure AD, it is best handled using the MSAL. See the Microsoft article Get Azure AD tokens by using the Microsoft Authentication Library for how to use a python script

with MSAL. These instructions are for the Databricks service, but the scope is changed to the database App ID URI and scope instead of the Databricks scope.

1. Bypass the steps to set up the client app registration, since you have already accomplished that step except make sure you add a Redirect URI (`http://localhost`) for your client app registration.
2. Go directly to **Get Azure AD tokens by using the MSAL Python library**.

   You will need the Directory (tenant) ID, Client ID for the public app client, and the database App ID URI and scope. You will see a code section for **scopes** with directions to not modify this variable. Because this python code was written for Databricks scope, you will need to change this scope variable to the scope of your database. For example:

   ```
   scopes = ['https://example.com/1111aa1a-a1aa-1a11-11aa-1a1a11aa1111/
   session:connect']
   ```

3. Modify the code to write the token to a file location.

   Use the following example code and append it to the print statements at the end. Note the extra lines to back up and restore the original `stdout`.

   ```
   stdout_backup = sys.stdout
   with open('token', 'w') as token_file:
       sys.stdout = token_file
       print(acquire_tokens_result['access_token'])

   sys.stdout = stdout_backup
   ```

Example: Using Curl with a Resource Owner Password Credential Flow

This example shows how to use the `curl` command against the Azure AD API uses a Resource Owner Password Credential (ROPC) flow with a public Azure AD client.

The cleartext password is part of this command so this is not so much for end-users as it is for applications. This would need to be protected.

- Enter the following `curl` command:

  ```
  curl -X POST -H 'Content-Type: application/x-www-form-urlencoded' https://
  login.microsoftonline.com/az207oracleoutlook.onmicrosoft.com/oauth2/v2.0/
  token
  -d 'client_id=571c3f0a-aa3c-4f0a-93ed-4f75748955ea' -d 'scope=https://
  example.com/383fe7ee-1433-4844-a2d5-5b80d811256d/session:scope:connect'
  -d 'username=peter.fitch@example.com' -d 'password=password' -d
  'grant_type=password'
  ```

The response is a JSON file with token type, scope, expiration, and then the actual token. This file will need to be parsed so only the access token is written and stored in a file.

Example: Azure CLI Using Authorization Flow

This example shows how to use the Azure CLI to retrieve an access token and then write the token to a file.

See the Microsoft Azure article Install the Azure CLI on Linux for information about installing the Azure CLI.

1. Log in to your Azure tenancy.

   ```
   $ az login
   ```

2. Get an access token and assign it to the token variable using the following syntax:

   ```
   token=$(az account get-access-token --resource=database_app_id_uri --query
   accessToken --output tsv)
   ```

   For example:

   ```
   token=$(az account get-access-token --resource=https://example.com/
   1111aa1a-a1aa-1a11-11aa-1a1a11aa1111 --query accessToken --output tsv)
   ```

   If you get an error saying that the Azure CLI client app ID does not have permission to access the database resource, then copy the Azure CLI client app ID from the error message and add it to the list of authorized client applications for the database resource. (Go to the database app registration in Azure AD, click **Expose an API** and then **Add a client application**).

3. Write the token to a file.

   ```
   $ echo "$token" >> token
   ```

Configuring SQL*Plus for Azure AD Access Tokens

You must configure SQL*Plus to retrieve the Azure AD database access token from a location and use it when the `/` slash login is used.

Use Oracle Database release 19.16 and above SQL*Plus and Instance Client. Oracle Database release 21c clients do not support the full range of features. There is no default location for the Azure AD token, so you must specify this location.

1. Ensure that you have an Azure AD user account.
2. Check with an Azure AD administrator or Oracle Database administrator for one of the following:
   - An application client ID that you can use to get Azure AD tokens. If you have Azure AD privileges to do so, then create your own client app registration, similar to registering the Oracle Database instance with an Azure AD tenancy.
   - You are mapped to a global schema in the database.
3. Ensure that you are using the latest release updates for the Oracle Database client releases 19c.

   This configuration only works with the Oracle Database client release 19.16 and above. Oracle Database release 21c clients do not support the full range of features.
4. Configure the database server and client to use a TLS connection with the ExaDB-D database server.
5. On the client, set the following parameters in the `sqlnet.ora` file:
   - Check for the parameter `SSL_SERVER_DN_MATCH = ON` to ensure that DN matching is enabled.

- Set the `TOKEN_AUTH` parameter to enable the client to use the Azure AD token. Include the `TOKEN_LOCATION` parameter to point to the token location. For example:

  ```
  TOKEN_AUTH=OAUTH
  TOKEN_LOCATION="token_location"
  ```

  Note that there is no default location. If the token is named `token`, then you only need to specify the file directory (for example, `/test/oracle/aad-token`). If the token name is different from `token` (for example, `azure.token`), then you must include this name in the path (for example, `/test/oracle/aad-token/azure.token`).

You can specify the `TOKEN_AUTH` and `TOKEN_LOCATION` parameters in `tnsnames.ora`, as well as in `sqlnet.ora`. The `TOKEN_AUTH` and `TOKEN_LOCATION` values in the `tnsnames.ora` connect strings take precedence over the `sqlnet.ora` settings for that connection. For example:

```
(description=
  (retry_count=20)(retry_delay=3)
  (address=(protocol=tcps)(port=1522)
  (host=example.us-phoenix-1.oraclecloud.com))

(connect_data=(service_name=aaabbbccc_exampledb_high.example.oraclecloud.com))
  (security=(ssl_server_cert_dn="CN=example.uscom-east-1.oraclecloud.com,
     OU=Oracle BMCS US, O=Example Corporation,
     L=Redwood City, ST=California, C=US")
  (TOKEN_AUTH=OAUTH)(TOKEN_LOCATION="/test/oracle/aad-token"))
```

After the connect string is updated with the `TOKEN_AUTH` and `TOKEN_LOCATION` parameters, the Azure user can log in to the Oracle Database instance by running the following command to start SQL*Plus. You can include the connect descriptor itself or use the name of the descriptor from the `tnsnames.ora` file.

```
connect /@exampledb_high
```

Or the user can use the connect string. For example:

```
connect /@(description=
  (retry_count=20)(retry_delay=3)
  (address=(protocol=tcps)(port=1522)
  (host=example.us-phoenix-1.oraclecloud.com))

(connect_data=(service_name=aaabbbccc_exampledb_high.example.oraclecloud.com))
  (security=(ssl_server_cert_dn="CN=example.uscom-east-1.oraclecloud.com,
     OU=Oracle BMCS US, O=Example Corporation,
     L=Redwood City, ST=California, C=US") (TOKEN_AUTH=OAUTH)
(TOKEN_LOCATION="/test/oracle/aad-token")
```

The database client is already configured to get an Azure `OAuth2` token because `TOKEN_AUTH` has already been set, either through the `sqlnet.ora` file or in a connect string. The database client gets the `OAuth2` token and then sends the token to the Oracle Database instance.

Related Topics

- Registering the Oracle Database Instance with a Microsoft Azure AD Tenancy
  A user with administrator privileges uses Microsoft Azure AD to register the Oracle Database instance with the Microsoft Azure AD tenancy.

Trace Files for Troubleshooting Oracle Database Client Connections with Azure AD

You can use trace files to troubleshoot Oracle Database client connections with Azure AD connections.

- About Trace Files Used for Troubleshooting Connections
  You can generate two levels of trace files to troubleshoot Microsoft Azure AD connections on client side.
- Setting Client Tracing for Token Authentication
  You can add `EVENT` settings to the client-side `sqlnet.ora` file to control client tracing.

About Trace Files Used for Troubleshooting Connections

You can generate two levels of trace files to troubleshoot Microsoft Azure AD connections on client side.

The two levels of trace files that you can generate are as follows:

- Low level tracing prints traces in case of failures:
  - If TCPS is not set up for the Azure AD connection, then it prints a message that the protocol has to be TCPS.
  - If `SSL_SERVER_DN_MATCH` is not set to `TRUE`, then it prints a message that the value is `FALSE`.
  - If `TOKEN_LOCATION` has not been specified, then it prints a message that the token location does not exist.
  - If the token is not present at the specified `TOKEN_LOCATION`, then it prints a message.
  - If the application has passed in the token without setting `OCI_ATTR_TOKEN_ISBEARER` to true, it prints a message for the missing attribute.
  - If the application has set `OCI_ATTR_TOKEN_ISBEARER` to `TRUE` and not passed in the token, it prints a message for the missing attribute.
  - If the token has expired, then it prints a message.
- High level tracing prints traces in case of failure as mentioned above. In addition, it prints traces in case of success, as follows:
  - It prints where `SSL_SERVER_DN_MATCH` is present, `tnsnames.ora` or `sqlnet.ora`. It also prints the value as `TRUE` if set to `TRUE`.
  - If both the token and `OCI_ATTR_TOKEN_ISBEARER=true` are set by the application, then it prints a message.
  - If `TOKEN_AUTH` has the correct value `OAUTH`, then it prints the value.
  - If the token is not expired, then it prints a message.

Setting Client Tracing for Token Authentication

You can add `EVENT` settings to the client-side `sqlnet.ora` file to control client tracing.

These `EVENT` settings can be used for both IAM and Azure AD connections with Oracle Database.

- Use either of the following methods:
  - Add the following settings to the client side `sqlnet.ora` file:
    - `EVENT_25701=14` for low level tracing
    - `EVENT_25701=15` for high level tracing
  - Set the environment variable `EVENT_25701`:
    - `EVENT_25701=14` for low level tracing
    - `EVENT_25701=15` for high level tracing

  Client trace files are created in the following locations:
  - **Linux:** `$ORACLE_HOME/log/diag/clients`
  - **Windows:** `%ORACLE_HOME%\log\diag\clients`

  You can use the `ADR_BASE` parameter in the client side `sqlnet.ora` to specify the directory in which tracing messages are stored. Ensure that the directory path is valid and has write permissions. Ensure that the `DIAG_ADR_ENABLED` parameter is not set to `FALSE`.

  An example of setting `ADR_BASE` is as follows:

```
ADR_BASE=/oracle/oauth2/trace
```

6

Reference Guides for Exadata Cloud Infrastructure

- Using the dbaascli Utility on Exadata Cloud Infrastructure
  Learn to use the `dbaascli` utility on Exadata Cloud Infrastructure.
- Monitoring and Managing Exadata Storage Servers with ExaCLI
  The ExaCLI command line utility allows you to perform monitoring and management functions on Exadata storage servers in an Exadata Cloud Infrastructure instance.
- Monitor Metrics for VM Cluster Resources
- Metrics for Oracle Exadata Database Service on Dedicated Infrastructure in the Monitoring Service
  This article describes the metrics emitted by the Exadata Cloud Infrastructure Database service in the `oci_database_cluster` and `oci_database` namespaces for Oracle Databases.
- Oracle Exadata Database Service on Dedicated Infrastructure Events
  Exadata Cloud Infrastructure resources emit events, which are structured messages that indicate changes in resources.
- Policy Details for Exadata Cloud Infrastructure
  This topic covers details for writing policies to control access to Exadata Cloud Infrastructure resources.
- Managing Exadata Resources with Oracle Enterprise Manager Cloud Control
  To manage and monitor Exadata Cloud Infrastructure and Exadata Database Service on Cloud@Customer resources, use Oracle Enterprise Manager Cloud Control.
- Observability and Management for Exadata Database Service on Dedicated Infrastructure
- Security Guide for Oracle Exadata Database Service on Dedicated Infrastructure
- Troubleshooting Exadata Cloud Infrastructure Systems
  These topics cover some common issues you might run into and how to address them.

Using the dbaascli Utility on Exadata Cloud Infrastructure

Learn to use the `dbaascli` utility on Exadata Cloud Infrastructure.

- About Using the dbaascli Utility on Exadata Cloud Infrastructure
  You can use the dbaascli utility to perform various database lifecycle and administration operations on Exadata Cloud Infrastructure such as changing the password of a database user, starting a database, managing pluggable databases (PDBs), and more.
- Creating Databases Using dbaascli
  Using `dbaascli`, you can create an Oracle Database by first creating an Oracle Database home of desired version, followed by creating a database in that Oracle Database home
- Changing the Database Passwords
  To change the SYS password, or to change the TDE wallet password, use this procedure.

- Managing Exadata Cloud Infrastructure Software Images Using the Dbaascli Utility
  You can list and download the Oracle database software images on an Exadata Cloud Infrastructure instance, which can then be used for provisioning a database home.
- Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli
  Learn to use the `dbaascli` utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata Cloud Infrastructure system.
- Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check Using dbaascli
  Using the dbaascli `diag` command allows you to collect Guest VM `dbaas` tooling logs for Exadata Database Service on Dedicated Infrastructure and Exadata Database Service on Cloud@Customer systems. You can use these logs to troubleshoot issues related to `dbaas` tooling.
- Updating Cloud Tooling Using dbaascli
  To update the cloud tooling release for Oracle Exadata Database Service on Dedicated Infrastructure, complete this procedure.
- Creating a Duplicate Database
- Release Notes
  Review the changes made in various releases of `dbaascli`.
- dbaascli Command Reference
  You must use `dbaascli` to create databases and integrate them with the cloud automation framework.

About Using the dbaascli Utility on Exadata Cloud Infrastructure

You can use the dbaascli utility to perform various database lifecycle and administration operations on Exadata Cloud Infrastructure such as changing the password of a database user, starting a database, managing pluggable databases (PDBs), and more.

You must use the Oracle Cloud Infrastructure console or command-line interface to scale resources. The capabilities of the `dbaascli` utility are in addition to, and separate from, the Console, API, or command-line interface (CLI). Unless specified differently, you need `root` access to `dbaascli` to run all administration commands.

To use the utility, you must be connected to an Exadata Cloud Infrastructure virtual machine. For detailed instructions, see *Connecting to an Exadata Cloud Infrastructure Instance* .

To get possible commands available with `dbaascli`, run `dbaascli --help`.

To get command-specific help, run `dbaascli` *`command`* `--help`. For example, `dbaascli database create --help`.

See *dbasscli Command Reference* in the document for commands and command specific information.

Related Topics

- Connecting to an Exadata Cloud Infrastructure Instance
  This topic explains how to connect to an Exadata Cloud Infrastructure instance using SSH or SQL Developer.
- dbaascli Command Reference
  You must use `dbaascli` to create databases and integrate them with the cloud automation framework.

Creating Databases Using dbaascli

Using `dbaascli`, you can create an Oracle Database by first creating an Oracle Database home of desired version, followed by creating a database in that Oracle Database home

- Listing Available Software Images and Versions for Database and Grid Infrastructure
  To produce a list of available supported versions for patching, use the `dbaascli cswlib showImages` command.
- Creating Oracle Database Home
  To create an Oracle Database home of desired version, use the `dbaascli dbhome create` command.
- Creating Oracle Database In the Specified Oracle Database Home
  To create an Oracle Database in the specified Oracle Database home of desired version, use the `dbaascli database create` command.

Listing Available Software Images and Versions for Database and Grid Infrastructure

To produce a list of available supported versions for patching, use the `dbaascli cswlib showImages` command.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the following command:

   ```
   dbaascli cswlib showImages --product database
   ```

   The command output lists the available database software images.

   ```
   dbaascli cswlib showImages --product grid
   ```

   The command output lists the available grid software images.
4. Exit the `root` user command shell:

   ```
   exit
   ```

   For more details on advanced supported options, see `dbaascli cswlib showImages`.

Example 6-1 dbaascli cswlib showImages

```
[root@dg11lrg1 dbhome_1]# dbaascli cswlib showImages
DBAAS CLI version <version>
Executing command cswlib
      showImagesJob id: 00e89b1a-1607-422c-a920-22f44bec1953Log file location:
      /var/opt/oracle/log/cswLib/showImages/dbaastools_2022-05-11_08-49-12-
AM_46941.log

############
```

```
List of Available Database Images
#############

17.IMAGE_TAG=18.17.0.0.0
   VERSION=18.17.0.0.0
   DESCRIPTION=18c JAN 2022 DB Image

18.IMAGE_TAG=19.10.0.0.0
   VERSION=19.10.0.0.0
   DESCRIPTION=19c JAN 2021 DB Image

19.IMAGE_TAG=19.11.0.0.0
   VERSION=19.11.0.0.0
   DESCRIPTION=19c APR 2021 DB Image

20.IMAGE_TAG=19.12.0.0.0
  VERSION=19.12.0.0.0
  DESCRIPTION=19c JUL 2021 DB Image

21.IMAGE_TAG=19.13.0.0.0
  VERSION=19.13.0.0.0
  DESCRIPTION=19c OCT 2021 DB Image

Images can be downloaded using their image tags. For details, see help using
'dbaascli cswlib download --help'.
dbaascli execution completed
```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- dbaascli cswlib showImages
  To view the list of available Database and Grid Infrastructure images, use the `dbaascli cswlib showImages` command.

Creating Oracle Database Home

To create an Oracle Database home of desired version, use the `dbaascli dbhome create` command.

> **Note:**
>
> You can create an Oracle Database home with a specified Oracle home name. If you do not specify, then this is computed automatically (recommended).

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the following command:

```
dbaascli dbhome create --version Oracle Home Version --imageTag image Tag
Value
```

Where:

- `--version` specifies the Oracle Database version
- `--imageTag` specifies the Image Tag of the image to be used

For example:

```
dbaascli dbhome create --version 19.9.0.0.0
```

> **Note:**
>
> Specifying `imageTag` is optional. To view the Image Tags, refer to command `dbaascli cswlib showImages`. Image Tags are typically same as the version of the database. However, it is kept as a provision for cases where multiple images may need to be released for the same version - each catering to a specific customer requirement.

4. Exit the `root` user command shell:

```
exit
```

For more details on advanced supported options, see `dbaascli dbhome create`.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- dbaascli dbhome create
  To create an Oracle Database home of desired version, use the `dbaascli dbhome create` command.

Creating Oracle Database In the Specified Oracle Database Home

To create an Oracle Database in the specified Oracle Database home of desired version, use the `dbaascli database create` command.

You can use the `dbaascli database create` command to:

- Create a Container Database (CDB) or non-Container Database
- Create a CDB with pluggable databases (PDBs)
- Create an Oracle Database with the specified Character Set
- Create Oracle Databases on a subset of cluster nodes

> **Note:**
>
> Databases created on a subset of nodes will not be displayed in the OCI console.

- Create Oracle Database version 12.1.0.2 or higher with the release update JAN 2021 or higher. For databases with lower versions, it is recommended to use the OCI Console based API.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the following command:

   ```
   dbaascli database create --dbName database name --oracleHome Oracle Home
   Path
   ```

   Where:

   - `--dbName` specifies the name of the database
   - `--oracleHome` specifies Oracle home location

   To create a CDB, run the following command:

   ```
   dbaascli database create --dbName database name --oracleHome Oracle Home
   Path
   ```

   To create a non-CDB, run the following command:

   ```
   dbaascli database create --dbName database name --oracleHome Oracle Home
   Path --createAsCDB false
   ```

   When prompted, enter the `sys` and `tde` passwords.
4. Exit the `root` user command shell:

   ```
   exit
   ```

   For more details on advanced supported options, see `dbaascli database create`.

- Running Prerequisite Checks Prior to Creating Oracle Database
  To run prerequisites checks, use the `--executePrereqs` command option. This will perform only the prerequisite checks without performing the actual Oracle Database creation.
- Resuming or Reverting Oracle Database Creation Operation
  To resume or revert a failed database creation operation, use the `--resume` or `--revert` command option.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- dbaascli database create
  To create Oracle Database, use the `dbaascli database create` command. When prompted, enter the `sys` and `tde` passwords.

Running Prerequisite Checks Prior to Creating Oracle Database

To run prerequisites checks, use the `--executePrereqs` command option. This will perform only the prerequisite checks without performing the actual Oracle Database creation.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the following command:

   ```
   dbaascli database create --dbName database name --oracleHome Oracle Home
   Path --executePrereqs
   ```

   Where:

   - `--dbName` specifies the name of the database
   - `--oracleHome` specifies the Oracle home location

4. Exit the `root` user command shell:

   ```
   exit
   ```

   For more details on advanced supported options, see `dbaascli database create`.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- dbaascli database create
  To create Oracle Database, use the `dbaascli database create` command. When prompted, enter the `sys` and `tde` passwords.

Resuming or Reverting Oracle Database Creation Operation

To resume or revert a failed database creation operation, use the `--resume` or `--revert` command option.

For example:

```
dbaascli database create --dbName database name --oracleHome Oracle Home Path
--resume
```

> **Note:**
>
> - While using the `--resume` or `--revert` command options, ensure that you use the same command from the same node that was used for actual create operation flow.
> - You can resume database creation only if there is a failure in the post database creation step.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- dbaascli database create
  To create Oracle Database, use the `dbaascli database create` command. When prompted, enter the `sys` and `tde` passwords.

Changing the Database Passwords

To change the SYS password, or to change the TDE wallet password, use this procedure.

The password that you specify in the **Database Admin Password** field when you create a new Exadata Cloud Infrastructure instance or database is set as the password for the SYS, SYSTEM, TDE wallet, and PDB administrator credentials. Use the following procedures if you need to change passwords for an existing database.

> **Note:**
>
> if you are enabling Data Guard for a database, then the SYS password and the TDE wallet password of the primary and standby databases must all be the same.

> **Note:**
>
> Using the `dbaascli` to change the SYS password will ensure the backup/restore automation can parallelize channels across all nodes in the cluster.

To Change the SYS Password for an Exadata Cloud Infrastructure Database

1. Log onto the Exadata Cloud Infrastructure virtual machine as `opc`.
2. Run the following command:

```
sudo dbaascli database changepassword --dbname database_name --user SYS
```

To Change Database Passwords in a Data Guard Environment

1. Run the following command on the primary database:

   ```
   dbaascli database changePassword –dbName <dbname> --user SYS --
   prepareStandbyBlob true --blobLocation <location to create the blob file>
   ```

2. Copy the blob file created to all the standby databases and update the file ownership to `oracle` user.
3. Run the following command on all the standby databases:

   ```
   dbaascli database changePassword –dbName <dbname> --user SYS --
   standbyBlobFromPrimary <location of copies the blob file>
   ```

To Change the TDE Wallet Password for an Exadata Cloud Infrastructure Database

1. Log onto the Exadata Cloud Infrastructure virtual machine as `opc`.
2. Run the following command:

   ```
   sudo dbaascli tde changepassword --dbname database_name
   ```

Managing Exadata Cloud Infrastructure Software Images Using the Dbaascli Utility

You can list and download the Oracle database software images on an Exadata Cloud Infrastructure instance, which can then be used for provisioning a database home.

> **Note:**
>
> You can create custom database software images for your Exadata Cloud Infrastructure instances using the Console or API. These images are stored in Object Storage, and can be used to provision a Database Home in your Exadata instance. See Oracle Database Software Images more information.

You can control the version of Oracle binaries that is installed when you provision a new database on an Exadata Cloud Infrastructure instance by maintaining the software images on the system. Oracle provides a library of cloud software images that you can view and download onto your instance by using the `dbaascli` utility.

- Listing Available Software Images and Versions for Database and Grid Infrastructure
  To produce a list of available supported versions for patching, use the `dbaascli cswlib showImages` command.
- To download a software image
  You can download available software images onto your Exadata Cloud Infrastructure instance by using the `cswlib download` subcommand of the `dbaascli` utility.

Listing Available Software Images and Versions for Database and Grid Infrastructure

To produce a list of available supported versions for patching, use the `dbaascli cswlib showImages` command.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the following command:

   ```
   dbaascli cswlib showImages --product database
   ```

   The command output lists the available database software images.

   ```
   dbaascli cswlib showImages --product grid
   ```

   The command output lists the available grid software images.
4. Exit the `root` user command shell:

   ```
   exit
   ```

   For more details on advanced supported options, see `dbaascli cswlib showImages`.

Example 6-2 dbaascli cswlib showImages

```
[root@dg11lrg1 dbhome_1]# dbaascli cswlib showImages
DBAAS CLI version <version>
Executing command cswlib
      showImagesJob id: 00e89b1a-1607-422c-a920-22f44bec1953Log file location:
      /var/opt/oracle/log/cswLib/showImages/dbaastools_2022-05-11_08-49-12-
AM_46941.log

############
List of Available Database Images
#############

17.IMAGE_TAG=18.17.0.0.0
   VERSION=18.17.0.0.0
   DESCRIPTION=18c JAN 2022 DB Image

18.IMAGE_TAG=19.10.0.0.0
   VERSION=19.10.0.0.0
   DESCRIPTION=19c JAN 2021 DB Image

19.IMAGE_TAG=19.11.0.0.0
   VERSION=19.11.0.0.0
   DESCRIPTION=19c APR 2021 DB Image

20.IMAGE_TAG=19.12.0.0.0
```

```
  VERSION=19.12.0.0.0
  DESCRIPTION=19c JUL 2021 DB Image

21.IMAGE_TAG=19.13.0.0.0
  VERSION=19.13.0.0.0
  DESCRIPTION=19c OCT 2021 DB Image

Images can be downloaded using their image tags. For details, see help using
'dbaascli cswlib download --help'.
dbaascli execution completed
```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- dbaascli cswlib showImages
  To view the list of available Database and Grid Infrastructure images, use the `dbaascli cswlib showImages` command.

To download a software image

You can download available software images onto your Exadata Cloud Infrastructure instance by using the `cswlib download` subcommand of the `dbaascli` utility.

1. Connect to a compute node as the `opc` user.For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a root-user command shell:

   ```
   $ sudo -s
   #
   ```

3. Execute the `dbaascli` command with the `cswlib download` subcommand:

   ```
   # dbaascli cswlib download [--version <software_version>] [--imageTag
   <image tag
       value>]
   ```

   The command displays the location of software images that are downloaded to your Exadata Cloud Infrastructure environment.
   The optional parameters are:

   - **version:** specifies an Oracle Database software version. For example, 19.14.0.0.0.
   - **imageTag:** specifies the image tag of the image.

4. Exit the root-user command shell:

   ```
   # exit
   $
   ```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli

Learn to use the `dbaascli` utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata Cloud Infrastructure system.

- Patching Databases using dbaascli
  Using `dbaascli`, you can choose to patch a database by patching Oracle home, or by moving the database to an Oracle home with the desired patch level.
- Patching Oracle Grid Infrastructure
  To apply a patch to Oracle Grid Infrastructure, use the `grid patch` command.
- Listing Available Software Images and Versions for Database and Grid Infrastructure
  To produce a list of available supported versions for patching, use the `dbaascli cswlib showImages` command.
- Performing a Precheck Before Patching Databases and Grid Infrastructure
  You can perform a prerequisites-checking operation (also called a "precheck") for the commands in this topic using the applicable precheck flag.
- Resuming or Rolling Back a Patching Operation
  You can resume or revert a failed patching operation. Reverting a patch is known as a rollback.

Patching Databases using dbaascli

Using `dbaascli`, you can choose to patch a database by patching Oracle home, or by moving the database to an Oracle home with the desired patch level.

- Patching an Oracle home (in-place patching). This updates all databases located in the Oracle home.
- Moving a database to a different Oracle home that has the desired Oracle Database software version (out-of-place patching).
- Patching a Database Home (In-Place Database Patching)
  To patch an Oracle home, use the `dbaascli dbHome patch` command.
- Moving a Database to a Different Oracle Home (Out-of-Place Patching)
  To patch an Oracle Database by moving it to an Oracle home that is already at the desired patch level, use the `dbaascli database move` command.

Patching a Database Home (In-Place Database Patching)

To patch an Oracle home, use the `dbaascli dbHome patch` command.

This will patch all databases running in the specified home, and the databases will remain in the home after the patching is complete. The following apply to using the `dbHome patch` command for in-place patching operations:

- You can patch all of your database nodes or a subset of nodes.
- Multi-node patching takes place in a rolling fashion.
- Optionally, you can perform a software-only patch operation. Then, when you are ready, you can run `datapatch` to perform post-patch SQL actions.
- You can patch an Oracle home containing one or more databases.

To patch an Oracle Home (dbhome):

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the following command:

   ```
   dbaascli dbhome patch --oracleHome dbhome_path --targetVersion
   Oracle_Database_version
   ```

   Where:

   - `--oracleHome` identifies the path of the Oracle home to be patched.
   - `--targetVersion` specifies the target Oracle Database version to use for patching, specified as five numeric segments separated by periods (e.g. 19.12.0.0.0).

   For example:

   ```
   dbaascli dbhome patch --oracleHome /u02/app/oracle/product/19.0.0.0/
   dbhome_2 --targetVersion 19.9.0.0.0
   ```

4. Exit the `root` user command shell:

   ```
   exit
   ```

   For more details on advanced supported options, see `dbaascli dbHome patch`.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- dbaascli dbHome patch
  To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.

Moving a Database to a Different Oracle Home (Out-of-Place Patching)

To patch an Oracle Database by moving it to an Oracle home that is already at the desired patch level, use the `dbaascli database move` command.

After the database move operation is complete, the database runs using the Oracle Database software version of the target Oracle Home.

To patch a database by moving it to a different Oracle Home:

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the following command:

   ```
   dbaascli database move --oracleHome path_to_target_oracle_home --dbname
   database_name
   ```

   Where:

   - `--oracleHome` identifies the path of the target Oracle home that uses the desired Oracle Database software version. Note that the target Oracle home must exist in your system prior to using the `database move` command.
   - `--dbname` specifies the name of the database that is being moved.

   For example:

   ```
   dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0.0/
   dbhome_2 --dbname xyz
   ```

4. Exit the `root` user command shell:

   ```
   exit
   ```

   For more details on advanced supported options, see `dbaascli database move`.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- dbaascli database move
  To move the database from one home to another, use the `dbaascli database move` command.

Patching Oracle Grid Infrastructure

To apply a patch to Oracle Grid Infrastructure, use the `grid patch` command.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the following command:

   ```
   dbaascli grid patch --targetVersion target_software_version_number
   ```

   Where `--targetVersion` identifies target software version that the Oracle Grid Infrastructure will be patched to.

   For example:

   ```
   dbaascli grid patch --targetVersion 19.11.0.0.0
   ```

4. Exit the `root` user command shell:

```
exit
```

For more details on advanced supported options, see `dbaascli grid patch`.

- Patching Oracle Grid Infrastructure (GI) Using GI Software Image
  To patch Oracle Grid Infrastructure (GI) using GI software image, use this procedure.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- dbaascli grid patch
  To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.

Patching Oracle Grid Infrastructure (GI) Using GI Software Image

To patch Oracle Grid Infrastructure (GI) using GI software image, use this procedure.

Oracle Grid Infrastructure can also be patched by first creating a patched software image, and then using that image to perform the patching operation. This provides the advantage that an image can be created ahead of time outside of the patching window. It also helps in conflict resolution as any conflicts among the patches are highlighted during the image creation process without impacting the patching window.

1. Create a patched software image.

```
dbaascli grid patch --targetVersion <target_software_version_number> --
createImage
```

Once the patched software image creation is completed, the image can then be used for performing the patching operation.

2. Perform the patching operation.

```
dbaascli grid patch --targetVersion <target_software_version_number> --
imageLocation <location_of_patched_software_image>
```

Listing Available Software Images and Versions for Database and Grid Infrastructure

To produce a list of available supported versions for patching, use the `dbaascli cswlib showImages` command.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

```
sudo -s
```

3. Run the following command:

```
dbaascli cswlib showImages --product database
```

The command output lists the available database software images.

```
dbaascli cswlib showImages --product grid
```

The command output lists the available grid software images.

4. Exit the `root` user command shell:

```
exit
```

For more details on advanced supported options, see `dbaascli cswlib showImages`.

Example 6-3 dbaascli cswlib showImages

```
[root@dg11lrg1 dbhome_1]# dbaascli cswlib showImages
DBAAS CLI version <version>
Executing command cswlib
      showImagesJob id: 00e89b1a-1607-422c-a920-22f44bec1953Log file location:
      /var/opt/oracle/log/cswLib/showImages/dbaastools_2022-05-11_08-49-12-
AM_46941.log

############
List of Available Database Images
#############

17.IMAGE_TAG=18.17.0.0.0
   VERSION=18.17.0.0.0
   DESCRIPTION=18c JAN 2022 DB Image

18.IMAGE_TAG=19.10.0.0.0
   VERSION=19.10.0.0.0
   DESCRIPTION=19c JAN 2021 DB Image

19.IMAGE_TAG=19.11.0.0.0
   VERSION=19.11.0.0.0
   DESCRIPTION=19c APR 2021 DB Image

20.IMAGE_TAG=19.12.0.0.0
  VERSION=19.12.0.0.0
  DESCRIPTION=19c JUL 2021 DB Image

21.IMAGE_TAG=19.13.0.0.0
  VERSION=19.13.0.0.0
  DESCRIPTION=19c OCT 2021 DB Image

Images can be downloaded using their image tags. For details, see help using
'dbaascli cswlib download --help'.
dbaascli execution completed
```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- dbaascli cswlib showImages
  To view the list of available Database and Grid Infrastructure images, use the `dbaascli cswlib showImages` command.

Performing a Precheck Before Patching Databases and Grid Infrastructure

You can perform a prerequisites-checking operation (also called a "precheck") for the commands in this topic using the applicable precheck flag.

Running prechecks allows you to run only the precheck portion of the patching operation without performing actual patching. Oracle recommends running prechecks to discover software issues that could prevent successful patching.

To perform patching prechecks, first, connect to a virtual machine in your Exadata Cloud Infrastructure instance as the `root` user.

- Precheck for Oracle Home Patching (In-Place Patching)
  Use the `--executePrereqs` flag with the `dbaascli dbhome patch` command.
- Precheck for Database Move Patching (Out-of-Place Patching)
  Use the `--executePrereqs` flag with the `dbaascli database move` command.
- Precheck for Oracle Grid Infrastructure Patching
  Use the `--executePrereqs` flag with the `dbaascli grid patch` command.

Precheck for Oracle Home Patching (In-Place Patching)

Use the `--executePrereqs` flag with the `dbaascli dbhome patch` command.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the following command:

   ```
   dbaascli dbhome patch --oracleHome dbhome_path --targetVersion
   Oracle_Database_version --executePrereqs
   ```

   Where:

   - `--oracleHome` identifies the path of the Oracle home to be prechecked.
   - `--targetVersion` specifies the target Oracle Database version to be patched to, specified as five numeric segments separated by periods (e.g. 19.12.0.0.0).
4. Exit the `root` user command shell:

   ```
   exit
   ```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- dbaascli dbHome patch
  To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.

Precheck for Database Move Patching (Out-of-Place Patching)

Use the `--executePrereqs` flag with the `dbaascli database move` command.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the following command:

   ```
   dbaascli database move --oracleHome path_to_target_oracle_home --dbname
   database_name --executePrereqs
   ```

   Where:

   - `--oracleHome` identifies the path of the target Oracle Home that uses the desired Oracle Database software version. Note that the target Oracle Home must exist in your system prior to using the `database move` command.
   - `--dbname` specifies the name of the database that is being moved

4. Exit the `root` user command shell:

   ```
   exit
   ```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- dbaascli database move
  To move the database from one home to another, use the `dbaascli database move` command.

Precheck for Oracle Grid Infrastructure Patching

Use the `--executePrereqs` flag with the `dbaascli grid patch` command.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the following command:

```
dbaascli grid patch --targetVersion target_software_version_number --
executePrereqs
```

Where `--targetVersion` identifies target software version that the Oracle Grid Infrastructure will be patched to, specified as five numeric segments separated by periods, for example, 19.12.0.0.0

4. Exit the `root` user command shell:

```
exit
```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- dbaascli grid patch
  To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.

Resuming or Rolling Back a Patching Operation

You can resume or revert a failed patching operation. Reverting a patch is known as a rollback.

- Resuming a Patch Operation
  To resume a patching operation, use the `--resume` flag with the original patching command.
- Rolling Back a Patch Operation
  Use the `--rollback` flag with the original patching command to roll back (revert) a patching operation.

Resuming a Patch Operation

To resume a patching operation, use the `--resume` flag with the original patching command.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

```
sudo -s
```

3. Run the original patching command to resume a patching operation:
   For example:

```
dbaascli dbhome patch --oracleHome /u02/app/oracle/product/19.0.0.0/
dbhome_2 --targetVersion 19.9.0.0.0 --resume
```

4. Exit the `root` user command shell:

```
exit
```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- dbaascli dbHome patch
  To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.
- dbaascli grid patch
  To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.

Rolling Back a Patch Operation

Use the `--rollback` flag with the original patching command to roll back (revert) a patching operation.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the original patching command to roll back (revert) a patching operation:
   For example:

   ```
   dbaascli grid patch --targetVersion 19.11.0.0.0 --rollback
   ```

   > **Note:**
   >
   > - Resume and Rollback operations are supported for Oracle Home patching, Oracle Grid Infrastructure patching, and database move operations.
   > - When resuming or rolling back a patching operation, you must run the resume or rollback command from the same node that was used to run the original patching command, and you must run the original command with the addition of the `--resume` or `--rollback` flag.

4. Exit the `root` user command shell:

   ```
   exit
   ```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- dbaascli dbHome patch
  To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.

- dbaascli grid patch
  To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.

Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check Using dbaascli

Using the dbaascli `diag` command allows you to collect Guest VM `dbaas` tooling logs for Exadata Database Service on Dedicated Infrastructure and Exadata Database Service on Cloud@Customer systems. You can use these logs to troubleshoot issues related to `dbaas` tooling.

You can use the `diag` command to collect dbaastools logs and perform a health check on all nodes in an Exadata cluster. Note that the `--waitForCompletion` options is supported starting in version 22.4.1

> **Note:**
>
> - dbaascli `diag` commands must be run as the `root` user
> - Running the `dbaascli diag collect` command on a single node will collect log data for all nodes
> - We recommend running the commands documented in this topic using the `--waitForCompletion` option for long-running commands. Refer to the examples for sample usage.

For information on updating Exadata Cloud Tooling, see *dbaascli admin updateStack*.

- Collecting Tooling Log Data Examples
  The dbaascli dbaascli diag collect command uses the syntax shown below to collect tooling log data:
- Performing a Health Check Examples
  Use dbaascli `dbaascli diag healthcheck` command to perform a health check on all system nodes.

Related Topics

- dbaascli diag collect
  To collect diagnostics, use the `dbaascli diag collect` command.
- dbaascli admin updateStack
  To install or update a dbaastools RPM, use the `dbaascli admin updateStack` command.

Collecting Tooling Log Data Examples

The dbaascli dbaascli diag collect command uses the syntax shown below to collect tooling log data:

See `dbaascli diag collect` In the *dbaascli Command Reference* for syntax details

NOT_SUPPORTED

```
# dbaascli diag collect
DBAAS CLI version 24.1.1.0.0
Executing command diag collect
Job id: 92f33125-aa70-4ce2-94fb-64d8f1cbdc93
Session log: /var/opt/oracle/log/diag/collect/dbaastools_2023-12-14_07-20-44-
PM_83383.log
Loading PILOT...
Session ID of the current execution is: 10
Log file location: /var/opt/oracle/log/diag/collect/pilot_2023-12-14_07-20-48-
PM_83856
-----------------
..
---------- DIAG COLLECT PLUGIN RESULT ----------
{
  "collectedArchive with SHA256 CheckSum" : "{/var/opt/oracle/dbaas_acfs/
diag_collect/artifacts_diag_cloudlogs_20231214-1920/
diag_cloudlogs_20231214-1920_node1.zip=a0d049b87ab9e9cec2ab7d95ded4903bac818c8
1c8b6a46d295e1e75f4630e19}"
}
dbaascli execution completed
```

NOT_SUPPORTED

```
# dbaascli diag collect --waitForCompletion false
DBAAS CLI version 24.1.1.0.0
Executing command diag collect --waitForCompletion false
Job id: 5b556976-dba1-4be9-a4fe-4b58e69c1d96
Session log: /var/opt/oracle/log/diag/collect/dbaastools_2023-12-14_07-23-26-
PM_98107.log
Job accepted. Use "dbaascli job getStatus --jobID 5b556976-dba1-4be9-
a4fe-4b58e69c1d96" to check the job status.
```

> **Note:**
>
> Use the job status command to monitor progress.

NOT_SUPPORTED

```
# dbaascli diag collect --dbnames myOracleDatabase19cName
DBAAS CLI version 24.1.1.0.0
Executing command diag collect --dbnames myOracleDatabase19cName
Job id: 8e1d2667-4649-4384-8610-b6348d6548ac
Session log: /var/opt/oracle/log/diag/collect/dbaastools_2023-12-14_08-41-41-
PM_88831.log
Loading PILOT...
Session ID of the current execution is: 12
Log file location: /var/opt/oracle/log/diag/collect/pilot_2023-12-14_08-41-45-
PM_89361
-----------------
..
```

```
---------- DIAG COLLECT PLUGIN RESULT ----------
{
  "collectedArchive with SHA256 CheckSum" : "{/var/opt/oracle/dbaas_acfs/
diag_collect/artifacts_diag_cloudlogs_20231214-2041/
diag_cloudlogs_20231214-2041_node1.zip=9e50500089a74ca7cd8ae08550c06868e26e1cd
9c52e808194256594f63397e4}"
}
dbaascli execution completed
```

NOT_SUPPORTED

```
# dbaascli diag collect --destLocation /tmp/test/
DBAAS CLI version 24.1.1.0.0
Executing command diag collect --destLocation /tmp/test/
Job id: f992afdf-415e-4b58-ab5b-9e38f8c2079d
Session log: /var/opt/oracle/log/diag/collect/dbaastools_2023-12-14_09-42-54-
PM_16270.log
Loading PILOT...
Session ID of the current execution is: 14
Log file location: /var/opt/oracle/log/diag/collect/pilot_2023-12-14_09-42-58-
PM_16777
-----------------
..
---------- DIAG COLLECT PLUGIN RESULT ----------
{
  "collectedArchive with SHA256 CheckSum" : "{/tmp/test/diag_collect/
artifacts_diag_cloudlogs_20231214-2143/
diag_cloudlogs_20231214-2143_node1.zip=8a26cffcfdd72c261660d4f736c615981856e35
7749d90751b94f3eda19a9a70}"
}
dbaascli execution completed
```

NOT_SUPPORTED

```
# dbaascli diag collect --startTime 2023-12-05T10:00:00 --endTime
2023-12-05T11:00:00
DBAAS CLI version 24.1.1.0.0
Executing command diag collect --startTime 2023-12-05T10:00:00 --endTime
2023-12-05T11:00:00
Job id: 70b03e50-98cc-4c2b-9684-1f82070bac88
Session log: /var/opt/oracle/log/diag/collect/dbaastools_2023-12-14_09-45-17-
PM_42856.log
Loading PILOT...
Session ID of the current execution is: 15
Log file location: /var/opt/oracle/log/diag/collect/pilot_2023-12-14_09-45-21-
PM_43526
-----------------
..
---------- DIAG COLLECT PLUGIN RESULT ----------
{
  "collectedArchive with SHA256 CheckSum" : "{/var/opt/oracle/dbaas_acfs/
diag_collect/artifacts_diag_cloudlogs_20231214-2145/
diag_cloudlogs_20231214-2145_node1.zip=b44cf3bfca1ab7a1629dd83098a7772790ab949
e50dbb3950f0017e427d7bd05}"
```

```
}
dbaascli execution completed
```

NOT_SUPPORTED

```
# dbaascli diag collect --nodes node1,node2
DBAAS CLI version 24.1.1.0.0
Executing command diag collect --nodes node1,node2
Job id: fa70da09-3de6-4cc8-854c-a739b4fc2ceb
Session log: /var/opt/oracle/log/diag/collect/dbaastools_2023-12-14_09-46-58-
PM_55884.log
Loading PILOT...
Session ID of the current execution is: 16
Log file location: /var/opt/oracle/log/diag/collect/pilot_2023-12-14_09-47-02-
PM_56418
-----------------
..
---------- DIAG COLLECT PLUGIN RESULT ----------
{
  "collectedArchive with SHA256 CheckSum" : "{/var/opt/oracle/dbaas_acfs/
diag_collect/artifacts_diag_cloudlogs_20231214-2147/
diag_cloudlogs_20231214-2147_node1.zip=de2805c9c6c2af2d602395a84d37747935327b7
3a6c73052282665a8410eb41f}"
}
```

NOT_SUPPORTED

```
# dbaascli diag collect --components dbaastools
DBAAS CLI version 24.1.1.0.0
Executing command diag collect --components dbaastools
Job id: da941d3c-5191-4ced-b1bb-9b083fa75865
Session log: /var/opt/oracle/log/diag/collect/dbaastools_2023-12-14_09-47-23-
PM_68256.log
Loading PILOT...
Session ID of the current execution is: 17
Log file location: /var/opt/oracle/log/diag/collect/pilot_2023-12-14_09-47-27-
PM_68729
-----------------
..
---------- DIAG COLLECT PLUGIN RESULT ----------
{
  "collectedArchive with SHA256 CheckSum" : "{/var/opt/oracle/dbaas_acfs/
diag_collect/artifacts_diag_cloudlogs_20231214-2147/
diag_cloudlogs_20231214-2147_node1.zip=d1f290fb42c981935e1142ec059c2dbba8be2e0
a9ffebc9eea83a6336abe2eed}"
}
dbaascli execution completed
```

NOT_SUPPORTED

```
# dbaascli diag collect --objectStoreBucketUri https://objectstorage.us-
phoenix-1.oraclecloud.com/p/aL-IbIKQ1j6lWNftJc2rLoLh6o9bJgbZm8z0S--
BeVuXaipSEEMISrSCfFrVEolG/n/intexadatateam/b/diag_collect_test/o/
DBAAS CLI version 24.1.1.0.0
```

```
Executing command diag collect --objectStoreBucketUri https://
objectstorage.us-phoenix-1.oraclecloud.com/p/aL-
IbIKQ1j6lWNftJc2rLoLh6o9bJgbZm8z0S--BeVuXaipSEEMISrSCfFrVEolG/n/
intexadatateam/b/diag_collect_test/o/
Job id: 028151b7-cbc4-409a-9ec6-69affe10f3bb
Session log: /var/opt/oracle/log/diag/collect/dbaastools_2023-12-14_09-51-36-
PM_2963.log
Loading PILOT...
Session ID of the current execution is: 20
Log file location: /var/opt/oracle/log/diag/collect/pilot_2023-12-14_09-51-40-
PM_3555
-----------------
..
---------- DIAG COLLECT PLUGIN RESULT ----------
{
  "collectedArchive with SHA256 CheckSum" : "{/var/opt/oracle/dbaas_acfs/
diag_collect/artifacts_diag_cloudlogs_20231214-2151/
diag_cloudlogs_20231214-2151_node1.zip=71633e13ccd06de15cb26850bb0266cf0d869e2
59550515c5b1fb734c487b470}"
}
dbaascli execution completed
```

Related Topics

- dbaascli diag collect
  To collect diagnostics, use the `dbaascli diag collect` command.

Performing a Health Check Examples

Use dbaascli `dbaascli diag healthcheck` command to perform a health check on all system nodes.

See *dbaascli diag healthcheck* for the syntax details in the dbaascli Command Reference.

NOT_SUPPORTED

```
# dbaascli diag healthcheck
DBAAS CLI version MAIN
Executing command diag healthcheck
INFO: Starting diag healthcheck
INFO: Collected diag logs at: /var/opt/oracle/dbaas_acfs/
diag_cloudlogs_20210322-2246.tar.gz
```

NOT_SUPPORTED

```
# dbaascli diag healthcheck --destLocation /tmp/test
DBAAS CLI version MAIN
Executing command diag healthcheck --destLocation /tmp/test
INFO: Starting diag healthcheck
INFO: Collected diag logs at: /tmp/test/diag_cloudlogs_20210322-2250.tar.gz
```

NOT_SUPPORTED

```
# dbaascli diag healthcheck --nodes rbcl1,rbcl2
DBAAS CLI version MAIN
```

```
Executing command diag healthcheck --nodes rbcl1,rbcl2
INFO: Starting diag healthcheck
INFO: Collected diag logs at: /var/opt/oracle/dbaas_acfs/
diag_cloudlogs_20210421-1915.tar.gz
```

NOT_SUPPORTED

```
# dbaascli diag healthcheck --objectStoreBucketUri https://objectstorage.us-
phoenix-1.oraclecloud.com/p/t0Z-kRV5pSmFzqnf-
y5XhaAbM4LS82epeBnulKnCr31IeHVjxI9tOkntLF2kq7fP/n/MyNamespace/b/MyParBucket/o/
DBAAS CLI version MAIN
Executing command diag healthcheck --objectStoreBucketUri https://
objectstorage.us-phoenix-1.oraclecloud.com/p/t0Z-kRV5pSmFzqnf-
y5XhaAbM4LS82epeBnulKnCr31IeHVjxI9tOkntLF2kq7fP/n/MyNamespace/b/MyParBucket/o/
INFO: Collected diag logs at: https://objectstorage.us-
phoenix-1.oraclecloud.com/p/t0Z-kRV5pSmFzqnf-
y5XhaAbM4LS82epeBnulKnCr31IeHVjxI9tOkntLF2kq7fP/n/MyNamespace/b/MyParBucket/o/
diag_cloudlogs_20210421-1839.tar.gz
```

Related Topics

- dbaascli diag collect
  To collect diagnostics, use the `dbaascli diag collect` command.
- dbaascli diag healthCheck
  To run diagnostic health checks, use the `dbaascli diag healthCheck` command.

Updating Cloud Tooling Using dbaascli

To update the cloud tooling release for Oracle Exadata Database Service on Dedicated Infrastructure, complete this procedure.

Cloud-specific tooling is used on the Exadata Cloud Infrastructure Guest VMs for local operations, including `dbaascli` commands.

The cloud tooling is automatically updated by Oracle when new releases are made available. If needed, you can follow the steps below to ensure you have the latest version of the cloud-specific tooling on all of the virtual machines in the VM cluster.

> **Note:**
>
> You can update the cloud-specific tooling by downloading and applying a software package containing the updated tools.

1. Connect to a virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. To update to the latest available cloud tooling release, run the following command:

```
dbaascli admin updateStack
```

The command takes care of updating the cloud tooling release on all the nodes of the cluster.

For more details and other available options, refer to `dbaascli admin updateStack --help`.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- dbaascli admin updateStack
  To install or update a dbaastools RPM, use the `dbaascli admin updateStack` command.

Creating a Duplicate Database

- Using dbaascli to Duplicate a Cloud Database
- Considerations When Using OCI Vault for the Key Management
- Duplicate an On-Premises Database

Using dbaascli to Duplicate a Cloud Database

You can create a duplicate database using `dbaascli`. This new database can be in the same cloud region as the source region or across the regions. The following steps describe how to create a duplicate database on cloud.

> **Note:**
>
> If a database is configured with OCI Vault for TDE encryption and you want to duplicate a database, then refer to the following sections.

Prepare for duplication

Ensure that the following prerequisites are ment:

- Make sure that there is a network path setup to access the source database through the `EZConnect` string.
- Copy the TDE wallet file (`ewallet.p12`) to the target database node. The node where you decide to run the `dbaascli` command.
- Create an Oracle home on the target node if required. Oracle home version must be the same version as the source or of higher RU version.

Run prerequisite checks

To run prerequisites checks, use the `--executePrereqs` command option. This will perform only the prerequisite checks without performing the actual Oracle Database duplication.

```
dbaascli database duplicate --dbName <database name> --oracleHome <Oracle
Home Path> --sourceDBConnectionString <source database EZConnect string> --
sourceDBTDEWalletLocation <location of copied wallet> --
sourceDBTdeConfigMethod FILE --tdeConfigMethod FILE --executePrereqs
```

Duplicate the database

```
dbaascli database duplicate --dbName <database name> --oracleHome <Oracle
Home Path> --sourceDBConnectionString <source database EZConnect string> --
sourceDBTDEWalletLocation <location of copied wallet> --
sourceDBTdeConfigMethod FILE --tdeConfigMethod FILE
```

Considerations When Using OCI Vault for the Key Management

This section is applicable only in the case of database is configured with OCI Vault for TDE encryption and you want to duplicate a database.

Duplicating a database within the same region

- Additional prerequisite steps
  Make sure to setup OCI Vault access policies for target database nodes. Target database nodes should be able to access both source database's OCI key vault along with its new key vault (if it is decided to use separate key vault).
- Run prerequisite checks

  ```
  dbaascli database duplicate --dbName <database name> --oracleHome <Oracle
  Home Path> --sourceDBConnectionString <source database EZConnect string> --
  sourceDBTDEWalletLocation <location of copied wallet> --
  sourceDBTdeConfigMethod KMS --sourceDBKMSKeyOCID <Source Database OCI
  Vault key OCID> --tdeConfigMethod KMS --kmsKeyOCID <OCI Vault key OCID> --
  executePrereqs
  ```

- Duplicate the database

  ```
  dbaascli database duplicate --dbName <database name> --oracleHome <Oracle
  Home Path> --sourceDBConnectionString <source database EZConnect string> --
  sourceDBTDEWalletLocation <location of copied wallet> --
  sourceDBTdeConfigMethod KMS --sourceDBKMSKeyOCID <Source Database OCI
  Vault key OCID> --tdeConfigMethod KMS --kmsKeyOCID <OCI Vault key OCID>
  ```

  Upon successful completion of this command, the database is duplicated.

Duplicating a database across regions

- Additional prerequisite steps

  - Setup a new OCI Vault for target database on the corresponding region by following the steps outlined in Prepare to Use Customer-Managed Keys in the Vault Service. Complete Tasks 1 through 3.
  - Setup OCI Vault replication from source region to target region. For more information, see Replicating Vaults and Keys.
  - Update Dynamic group policy, which is created in step 2 to allow access to replicated OCI Vault key.

- Run prerequisite checks

```
dbaascli database duplicate --dbName <database name> --oracleHome <Oracle
Home Path> --sourceDBConnectionString <source database EZConnect string> --
sourceDBTDEWalletLocation <location of copied wallet> --
sourceDBTdeConfigMethod KMS --sourceDBKMSKeyOCID <Source Database OCI
Vault key OCID> --tdeConfigMethod KMS --kmsKeyOCID <OCI Vault key OCID> --
executePrereqs
```

- Duplicate the database

```
dbaascli database duplicate --dbName <database name> --oracleHome <Oracle
Home Path> --sourceDBConnectionString <source database EZConnect string> --
sourceDBTDEWalletLocation <location of copied wallet> --
sourceDBTdeConfigMethod KMS --sourceDBKMSKeyOCID <Source Database OCI
Vault key OCID> --tdeConfigMethod KMS --kmsKeyOCID <OCI Vault key OCID>
```

  Upon successful completion of this command, the database is duplicated.

Duplicate an On-Premises Database

Using `dbaascli`, you can duplicate an on-prem database onto the cloud. This can be done with the `dbaascli database duplicate` command. This command creates a new database on the cloud, which is a duplicate of an on-prem database along with its data. While this process is going on, the on-prem database remains still operational. You can migrate your applications to the duplicated database on the cloud after due verification.

Prepare for duplication

The migration process includes the following prerequisites to be met.

- Make sure that there is a network path setup to access an on-prem database from the OCI node through the `EZConnect` string.
- If an on-prem database is configured with TDE, copy the TDE wallet file (`ewallet.p12`) to the OCI node, where you decide to run the `dbaascli` command.
- Create an Oracle home on the OCI node if required. The Oracle home version must be the same as the source or of a higher RU version.

Verify the necessary RPMs

This process requires a minimum `dbaastools` RPM version of 23.3.2.0.0 but updating to the latest `dbaastools` rpm is always recommended.

- To check the currently installed version, run:

```
dbaascli --version
DBAAS CLI version 23.3.2.0.0
```

- To apply the latest tools RPM, as the `root` user, run:

```
# dbaascli admin updateStack
```

Run the prerequisite checks

To run the prerequisite checks, use the `--executePrereqs` command option. This will perform only the prerequisite checks without performing the actual Oracle Database duplication.

```
dbaascli database duplicate --dbName <database name> --oracleHome <Oracle Home
Path> --sourceDBConnectionString <source database EZConnect string> --
sourceDBTDEWalletLocation <location of copied wallet> --executePrereqs
```

Duplicate the database

Duplicate the database using the following command:

```
dbaascli database duplicate --dbName <database name> --oracleHome <Oracle Home
Path> --sourceDBConnectionString <source database EZConnect string> --
sourceDBTDEWalletLocation <location of copied wallet>
```

For example:

```
dbaascli database duplicate --sourceDBConnectionString xyzhost.oracle.com:1521/
dbuniquename.oracle.com --dbName orcl --oracleHome /u02/app/oracle/product/
19.0.0.0/dbhome_1 --sourceDBTDEWalletLocation /tmp/wallet_copy/tde --
waitForCompletion false
```

Upon successful completion of this command, the database is duplicated to Cloud and ready for sanity checks for application usage. Once verification is done, application connections can be migrated to the Cloud database.

Refer to `dbaascli database duplicate -help` for additional configuration options.

Few considerations for migration

- If you prefer to allocate multiple channels for RMAN duplicate, you could do so by specifying the `--rmanParallelism` argument.
- Exadata Cloud Service configures database memory as Automatic Shared Memory Management (ASMM). If your on-prem database is configured with different memory management, make sure to adjust memory parameter values accordingly on the OCI side by providing values for `--sgaSizeInMB` and `--pgaSizeInMB`.
- Verify that the on-prem database does not contain any deprecated or invalid initialization parameters.
- Database initialization parameters related to database storage (datafile location, redo location, recovery area destination, control file multiplexing) may be changed using the `--initParams` argument.
  For example, to override `db_create_online_log_dest` value for the duplicate database: `--initParams db_create_online_log_dest_1=+DATAC1,db_create_online_log_dest_2=+RECOC1`

Troubleshooting the database duplication

- `dbaascli` operation log file can be found under `/var/opt/oracle/log/<dbname>/database/duplicate`
- One of the jobs of the duplicate is to run `dbca`. Its log file can be found under `/u02/app/oracle/cfgtoollogs/dbca` and `/u02/app/oracle/cfgtoollogs/dbca/<dbuniquename>`.

If the operation fails, you will have an option to resume the operation by providing the `--resume` argument to the same command. Alternatively, clean up the database using `dbaascli database delete -dbname <dbname> -force`, and then rerun the database duplicate command.

Release Notes

Review the changes made in various releases of `dbaascli`.

- Release 24.2.1.0.0 (240530)
- Release 24.1.2.0.0 (240306)
- Release 24.1.1.0.0 (231219)
- Release 23.4.1.0.0 (231102)
- Release 23.3.2.0.0 (230503)
- Release 23.3.1.0.0 (230712)
- Release 23.2.1.0.0 (230503)
- Release 23.1.2.0.0 (230305)
- Release 23.1.1.0.1 (230113)
- Release 22.4.1.0.1 (221122)
- Release 22.3.1.1.0 (221003)
- Release 22.3.1.0.1 (220721)
- Release 22.2.1.1.0 (220623)
- Release 22.2.1.1.0 (220609)
- Release 22.2.1.0.1 (220423)
- Release 22.1.1.2.0 (220405)
- Release 22.1.1.1.0 (220317)
- Release 22.1.1.0.1 (220223)
- Release 21.4.1.1.0 (220209)
- Release 21.4.1.1.0
- Release 21.3.1.2.0
- Release 21.3.1.1.0
- Release 21.3.1.0.1
- Release 21.2.1.x.x

Release 24.2.1.0.0 (240530)

- Added support for Oracle Database 23ai.
- Improvements in backup and recovery area with Zero Data Loss Autonomous Recovery Service (ZRCV) as backup destination.
- Various bug fixes and stability improvements.

Release 24.1.2.0.0 (240306)

- Introduced a new optimized workflow for Data Guard operations
- Various bug fixes and stability improvements

Release 24.1.1.0.0 (231219)

- Various bug fixes and stability improvements

Release 23.4.1.0.0 (231102)

- **Backup and Recovery:** Minimum Backup recovery window has been changed to 7 days. While obsoleting backup pieces automation considers recovery window as 7 days if it discovers any value less than 7 from the system.
- Various bug fixes and stability improvements

Release 23.3.2.0.0 (230503)

- Pluggable Database Operations
  - Added support to set custom key version OCID (Bring Your Own Key - BYOK) of OCI Vault during create and clone operations. For details, see respective PDB commands help.
- Grid Infrastructure (GI) Patching
  - Enhanced the patching workflow to improve patching time, especially in environments having high number of databases.
- Database Patching
  - Provided option to run datapatch on a specific node of cluster.
- Various bug fixes and stability improvements

Release 23.3.1.0.0 (230712)

- New `dbaascli` commands
  - `dbaascli gridHome create` - This command can be used to create a Grid Infrastructure home of a supported version. For details, see `dbaascli gridHome create --help`.
  - `dbaascli system getGridHomes` - This command gives details on the available Grid Infrastructure homes on the system. For details, see `dbaascli system getGridHomes --help`.
  - `dbaascli admin updateAHF` - This command can be used to update the AHF to a specified cloud certified version of AHF release. It is however recommended that AHF

updates be implicitly handled by cloud automation. For details, see `dbaascli admin updateAHF --help`.

- Pluggable Database Operations
  - Improvements in the area of refreshable pluggagble database lifecycle.
- Database Backup and Recovery
  - Added support to configure backups on standby sites in case of dataguard configurations. The backups configuration are Data Guard site-specific, that is, the change of roles (for example, with Data Guard switchover operation) will not impact the backup operations of the database on primary or standby sites. Backups, if configured on primary site or stand-by site, will continue regardless of the role-change.
- Various bug fixes and stability improvements

Release 23.2.1.0.0 (230503)

- Database Lifecycle related improvements
  - Introduced `dbaascli grid removeTCPSCert` to remove expired TCPS certificates. For details, see `dbaascli grid removeTCPSCert --help`.
  - Added option to exclude specific PDBs during database duplicate. For details, see `skipPDBs` argument in `dbaascli database duplicate --help`.
- Database Backup and Recovery
  - Changed the default for `FILES_PER_SET` to 64 for OSS backups. This can be changed with `dbaascli database backup --configure`. For details, see `dbaascli database backup --help`.
  - Archive log backups continue from the standby site after the role switchover in data guard environments.
  - For backups that are not managed by Oracle, the schedules for L0 and L1 backups are not created by default. They must be be created explicitly by using `dbaascli database backup --configure` command.
- sysLens
  A framework that collects, analyzes, and reports system resource data for ExaDB-D fleets is included in 23.2.1.0.0 (235503). For more information, see Manage sysLens.
- Various bug fixes and stability improvements

Release 23.1.2.0.0 (230305)

- Database Lifecycle related improvements
  - Added option to create database templates (DBCA temapltes) to object store. DBCA templates can subsequently be used to create databases. For details, see `dbaascli database createTemplate --help`.
- Pluggable Database Operations
  - Introduced `dbaascli pdb refresh` to refresh a pluggable database that was created using manual refresh option. For details, see `dbaascli pdb refresh --help`.
  - Added option to convert refreshable pluggable database to a regular pluggable database. For details, see `dbaascli pdb open --help`.

  - Creation of a refreshable pluggable database now requires existing source database user for creation of database link to the source pluggable database. For details, see `dblinkUserName` argument in `dbaascli pdb remoteClone --help`.
- Various bug fixes and stability improvements

Release 23.1.1.0.1 (230113)

- Database Lifecycle related improvements
  - Added support to create a duplicate database from a source database which uses OCI Vault Services for encryption key management.
- Various bug fixes and stability improvements

Release 22.4.1.0.1 (221122)

- Pluggable Database Operations
  - Added option to not open the PDB at the end of relocate. For details, see `skipOpenPDB` argument in `dbaascli pdb relocate --help`. After using this option, the pdb relocate can be completed by running the command using `completePDBRelocate` argument.
  - Added option to clean up the relocated PDB metadata/services at the source location. For details, see `cleanupRelocatedPDB` argument in `dbaascli pdb delete --help`
- New `dbaascli` commands
  - `dbaascli database createTemplate` - This command can be used to create database templates (DBCA templates) that can subsequently be used to create databases. DBCA templates are widely used for creating a clone database with DBCA - a tool that is shipped with Oracle Database server software. For details, see `dbaascli database createTemplate --help`
  - Introduced `dbaascli tde rotateMasterKey` to rotate the master key for database encryption. For details, see `dbaascli tde rotateMasterKey --help`. The command `dbaascli tde rotate masterkey` is now deprecated.
- Database Lifecycle related improvements
  - Added support to use dbca templates in database creation workflows. For details, see `dbcaTemplateFilePath` argument in `dbaascli database create --help`
  - Improved performance for duplicate database creation. For details on how to create duplicate database, see `dbaascli database duplicate --help`
  - Added support to create a duplicate database from a source database which is not TDE-encrypted.
- TDE management
  - Introduced `dbaascli tde rotateMasterKey` to rotate the master key for database encryption. For details, see `dbaascli tde rotateMasterKey --help`. The command `dbaascli tde rotate masterkey` is now deprecated.
  - Revamped workflow for all TDE operations. For details, see `dbaascli tde --help`
- Grid Infrastructure (GI) Patching
  - Added support to allow parallel execution of patching operation on nodes. This option needs to be carefully exercised as it results into reduced database availability.
- Database Backup and Recovery

  - Revamped workflow for creating database from standalone backups
- Includes AHF version 22.2.4
- Various bug fixes and stability improvements

Release 22.3.1.1.0 (221003)

- New `dbaascli` commands
  - `dbaascli database getDetails` - This command shows the detailed information of a given database, for example, dbname, node information, pluggable databases information, and so on. For details, see `dbaascli database getDetails --help`.
- Pluggable Database Operations
  - Added support for creating pluggable databases as refreshable clone using `refreshablePDB` argument. For details, see `dbaascli pdb remoteClone --help`
- Various bug fixes and stability improvements

Release 22.3.1.0.1 (220721)

- New database lifecycle commands
  - `dbaascli database addInstance` - This command can be used to add a database instance to one of the nodes of the cluster where database is not already configured. For details, see `dbaascli database addInstance --help`.
  - `dbaascli database deleteInstance` - This command can be used to delete a database instance from one of the nodes of the cluster where database is configured. For details, see `dbaascli database deleteInstance --help`.
  - `dbaascli database duplicate` - This command can be used to create a new database from an already existing database within a cluster, or across clusters, provided network connection exists between the clusters. For details, see `dbaascli database duplicate --help`.
- Cloud Software Library
  - Introduced `dbaascli cswlib listLocal` command to list images that are downloaded from software library locally on the system. For details, see `dbaascli cswlib listLocal --help` The command `dbaascli dbimage list` is now deprecated.
  - Introduced `dbaascli cswlib deleteLocal` command to delete images that are downloaded from cloud software library. For details, see `dbaascli cswlib deleteLocal --help` The command `dbaascli dbImage purge` is now deprecated.
- The log location for the command `dbaascli admin updateStack` has been changed to follow the convention of other `dbaascli` commands. The logs can be conveniently found under `/var/opt/oracle/log/admin/updateStack` directory. The earlier location was `/var/opt/oracle/log/tooling/Update`.
- `dbaascli` help is now cloud platform aware in that it will list help output for commands applicable for the cloud environment it is operating on.
- Added support for changing TDE password in dataguard environments. For details, see dbaascli tde changePassword --help. This support is currently not available for 11.2.0.4 release.
- Included AHF version 22.1.5.
- Revamped workflow for database upgrade operation.

- Revamped workflow for database home create operation.
- Various bug fixes and stability improvements

Release 22.2.1.1.0 (220623)

- Included AHF version 22.1.1
- Fixed an issue where update of dbaastools rpm on the system may have resulted into database downtime with ORA-600 error
- Various bug fixes and stability improvements

Release 22.2.1.1.0 (220609)

- New `dbaascli` commands:
  - `dbaascli dbHome getDatabases` - This command lists all the databases running from a given database Oracle home. The output is returned in JSON format to facilitate automation. For details, see `dbaascli dbHome getDatabases --help`.
  - `dbaascli database getPDBs` - This command lists all the pluggable databases of a given container database. The output is returned in JSON format to facilitate automation. For details, see `dbaascli database getPDBs --help`.
  - `dbaascli dbHome delete` - This command deletes a given database Oracle home. For details, see `dbaascli dbHome delete --help`.
  - `dbaascli dataguard prepareStandbyBlob` - This command generates a blob file containing various files that are required on the standby site for a Data Guard environment. For details, see `dbaascli dataguard prepareStandbyBlob --help`.
- Grid Infrastructure (GI) Patching:
  - New optimized workflow
  - Introduced a way to create the Grid Infrastructure (GI) software image prior to patching. This GI image can be subsequently used for performing the GI patching operation. The advantage of this approach is that it results in reduced patching window as the image is already prepared. The GI stack on the node is not brought down to create the image. For details, see `createImage` option in `dbaascli grid patch --help`
  - Introduced a way to perform the Grid Infrastructure patching through the use of user specified GI software image, created using `createImage` option of the `dbaascli grid patch` command. For details, see `imageLocation` option in `dbaascli grid patch --help`.
- Change Password support in Data Guard environment:
  - Added support to change password in Data Guard environments. For details, see `dbaascli database changePassword --help` and `dbaascli dataguard prepareStandbyBlob --help`
- Data Guard configuration:
  - Added support to update Data Guard Automation Attributes (in the `/var/opt/oracle/dg/dg.conf` file). For details, see `dbaascli dataguard --help`.
- Various bug fixes and stability improvements

Release 22.2.1.0.1 (220423)

- New dbaascli commands
  - Introduced *dbaascli admin showLatestStackVersion* to show the latest dbaastools version available for customers to download and install. The installation of dbaastools rpm can be performed by using the command *dbaascli admin updateStack*. For details see “dbaascli Command Reference” section.
- Cloud Software Library
  - Deprecated the support for BP activation (*dbaascli cswlib activateBP)* as BPs (Bundle Patches) are now replaced with RUs (“Release Updates”). Cloud deployment consumes RUs in the form of software images, identified with “Image Tags”. It is therefore recommended to use image tags while interfacing with Cloud Software Library (cswlib) commands. For details, see *dbaasscli cswlib download –help*.
  - Eliminated the need to download Non-CDB images to create nonCDB databases. Now users can create the nonCDB database using regular images. For details, see *createAsCDB* option in *dbaascli database create –help*.
- Non-CDB Database Creation
  - Enhanced database creation workflow to create a nonCDB database using standard database software image. For details, see *createAsCDB* option in *dbaascli database create –help*.
- Database Home Patching
  - New optimized workflow
- Grid Infrastructure Upgrade
  - New optimized workflow
- Pluggable Database (PDB) Operations
  - Deletion of PDB in DataGuard environments requires explicit acknowledgement to indicate that operations necessary on standby site are completed, by passing of additional argument –allStandByPrepared. For details, see *dbaascli pdb delete --help*
- Provided rolling capability for database bounce operation. For details, see *dbaascli database bounce –help*.
- Various bug fixes and stability improvements

Release 22.1.1.2.0 (220405)

- Added support for ExaDB-D X9M
- Various bug fixes and stability improvements

Release 22.1.1.1.0 (220317)

- New `dbaascli` commands:
  - Introduced `dbaascli system getDBHomes` to get all the database Oracle homes on the cluster. The output is returned in JSON format to facilitate automation.
  - Introduced `dbaascli dbhome getDetails` to get detailed information on a specific Oracle home. The output is returned in JSON format to facilitate automation.

- Cloud Software Library (cswlib):
  - Deprecated the support for `dbaascli cswlib list` command for cloud software library listing operations. The new command is `dbaascli cswlib showImages` that lists the images along with its of `ImageTag`. It is recommended to use `Image tags` to download the images from the cloud software library. For details on downloads using image tags, see `dbaascli cswlib download -help`.
  - Various bug fixes and stability improvements

Release 22.1.1.0.1 (220223)

- Grid Infrastructure Upgrade
  - New optimized workflow
- Database Backup And Recovery
  - Internal update to metadata repository for backup metadata
  - Introduced deprecation messages for bkup_api commands as they are now replaced with dbaascli commands. For details, see 'dbaascli database backup --help' and 'dbaascli database recover –help'
- Pluggable Database (PDB) Operations
  - Relocate operation of PDB is now supported. For details, see 'dbaascli pdb relocate –help'.
  - Revamped workflow for nonCDB to PDB conversion. For details, see 'dbaascli database convertToPDB –help'.
- Encryption Key Management
  - Transparent Data Encryption (TDE) heartbeat specific initialization parameters are set to the cloud recommended values for databases with 'Customer Managed Keys'.
- Cloud Software Library Management
  - Revamped software library download of artifacts through imageTags. It is recommended to use imageTags to download the database and grid software images. For details, see 'dbaascli cswlib showimages' and 'dbaascli cswlib download –help'
- Included AHF version 21.4.2
- Various bug fixes and stability improvements

Release 21.4.1.1.0 (220209)

- Included AHF version 21.4.1
- Bug fixes and stability improvements

Release 21.4.1.1.0

- Enabled encryption of the system level tablespaces (`SYSTEM`, `SYSAUX`, `UNDO`, and `TEMP`) for databases that will get created with this version of dbaastools onwards. This feature is enabled for Oracle Database version 19.6.0.0.0 and above.
- Grid Patching:
  - Prerequisite condition added to check for following file ownership to be owned by `grid` user.

    * `<gi_home>/suptools/tfa/release/tfa_home/jlib/jdev-rt.jar`

    * `<gi_home>/suptools/tfa/release/tfa_home/jlib/jewt4.jar`

- Database Patching:
  - Simultaneous `database move` operation is disallowed by default. A new option `-allowParallelDBMove` is introduced that can be used to override the default behavior for Oracle Database releases 12.2 and above.
  - Fixed issues related to move of standby databases being in `MOUNT` mode.
- Database Backup and Recovery:
  - Added new command-line options for database backup. For more details, refer to *dbaascli database backup* command reference.
  - Added new command-line options for database recovery. For more details, refer to *dbaascli database recover* command reference.
  - `bkup_api` usage for backup and recovery operations will be deprecated in future.
  - To align with the Oracle recommended practice of using `SYSBACKUP` administrative privilege for Backup and Recovery operations, cloud automation creates a common administrative user `C##DBLCMUSER` with `SYSBACKUP` role at the `CDB$ROOT` container level. Backup and Recovery operations are therefore performed with the user having the least required privileges. Credentials for this user are randomly generated and securely managed by cloud automation. If the user is not found or is `LOCKED` and `EXPIRED`, then cloud automation will recreate or unlock this user during the backup or recovery operation. This change in the cloud automation is made starting with dbaastools version 21.4.1.1.0.
- Enhanced `dbaascli resume` functionality to resume any previous session by specifying the `-sessionID <value>` argument to the resume command. The session ID is shared in the `dbaascli` output as well as in the logs.
- Enhanced `dbaascli help` output to show the command usage.
- Deprecated the usage of `dbaascli` shell (interactive session). This will be completely unsupported after March 2022. It is recommended to execute complete `dbaascli` commands on command prompt as suggested in all document examples.
- Included Autonomous Health Framework (AHF) version 21.2.8.
- Various bug fixes and stability improvements.

Release 21.3.1.2.0

- Improved the timing of `dbaascli` operations with enhanced Control Plane metadata synchronization logic.
- Enhanced `dbaascli` logs to have millisecond-level information along with the associated thread.
- Introduced more prerequisite checks in database home patching and database move operations to catch potential failures scenarios with suggestions to corrective action.
- Database patching operations now retain the state of the databases to be same as it was prior to patching. For pluggable databases, pdb saved state is honored.
- Various bug fixes and stability improvements.

Release 21.3.1.1.0

- Added support to unlock PDB Admin user account as part of PDB creation, `localClone`, or `remoteClone` operation. For details, see option `--lockPDBAdminAccount` in `dbaascli pdb create --help`.
- Fixed an issue that updates the database resource registered with Oracle Grid Infrastructure in existing environments with the correct value of database name.
- Enhanced PDB lifecycle operations.
- Various bug fixes and stability improvements.

Release 21.3.1.0.1

- Support for the following `dbaascli` commands to be run as `oracle` user.
  - `dbaascli pdb bounce`
  - `dbaascli pdb close`
  - `dbaascli pdb connectString`
  - `dbaascli pdb create`
  - `dbaascli pdb delete`
  - `dbaascli pdb getDetails`
  - `dbaascli pdb list`
  - `dbaascli pdb localClone`
  - `dbaascli pdb open`
  - `dbaascli pdb remoteClone`
- Revamped out-of-place patching of database. For details, see `dbaascli database move -help`.
- Timing related enhancements in Oracle Grid Infrastructure patching workflow. For details, see `dbaascli grid patch -help`.
- Deprecated the support for `exadbcpatchmulti` / `dbaascli patch` for patching operations. The `dbaascli dbhome patch` and `dbaascli grid patch` commands are provided for patching operation for database homes and Oracle Grid Infrastructure. Refer to the *Patching Oracle Grid Infrastructure* and *Oracle Database Using dbaascli* section for details. Also see, *dbaascli Command Reference* section.
- Deprecated the support for `dbaascli` tools patch command to bring consistency in the `dbaascli` command conventions. The new command is `dbaascli admin updateStack`. For details, see section *Updating Cloud Tooling using dbaascli*.
- Ability to run `dbaascli` in disconnected mode for long running operations. Executing `dbaascli` command with `--waitForCompletion false` gets you a job ID that can be queried later to get the status of the operation, using `dbaascli job getStatus -jobid job_id`. This is useful for long running operations where users may want to get the control back immediately after command execution. In this release, this option is available only for `dbaascli database create` command. More commands will be added in subsequent releases to have this support. The help output for those commands will reflect the support of `--waitForCompletion` option.

- Deprecated the support for `dbaascli` shell. It is recommended that users run the complete `dbaascli` commands on the command prompt as suggested in all the document examples. Execution of just `dbaascli` will show the output of its usage help instead of entering into a `dbaascli` shell.
- Various bug fixes and stability improvements.

Release 21.2.1.x.x

- Redesigned Oracle Grid Infrastructure patching operation and added ability to resume from failed point, patch on subset of nodes, instance draining, and other enhancements. For details, see `dbaascli grid patch --help`. Also refer to the *Patching Oracle Grid Infrastructure and Oracle Database Using dbaascli* section.
- Deprecated the support for `exadbcpatchmulti` / `dbaascli patch` for patching operations. `dbaascli dbhome patch` and `dbaascli grid patch` commands are provided for patching operation for database homes and Oracle Grid Infrastructure. Refer to the *Patching Oracle Grid Infrastructure and Oracle Database Using dbaascli* section for details. Also see, *dbaascli Command Reference* section.
- Deprecated the support for `dbaascli tools patch` command to bring consistency in the command conventions. The new command is `dbaascli admin updateStack`.
- Redesigned PDB management APIs for create, local clone, and remote clone operations. For details, see `dbaascli pdb --help`.
- Redesigned database delete API. For details, see `dbaascli database delete --help`.
- Revamped dbhome creation (support for custom software image, scale-out operation). For details, see `dbaascli dbhome create --help`.
- Support for database creation on subset of cluster nodes. For details, see `dbaascli database create --help`.
- Ability to run `dbaascli` in disconnected mode for long running operations. Executing `dbaascli` command with `--waitForCompletion false` gets you a job ID that can be queried later to get the status of the operation, using `dbaascli job getStatus –jobid job_id`. This is useful for long running operations where users may want to get the control back immediately after command execution. In this release, this option is available only for `dbaascli database create` command. More commands will be added in subsequent releases to have this support. The help output for those commands will reflect the support of `--waitForCompletion` option.
- Enhanced dbhome patching experience with introduction of multiple options like `skipPDBs`, `continueWithDowntime`, and so on. For details, see `dbaascli dbhome patch --help`.
- Support for better diagnostic collection. For details, see `dbaascli diag collect --help`.
- Minor improvements in the area of database upgrade automation.
- Various bug fixes and stability improvements.

dbaascli Command Reference

You must use `dbaascli` to create databases and integrate them with the cloud automation framework.

`dbaascli` is a cloud native interface that can take DBCA templates as inputs, calls the functionality of DBCA to create databases, and then calls OCI APIs to integrate the database

into the cloud automation framework. Customers using DBCA in scripts today can update their existing scripts to call `dbaascli` instead of DBCA. If `dbaascli` cannot be used due to a particular feature of DBCA being unavailable in dbaascl, then customers should open a My Oracle Support (MOS) request to add that functionality to `dbaascli`.

To use the `dbaascli` utility, you must be connected to an Exadata Cloud Infrastructure compute node. See Connecting to an Exadata Cloud Infrastructure Instance for instructions.

Some `dbaascli` commands can be run as the `oracle` or the `opc` user, but many commands require `root` administrator privileges. Refer to each command for specific requirements.

- dbaascli admin updateAHF
  To install or update Autonomous Health Framework (AHF), use the `dbaascli admin updateAHF` command.
- dbaascli admin updateStack
  To install or update a dbaastools RPM, use the `dbaascli admin updateStack` command.
- dbaascli cswlib deleteLocal
  To delete the local image, use the `dbaascli cswlib deleteLocal` command.
- dbaascli cswlib download
  To download available software images and make them available in your Exadata Cloud Infrastructure environment, use the `dbaascli cswlib download` command.
- dbaascli cswlib listLocal
  To view the list of locally available Database and Grid Infrastructure images, use the `dbaascli cswlib listLocal` command.
- dbaascli cswlib showImages
  To view the list of available Database and Grid Infrastructure images, use the `dbaascli cswlib showImages` command.
- dbaascli database addInstance
  To add the database instance on the specified node, use the `dbaascli database addInstance` command.
- dbaascli database backup
  To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.
- dbaascli database bounce
  To shut down and restart a specified Exadata Cloud Infrastructure database, use the `dbaascli database bounce` command.
- dbaascli database changepassword
  To change the password of a specified Oracle Database user, use the `dbaascli database changePassword` command. When prompted enter the user name for which you want to change the password and then enter the password.
- dbaascli database convertToPDB
  To convert the specified non-CDB database to PDB, use the `dbaascli database convertToPDB` command.
- dbaascli database create
  To create Oracle Database, use the `dbaascli database create` command. When prompted, enter the `sys` and `tde` passwords.
- dbaascli database delete
  To delete an Oracle Database, use the `dbaascli database delete` command.

- dbaascli database deleteInstance
  To delete the database instance on the specified node, use the `dbaascli database deleteInstance` command.
- dbaascli database duplicate
  To create a database from an active database, use the `dbaascli database duplicate` command.
- dbaascli database getDetails
  This command shows the detailed information of a given database e.g. dbname, node information, pluggable databases information etc.
- dbaascli database getPDBs
  To view the list of all pluggable databases in a container database, use the `dbaascli database getPDBs` command.
- dbaascli database modifyParameters
  To modify or reset initialization parameters for an Oracle Database, use the `dbaascli database modifyParameters` command.
- dbaascli database move
  To move the database from one home to another, use the `dbaascli database move` command.
- dbaascli database recover
  To recover a database, use the `dbaascli database recover` command.
- dbaascli database runDatapatch
  To patch an Oracle Database, use the `dbaascli database runDatapatch` command.
- dbaascli database createTemplate
  Use this command to create database templates (DBCA templates) that can subsequently be used to create databases.
- dbaascli database start
  To start an Oracle Database, use the `dbaascli database start` command.
- dbaascli database status
  To check the status of an Oracle Database, use the `dbaascli database status` command.
- dbaascli database stop
  To stop an Oracle Database, use the `dbaascli database stop` command.
- dbaascli database upgrade
  To upgrade an Oracle Database, use the `dbaascli database upgrade` command.
- dbaascli dataguard prepareStandbyBlob
  To generate a blob file containing various files that are required on the standby site in case of a dataguard environment, use the `dbaascli dataguard prepareStandbyBlob` command.
- dbaascli dataguard updateDGConfigAttributes
  To update Data Guard automation attributes across all the cluster nodes, use the `dbaascli dataguard updateDGConfigAttributes` command.
- dbaascli dbhome create
  To create an Oracle Database home of desired version, use the `dbaascli dbhome create` command.
- dbaascli dbHome delete
  To delete a given Oracle Database home, use the `dbaascli dbHome delete` command.

- dbaascli dbhome getDatabases
  To view information about all Oracle Databases running from a given database Oracle home, use the `dbaascli dbHome getDatabases` command. Specify either the Oracle home location or Oracle home name.
- dbaascli dbHome getDetails
  To view information about a specific Oracle home, use the `dbaascli dbHome getDetails` command. Specify either the Oracle home location or Oracle home name.
- dbaascli dbHome patch
  To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.
- dbaascli dbimage purge
  The `dbimage purge` command removes the specified software image from your Exadata Cloud Infrastructure environment.
- dbaascli diag collect
  To collect diagnostics, use the `dbaascli diag collect` command.
- dbaascli diag healthCheck
  To run diagnostic health checks, use the `dbaascli diag healthCheck` command.
- dbaascli gridHome create
  To configure Grid Infrastructure home, use the `dbaascli gridHome create` command.
- dbaascli grid configureTCPS
  To configure TCPS for the existing cluster, use the `dbaascli grid configureTCPS` command.
- dbaascli grid patch
  To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.
- dbaascli grid removeTCPSCert
  To remove existing TCPS certificates from Grid Infrastructure wallet, use the `dbaascli grid removeTCPSCert` command.
- dbaascli grid rotateTCPSCert
  To rotate TCPS certificates, use the dbaascli grid rotateTCPSCert command.
- dbaascli grid upgrade
  To upgrade Oracle Grid Infrastrucure from one major version to another, use the `dbaascli grid upgrade` command.
- dbaascli job getStatus
  To view the status of a specified job, use the `dbaascli job getStatus` command.
- dbaascli patch db apply
- dbaascli patch db prereq
- dbaascli pdb backup
  To backup a pluggable database (PDB), query PDB backups, and delete a PDB backup, use the `dbaascli pdb backup` command.
- dbaascli pdb bounce
  To bounce a pluggable database (PDB), use the `dbaascli pdb bounce` command.
- dbaascli pdb close
  To close a pluggable database (PDB), use the `dbaascli pdb close` command.

- dbaascli pdb getConnectString
  To display Oracle Net connect string information for a pluggable database (PDB) run the `dbaascli pdb getConnectString` command.
- dbaascli pdb create
  To create a new pluggable database (PDB), use the `dbaascli pdb create` command.
- dbaascli pdb delete
  To delete a pluggable database (PDB) run the `dbaascli pdb delete` command.
- dbaascli pdb getDetails
  To view details of a pluggable database (PDB), use the `dbaascli pdb getDetails` command.
- dbaascli pdb list
  To view the list of pluggable databases (PDB) in a container database, use the `dbaascli pdb list` command.
- dbaascli pdb localClone
  To create a new pluggable database (PDB) as a clone of an existing PDB in the same container database (CDB), use the `dbaascli pdb localClone` command.
- dbaascli pdb open
  To open a pluggable database (PDB), use the `dbaascli pdb open` command.
- dbaascli pdb recover
  To recover a pluggable database (PDB), use the `dbaascli pdb recover` command.
- dbaascli pdb refresh
  To refresh a specified pluggable database (PDB), use the `dbaascli pdb refresh` command.
- dbaascli pdb relocate
  To relocate the specified PDB from the remote database into local database, use the `dbaascli pdb relocate` command.
- dbaascli pdb remoteClone
  To create a new pluggable database (PDB) as a clone of an existing PDB in another container database (CDB), use the `dbaascli pdb remoteClone` command.
- dbaascli system getDBHomes
  To view information about all the Oracle homes, use the `dbaascli system getDBHomes` command.
- dbaascli system getGridHomes
  To list the details of all Grid homes, use the `dbaascli system getGridHomes` command.
- dbaascli tde changePassword
  To change TDE keystore password as well as DB wallet password for the alias `tde_ks_passwd`, use the `dbaascli tde changePassword` command.
- dbaascli tde addSecondaryHsmKey
  To add a secondary HSM (KMS) key to the existing HSM (KMS) configuration, use the `dbaascli tde addSecondaryHsmKey` command.
- dbaascli tde enableWalletRoot
  To enable `wallet_root` spfile parameter for the existing database, use the `dbaascli tde enableWalletRoot` command.
- dbaascli tde encryptTablespacesInPDB
  To encrypt all the tablespaces in the specified PDB, use the `dbaascli tde encryptTablespacesInPDB` command.

- dbaascli tde fileToHsm
  To convert FILE based TDE to HSM (KMS/OKV) based TDE, use the `dbaascli tde fileToHsm` command.
- dbaascli tde getHsmKeys
  To get TDE active key details, use the `dbaascli tde getHsmKeys` command.
- dbaascli tde getMkidForKeyVersionOCID
  To get Master Key ID associated with the KMS key version OCID, use the `dbaascli tde getMkidForKeyVersionOCID` command.
- dbaascli tde getPrimaryHsmKey
  To get primary HSM (KMS) key from the existing HSM (KMS) configuration, use the `dbaascli tde getPrimaryHsmKey` command.
- dbaascli tde hsmToFile
  To convert HSM (KMS/OKV) based TDE to FILE based TDE, use the `dbaascli tde hsmToFile` command.
- dbaascli tde listKeys
  To list TDE master keys, use the `dbaascli tde listKeys` command.
- dbaascli tde removeSecondaryHsmKey
  To remove secondary HSM (KMS) key from the existing HSM (KMS) configuration, use the `dbaascli tde removeSecondaryHsmKey` command.
- dbaascli tde rotateMasterKey
  Rotate the master key for database encryption.
- dbaascli tde setKeyVersion
  To set the version of the primary key to be used in DB/CDB or PDB, use the `dbaascli tde setKeyVersion` command.
- dbaascli tde setPrimaryHsmKey
  To change the primary HSM (KMS) key for the existing HSM (KMS) configuration, use the `dbaascli tde setPrimaryHsmKey` command.
- dbaascli tde status
  To display information about the keystore for the specified database, use the `dbaascli tde status` command.

dbaascli admin updateAHF

To install or update Autonomous Health Framework (AHF), use the `dbaascli admin updateAHF` command.

Prerequisites

Run the command as the `root` user.

Syntax

```
dbaascli admin updateAHF
 {
   --targetVersion value | --imageTag value
}
[--resume [--sessionID value]] [--executePrereqs]
```

Where:

- `--targetVersion` specifies the target version to update AHF to
- `--imageTag` specifies the image tag of the AHF artifact to be installed
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--executePrereqs` runs prerequisite checks and reports the results

dbaascli admin updateStack

To install or update a dbaastools RPM, use the `dbaascli admin updateStack` command.

Prerequisites

Run the command as the `root` user.

To use the utility, you must connect to an Exadata Cloud Infrastructure virtual machine.

See, *Connecting to a Virtual Machine with SSH*.

Syntax

```
dbaascli admin updateStack
[--resume]
[--prechecksOnly]
[--nodes]
```

Where:

- `--resume` resumes the previous execution
- `--prechecksOnly` runs only the prechecks for this operation
- `--nodes` specifies a comma-delimited list of nodes to install the RPM on. If you do not pass this argument, then the RPM will be installed on all of the cluster nodes

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

dbaascli cswlib deleteLocal

To delete the local image, use the `dbaascli cswlib deleteLocal` command.

Run the command as the `root` user.

Syntax

```
dbaascli cswLib deleteLocal --imageTag <value>
```

Where:

- `--imageTag` specifies Oracle home image tag

Example 6-4 dbaascli cswlib deletelocal

```
dbaascli cswlib deletelocal --imagetag 19.15.0.0.0
DBAAS CLI version MAIN
Executing command cswlib deletelocal --imagetag 19.15.0.0.0
Job id: 8b3e71de-4b81-4832-b49c-7f892179bb4f
Log file location: /var/opt/oracle/log/cswLib/deleteLocal/
dbaastools_2022-07-18_10-00-02-AM_73658.log
dbaascli execution completed
```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

dbaascli cswlib download

To download available software images and make them available in your Exadata Cloud Infrastructure environment, use the `dbaascli cswlib download` command.

Prerequisites

Run the command as the `root` user.

To use the utility, you must connect to an Exadata Cloud Infrastructure virtual machine.

See, *Connecting to a Virtual Machine with SSH*.

Syntax

```
dbaascli cswlib download --version | --imageTag
[--product]
```

Where:

- `--version` specifies an Oracle home image version
- `--imageTag` specifies the image tag of the image
- `--product` specifies the image type. Valid values: `database` or `grid`

Example 6-5 dbaascli cswlib download --product --imageTag

```
dbaascli cswlib download --product database --imageTag 19.14.0.0.0
```

Example 6-6 dbaascli cswlib download --version 19.9.0.0.0

```
dbaascli cswlib download --product database --imageTag 19.14.0.0.0
```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

dbaascli cswlib listLocal

To view the list of locally available Database and Grid Infrastructure images, use the `dbaascli cswlib listLocal` command.

Run the command as the `root` user.

Syntax

```
dbaascli cswLib listLocal [--product <value>]
```

Where:

- `--product` identifies Oracle home product type. Valid values: `database` or `grid`.

Example 6-7 dbaascli cswlib listlocal

```
dbaascli cswlib listlocal
DBAAS CLI version MAIN
Executing command cswlib listlocal
Job id: bc4f047c-0a34-4d4d-a1ea-21ddc2a9c627
Log file location: /var/opt/oracle/log/cswLib/listLocal/
dbaastools_2022-07-18_10-29-53-AM_16077.log
############ List of Available Database Images  #############
1.IMAGE_TAG=12.2.0.1.220419
  IMAGE_SIZE=5GB
  VERSION=12.2.0.1.220419
  DESCRIPTION=12.2 APR 2022 DB Image
2.IMAGE_TAG=18.16.0.0.0
  IMAGE_SIZE=6GB
  VERSION=18.16.0.0.0
  DESCRIPTION=18c OCT 2021 DB Image
3.IMAGE_TAG=19.14.0.0.0
  IMAGE_SIZE=5GB
  VERSION=19.14.0.0.0
  DESCRIPTION=19c JAN 2022 DB Image
dbaascli execution completed
```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

dbaascli cswlib showImages

To view the list of available Database and Grid Infrastructure images, use the `dbaascli cswlib showImages` command.

Run the command as the `root` user.

Syntax

```
dbaascli cswlib showImages
[--product]
```

Where:

- `--product` identifies Oracle home product type. Valid values: `database` or `grid`.

Example 6-8 dbaascli cswlib showImages

```
dbaascli cswlib showImages
```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

dbaascli database addInstance

To add the database instance on the specified node, use the `dbaascli database addInstance` command.

Prerequisite

- Run the command as the `root` user.

Syntax

```
dbaascli database addInstance --dbname <value> --node <value> [--newNodeSID
<value>]
```

Where:

- `--dbname` specifies Oracle Database name
- `--node` specifies the node name for the database instance
  - `--newNodeSID` specifies SID for the instance to add in the new node

dbaascli database backup

To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

Prerequisite

- Run the command as the `root` user.

Syntax

```
dbaascli database backup --dbname <value>
        {
            --list
                {
```

```
                    [--backupType <value>]
                    | [--json <value>]
                }
            | --start [--level0] [--level1]
                {
                    [--archival --tag <value>]
                    | [--archivelog]
                }
            | --delete --backupTag <value>
            | --status --uuid <value>
            | --getBackupReport
                {
                    --tag <value>
                    | --latest
                }
                --json <value>
            | --configure
                {
                    --configFile <value>
                    | --enableRTRT
                    | --disableRTRT
                }
            | --getConfig [--configFile <value>]
            | --validate [--untilTime <value>]
            | --showHistory [--all]
        }
```

Where:

```
--dbname: Oracle Database name.
--list | --start | --delete | --status | --getBackupReport | --configure | --
getConfig
--list: Returns database backup information.
    [--json: Specify the file name for JSON output.]
--start: Begins database backup.
        [--level0 | --level1 | --archival]
        [--level0: Creates a Level-0 (full) backup. ]
        [--level1: Creates a Level-1 (incremental) backup. ]
        [--archival: Creates an Archival full backup. ]
             --tag: Specify backup tag.
--delete: Deletes Archival backup.
           --backupTag <value>
--status
           --uuid <value>
--getBackupReport: Returns backup report.
           --tag: Specify backup tag.
           --latest: Returns latest backup report (all types of database backup).
           --json: Specify the file name for JSON output.
--configure: Configures database for backup.
           --configFile | --enableRTRT | --disableRTRT
           --configFile: Specify database backup configuration file.
           --enableRTRT: Enables Real Time Redo Transport.
           --disableRTRT: Disables Real Time Redo Transport.
--getConfig: Returns database backup configuration.
           [--configFile: Specify the database backup configuration file.]
--validate: Validates that backups are complete and corruption-free.
```

```
            [--untilTime: Validates from closest Level-0 (full) backup until time
provided. Input format: DD-MON-YYYY HH24:MI:SS.]
--showHistory: Displays the history of backup operations.
            [--all: Displays all backup operations.]
```

> **Note:**
>
> `enableRTRT` and `disableRTRT` are applicable only for ZDLRA backup destination on Exadata Database Service on Cloud@Customer.

Example 6-9 Examples

- To change the archive log retention period follow the below steps:

  ```
  dbaascli database backup --getConfig --dbname <dbname>
  ```

  This will generate a backup config file `.cfg`.

  Update `bkup_archlog_fra_retention` value in this config file.

  Run the configure command:

  ```
  dbaascli database backup --configure --dbname <dbname> --configfile
  <config file generated above>
  ```

- To get backup configuration for a database *myTestDB*:

  ```
  dbaascli database backup --dbName myTestDB --getConfig --configFile /tmp/
  configfile_1.txt
  ```

- To set backup configuration for a database *myTestDB* by modifying the config file with configuration details:

  ```
  dbaascli database backup --dbName myTestDB --configure --configFile /tmp/
  configfile_1_modified.txt
  ```

- To take backup of the database *myTestDB*:

  ```
  dbaascli database backup --dbName myTestDB --start
  ```

- To query the status of backup request submitted with `uuid` *58fdcae0bd1c11eb92bc020017075151*:

  ```
  dbaascli database backup --dbName myTestDB --status --uuid
  58fdcae0bd1c11eb92bc020017075151
  ```

- To enable RTRT for the database *myTestDB*:

  ```
  dbaascli database backup --dbName myTestDB --configure —enableRTRT
  ```

dbaascli database bounce

To shut down and restart a specified Exadata Cloud Infrastructure database, use the `dbaascli database bounce` command.

Prerequisites

Run the command as the `oracle` user.

Syntax

```
dbaascli database bounce
[--dbname][--rolling <value>]
```

Where:

- `--dbname` specifies the name of the database
- `--rolling` specifies `true` or `false` to bounce the database in a rolling manner. Default value is `false`.

The command performs a database shutdown in immediate mode. The database is then restarted and opened. In Oracle Database 12c or later, all of the PDBs are also opened.

Example 6-10 dbaascli database bounce

```
dbaascli database bounce --dbname dbname
```

dbaascli database changepassword

To change the password of a specified Oracle Database user, use the `dbaascli database changePassword` command. When prompted enter the user name for which you want to change the password and then enter the password.

Prerequisites

Run the command as the `root` or `oracle` user.

Syntax

```
dbaascli database changePassword [--dbname <value>] [--user <value>]
{
  [--prepareStandbyBlob <value> [--blobLocation <value>]] | [--
standbyBlobFromPrimary <value>]
}
[--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the name of the Oracle Database that you want to act on
- `--user` specifies the user name whose password change is required
- `--prepareStandbyBlob` specifies `true` to generate a blob file containing the artifacts needed to change the password in a Data Guard environment. Valid values: `true|false`

- `--blobLocation` specifies the custom path where blob file will be generated
- `--standbyBlobFromPrimary` specifies the standby blob file, which is prepared from the primary database
- `--resume` specifies to resume the previous execution
  - `--sessionID` specifies to resume a specific session ID

Example 6-11 dbaascli database changePassword

```
dbaascli database changepassword --dbname db19
```

dbaascli database convertToPDB

To convert the specified non-CDB database to PDB, use the `dbaascli database convertToPDB` command.

Syntax

```
dbaascli database convertToPDB --dbname <value> [--cdbName <value>] [--
executePrereqs]
        {
            [--copyDatafiles [--keepSourceDB]]|[backupPrepared]
        }
        [--targetPDBName <value>] [--waitForCompletion <value>] [--resume [--
sessionID <value>]]
```

Where:

- `--dbname` specifies the name of Oracle Database
- `--cdbName` specifies the name of the target CDB in which the PDB will be created. If the CDB does not exist, then it will be created in the same Oracle home as the source non-CDB
- `--executePrereqs` specifies to run only the pre-conversion checks
- `--copyDatafiles` specifies to create a new copy of the data files instead of using the ones from the source database
  `--keepSourceDB` - to preserve the source database after completing the operation.
- `--backupPrepared` - flag to acknowledge that a proper database backup is in place for the non CDB prior to performing the conversion to PDB.
- `--backupPrepared` flag to acknowledge that a proper database backup is in place for the non-CDB prior to performing the conversion to PDB
- `--targetPDBName` specifies the name of the PDB that will be created as part of the operation
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true|false`
- `--resume` specifies to resume the previous execution
  - `--sessionID` specifies to resume a specific session ID

Example 6-12 dbaascli database convertToPDB

To run pre-conversion prechecks:

```
dbaascli database convertToPDB --dbname ndb19 --cdbname cdb19 --
backupPrepared --executePrereqs
```

To run a full conversion with a copy of the data files from the non-CDB:

```
dbaascli database convertToPDB --dbname tst19 --cdbname cdb19 --copyDatafiles
```

dbaascli database create

To create Oracle Database, use the `dbaascli database create` command. When prompted, enter the `sys` and `tde` passwords.

Use this command to create Oracle Database version 12.1.0.2 or higher with the release update JAN 2021 or higher. For databases with lower versions, it is recommended to use the OCI Console based API.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli database create --dbName {--oracleHome | --oracleHomeName}
[--dbUniqueName <value>]
[--dbSID <value>]
[--createAsCDB <value>]
[--pdbName <value>]
[--pdbAdminUserName <value>]
[--dbCharset <value>]
[--dbNCharset <value>]
[--dbLanguage <value>]
[--dbTerritory <value>]
[--sgaSizeInMB <value>]
[--pgaSizeInMB <value>]
[--datafileDestination <value>]
[--fraDestination <value>]
[--fraSizeInMB <value>]
[--nodeList <value>]
[--tdeConfigMethod <value>]
[--kmsKeyOCID <value>]
{
            [--resume [--sessionID <value>]]
            | [--revert [--sessionID <value>]]
        }
[--executePrereqs]
[--honorNodeNumberForInstance <value>]
[--lockPDBAdminAccount <value>]
[--dbcaTemplateFilePath <value>]
[--waitForCompletion]
```

Where:

- `--dbname` specifies the name of the database
- `--oracleHome` specifies the location of the Oracle home
- `--oracleHomeName` specifies the name of the Oracle home
- `--dbUniqueName` specifies database unique name
- `--dbSID` specifies the SID of the database
- `--createAsCDB` specify `true` or `false` to create database as CDB or Non-CDB
- `--pdbName` specify PDB name
- `--pdbAdminUserName` specify PDB admin user name
- `--dbCharset` specifies database character set
- `--dbNCharset` specifies database national character set
- `--dbLanguage` specifies the database language
- `--dbTerritory` specifies the database territory
- `--sgaSizeInMB` specifies the `sga_target` value in megabyte unit
- `--pgaSizeInMB` specifies the `pga_aggregate_target` value in megabyte unit
- `--datafileDestination` specifies the ASM disk group name to use for database datafiles
- `--fraDestination` specifies ASM disk group name to use for database Fast Recovery Area
- `--fraSizeInMB` specifies the Fast Recovery Area size value in megabyte unit
- `--nodeList` specifies a comma-delimited list of nodes for the database
- `--tdeConfigMethod` specifies TDE configuration method. Valid values: `FILE`, `KMS`
- `--kmsKeyOCID` specifies KMS key OCID to use for TDE. This is applicable only if KMS is selected for TDE
- `--resume` resumes the previous execution
- `--revert` rolls back the previous run
- `--sessionID` to resume or revert a specific session id.
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes` or `no`
- `--honorNodeNumberForInstance` specifies `true` or `false` to indicate instance name to be suffixed with the cluster node numbers. Default value: `true`
- `--lockPDBAdminAccount` specify `true` or `false` to lock the PDB admin user account. Default value is `true`
- `--dbcaTemplateFilePath` specify the absolute path of the dbca template name to create the database.
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`

Example 6-13 dbaascli database create

```
dbaascli database create --dbName db19 --oracleHomeName myhome19 --dbSid
db19sid --nodeList node1,node2 --createAsCDB true
```

dbaascli database delete

To delete an Oracle Database, use the `dbaascli database delete` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli database delete --dbname <value>
[--deleteArchiveLogs <value>]
[--deleteBackups <value>]
[--precheckOnly <value>]
[--waitForCompletion <value>]
[--force]
[--dbSID <value>]
[--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the name of the database.
- `--deleteArchiveLogs` specifies `true` or `false` to indicate deletion of database archive logs.
- `--deleteBackups` specifies `true` or `false` to indicate deletion of database backups.
- `--precheckOnly` specifies `yes` to run only the prechecks for this operation. Valid values: `yes` or `no`.
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`.
- `--force` flag to force delete database.
- `--dbSID` specify database SID.
- `--resume` to resume the previous execution.
- `--sessionID` to resume a specific session id.

Example 6-14 dbaascli database delete

```
dbaascli database delete --dbname db19
```

dbaascli database deleteInstance

To delete the database instance on the specified node, use the `dbaascli database deleteInstance` command.

Prerequisite

- Run the command as the `root` user.

Syntax

```
dbaascli database deleteInstance --dbname <value> --node <value> [--
continueOnUnreachableNode]
```

Where:

- `--dbname` specifies Oracle Database name
- `--node` specifies the node name for database instance
- `--continueOnUnreachableNode` specifies to perform the operation even if the node is unreachable

Example 6-15 database deleteinstance

```
database deleteinstance --node test-node
```

dbaascli database duplicate

To create a database from an active database, use the `dbaascli database duplicate` command.

Prerequisite

- Run the command as the `root` user.

Syntax

```
dbaascli database duplicate --dbName <value> --sourceDBConnectionString
<value>
        {
            --oracleHome <value>
            | --oracleHomeName <value>
        }
[--dbSID <value>]
[--dbUniqueName <value>]
[--sgaSizeInMB <value>]
[--pgaSizeInMB <value>]
[--datafileDestination <value>]
[--fraDestination <value>]
[--fraSizeInMB <value>]
[--sourceDBWalletLocation <value>]
[--nodeList <value>]
        {
            [--resume [--sessionID <value>]]
            | [--revert [--sessionID <value>]]
        }
[--rmanParallelism <value>]
[--rmanSectionSizeInGB <value>]
[--tdeConfigMethod <value>]
[--kmsKeyOCID <value>]
[--sourceDBTdeConfigMethod <value>]
[--sourceDBKmsKeyOCID <value>]
[--executePrereqs <value>]
```

```
[--waitForCompletion <value>]
[--skipPDBs <value>]
```

Where:

- `--dbName` specifies Oracle Database name
- `--sourceDBConnectionString` specifies source database connection string in the format of `<scan_name>:<scan_port>/<database_service_name>`
- `--oracleHome` specifies Oracle home location
- `--oracleHomeName` specifies Oracle home name
- `--dbSID` specifies database SID
- `--dbUniqueName` specifies database unique name
- `--sgaSizeInMB` specifies `sga_target` value in mega byte unit
- `--pgaSizeInMB` specifies `pga_aggregate_target` value in mega byte unit
- `--datafileDestination` specifies ASM disk group name to use for database datafiles
- `--fraDestination` specifies ASM disk group name to use for database fast recovery area
- `--fraSizeInMB` specifies fast recovery area size value in mega byte unit
- `--sourceDBWalletLocation` specifies source database TDE wallet file location. This is required to duplicate database from active database
- `--nodeList` specifies a comma-delimited list of nodes for the database
- `--resume` specifies to resume the previous execution
  - `--sessionID` specifies to resume a specific session ID
- `--revert` specifies to rollback the previous execution
  - `--sessionID` specifies to rollback a specific session ID
- `--rmanParallelism` specifies parallelsim value
- `--rmanSectionSizeInGB` specifies RMAN section size in GB
- `--tdeConfigMethod` specifies TDE configuration method. Allowed values are `FILE` and `KMS`.
- `--kmsKeyOCID` specifies KMS key OCID to use for TDE. This is applicable only if KMS is selected for TDE.
- `--sourceDBTdeConfigMethod` specifies source database TDE configuration method. Allowed values are `FILE` and `KMS`.
- `--sourceDBKmsKeyOCID` specifies source database KMS key OCID to use for TDE. This is applicable only if KMS is selected for TDE.
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes|no`
- `--waitForCompletion` specifies `false` to run the operation in background. Valid values: `true|false`
- `--skipPDBs` specifies a comma-delimited list of source database PDB names, which needs to be excluded for the duplicate database operation. Example: pdb1,pdb2...

Example 6-16 dbaascli database duplicate

```
dbaascli database duplicate --sourceDBConnectionString test-user-
scan.dbaastoolslrgsu.dbaastoolslrgvc.oraclevcn.com:1521/
mynew.dbaastoolslrgsu.dbaastoolslrgvc.oraclevcn.com --oracleHome /u02/app/
oracle/product/19.0.0.0/dbhome_2 --dbName newdup --
sourceDBWalletLocation /var/opt/oracle/dbaas_acfs/tmp/prim_wallet
```

dbaascli database getDetails

This command shows the detailed information of a given database e.g. dbname, node information, pluggable databases information etc.

Prerequisites

Run the command as the `root` user or the `oracle` user

Syntax

```
dbaascli database getDetails --dbname <value>
```

Where :

- `--dbname` - Oracle database name.

dbaascli database getPDBs

To view the list of all pluggable databases in a container database, use the `dbaascli database getPDBs` command.

Run the command as the `root` or `oracle`user.

Syntax

```
dbaascli database getPDBs --dbname <value>
```

Where:

- `--dbname` specifies the name of the container database

Example 6-17 dbaascli database getPDBs --dbname

```
dbaascli database getPDBs --dbname apr_db1
```

dbaascli database modifyParameters

To modify or reset initialization parameters for an Oracle Database, use the `dbaascli database modifyParameters` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli database modifyParameters --dbname <value> --setParameters <values>|
--resetParameters <values> | --responseFile
[--backupPrepared]
[--instance]
[--allowBounce]
```

Where:

- `--dbname` specifies the name of the database.
- `--setParameters` specifies a comma-delimited list of parameters to modify with new values. For example: `parameter1=`*`valueA`*`,parameter2=`*`valueB`*, and so on. For blank values use parameter1=valueA,parameter2='',etc.
- `--resetParameters` specifies a comma-delimited list of parameters to be reset to their corresponding default values. For example, `parameter1,parameter2`, and so on.
- `--responseFile` specifies the absolute location of the response JSON file to modify the database parameters
- `--backupPrepared` acknowledges that a proper database backup is in place prior to modifying critical or sensitive parameters.
- `--instance` specifies the name of the instance on which the parameters will be processed. If not specified, then the operation will be performed at the database level.
- `--allowBounce` grants permission to bounce the database in order to reflect the changes on applicable static parameters.

Example 6-18 dbaascli database modifyParameters

```
dbaascli database modifyParameters --dbname dbname --setParameters
"log_archive_dest_state_17=ENABLE"
```

dbaascli database move

To move the database from one home to another, use the `dbaascli database move` command.

Prerequisites

- Before performing a move operation, ensure that all of the database instances associated with the database are up and running.
- Run the command as the `root` user.

Syntax

```
dbaascli database move
{
  --oracleHome <value> | --oracleHomeName <value>
}
--dbname <value>
[--executePrereqs]
[--resume [--sessionID <value>]]
[--rollback [--sessionID <value>]]
```

```
[--skipDatapatch]
[--skipPDBs <value>]
[--skipClosedPDBs]
[--continueWithDbDowntime]
[--allowParallelDBMove]
[--waitForCompletion <value>]
[--nodeList <value>]
```

Where:

- `--oracleHome` specifies Oracle home path
- `--oracleHomeName` specifies the name of Oracle home
- `--dbname` specifies the name of the database
- `--executePrereqs` runs the prerequisite checks and report the results
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--rollback` rolls the database back to previous home
  - `--sessionID` specifies to resume a specific session ID
- `--skipDatapatch` skips running the datapatch on the databases
- `--skipPdbs` skips running the datapatch on a specified comma-delimited list of PDBs. For example: *pdb1*,*pdb2*...
- `--skipClosedPDBs` skips patching closed PDBs
- `--continueWithDbDowntime` continues patching with database downtime. This option can be used in environments wherein there is only one active instance up and the patching operation can be continued even with a downtime.
- `--allowParallelDBMove` allows database move in parallel.
- `--waitForCompletion` specifies `false` to run the operation in the background. **Valid values:** `true`|`false`
- `--nodeList` specifies a comma-delimited list of nodes if operation has to be performed on a subset of nodes

Example 6-19 dbaascli database move

```
dbaascli database move --dbname testdb1 --oracleHome /u02/app/oracle/product/
12.1.0/dbhome_2
```

dbaascli database recover

To recover a database, use the `dbaascli database recover` command.

Prerequisite

- Run the command as the `root` user.
- Database must have been configured with backup storage destination details where backups are stored.

Syntax

```
dbaascli database recover --dbname <value>
        {
            --start
                {
                    --untilTime <value>
                    | --untilSCN <value>
                    | --latest
                    | --tag <value>
                }
            | --status --uuid <value>
        }
```

Where:

```
--dbname: Oracle Database name.
      --start | --status
--start: Begins database recovery.
      --untilTime | --untilSCN | --latest | --tag
      --untilTime: Recovers database until time. Input format: DD-MON-YYYY
HH24:MI:SS.
      --untilSCN: Recovers database until SCN.
      --latest: Recovers database to last known state.
      --tag: Recovers database to archival tag.
--status
      --uuid <value>
```

Example 6-20 Examples

- To recover the database *myTestDb* to latest:

  ```
  dbaascli database recover --dbname myTestDb --start --latest
  ```

- To query the status of recovery request submitted with `uuid` *2508ea18be2911eb82d0020017075151*:

  ```
  dbaascli database recover --dbname myTestDb --status --uuid
  2508ea18be2911eb82d0020017075151
  ```

dbaascli database runDatapatch

To patch an Oracle Database, use the `dbaascli database runDatapatch` command.

Prerequisites

- Before performing a `runDatapatch` operation, ensure that all of the database instances associated with the database are up and running.
- Run the command as the `root` user.

Syntax

```
dbaascli database runDatapatch --dbname
[--resume]
```

```
    [--sessionID]
[--skipPdbs | --pdbs]
[--executePrereqs]
[--patchList]
[--skipClosedPdbs]
[--rollback]
```

Where:

- `--dbname` specifies the name of the database
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--skipPdbs` skips running the datapatch on a specified comma-delimited list of PDBs. For example: *pdb1*,*pdb2*...
- `--pdbs` runs the datapatch only on a specified comma-delimited list of PDBs. For example: *pdb1*,*pdb2*...
- `--executePrereqs` runs prerequisite checks
- `--patchList` applies or rolls back the specified comma-delimited list of patches. For example: *patch1*,*patch2*...
- `--skipClosedPdbs` skips running the datapatch on closed PDBs
- `--rollback` rolls back the patches applied

```
dbaascli database runDatapatch --dbname db19
```

dbaascli database createTemplate

Use this command to create database templates (DBCA templates) that can subsequently be used to create databases.

Run the command as the `root` or `oracle` user.

Syntax

Create a new DBCA template from the specified database.

```
dbaascli database createTemplate --dbname <value>
 {
   --templateLocation <value> | --uploadToObjectStorage --
objectStorageLoginUser <value> --objectStorageBucketName <value> [--
objectStorageUrl <value>]
 }
 [--templateName <value>] [--rmanParallelism <value>]
```

Where:

- `--dbname` specifies the name of the database
- `--templateLocation` specifies the template name
- `--uploadToObjectStorage`: specifies to upload the template to Object Storage
  - `--objectStorageLoginUser`: specifies the Object Storage login user

  - `--objectStorageBucketName`: specifies the Object Storage bucket name
  - `--objectStorageUrl`: specifies the Object Storage URL
- `--templateName`: specifies the name of the template
- `--rmanParallelism` specifies the parallelsim value

dbaascli database start

To start an Oracle Database, use the `dbaascli database start` command.

Prerequisites

Run the command as the `root` user.

Syntax

```
dbaascli database start
[--dbname]
[--mode]
```

Where:

- `--dbname` specifies the name of the database
- `--mode` specifies mount or nomount to start database in the corresponding mode

The command starts and opens the database. In Oracle Database 12c or later, all of the PDBs are also opened.

Example 6-21 dbaascli database start

```
dbaascli database start --dbname dbname --mode mount
```

dbaascli database status

To check the status of an Oracle Database, use the `dbaascli database status` command.

Prerequisites

Run the command as the `root` user.

Syntax

```
dbaascli database status
[--service][--dbname]
[--user]
[--password]
```

Where:

- `--service` specifies the name of the service
- `--dbname` specifies the name of the database
- `--user` specifies the user name of the service
- `--password` specifies the password of the user

Output from the command includes the open mode of the database, the software release and edition of the database, and release version of other software components.

Example 6-22 dbaascli database status

```
dbaascli database status --dbname db19
```

dbaascli database stop

To stop an Oracle Database, use the `dbaascli database stop` command.

Prerequisites

Run the command as the `root` user.

Syntax

```
dbaascli database stop
[--dbname <value>]
[--mode <value>]
```

Where:

- `--dbname` specifies the name of the database that you want to stop
- `--mode` specifies the mode of the database. Valid values: `abort`, `immediate`, `normal`, `transactional`

The command performs a database shutdown in immediate mode. No new connections or new transactions are permitted. Active transactions are rolled back, and all connected users are disconnected.

Example 6-23 dbaascli database stop

```
dbaascli database stop --dbname db19
```

dbaascli database upgrade

To upgrade an Oracle Database, use the `dbaascli database upgrade` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli database upgrade --dbname <value>
{--targetHome <value> | --targetHomeName <value>}
{ [--executePrereqs | --postUpgrade | --rollback]}
{[--standBy | --allStandbyPrepared]}
{[--upgradeOptions <value>]  | [--standBy]}
[--removeGRP]
[--increaseCompatibleParameter]
[--resume [--sessionID <value>]]
[--waitForCompletion <value>]
```

Where:

- `--dbname` (mandatory) specifies the name of the database.
- `--targetHome` specifies the target Oracle home location
- `--targetHomeName` specifies the name of the target Oracle Database home
- `--standBy` use this option to upgrade standby databases in Data Guard configurations
- `--allStandbyPrepared` required for Data Guard configured primary databases. Flags to acknowledge that all the required operations are performed on the standby databases prior to upgrading primary database
- `--removeGRP` automatically removes the Guaranteed Restore Point (GRP) backup only if the database upgrade was successful
- `--increaseCompatibleParameter` automatically increases the compatible parameter as part of the database upgrade. The parameter will get increased only if the database upgrade was successful
- `--executePrereqs` runs only the preupgrade checks
- `--postUpgrade` use this option if postupgrade fails and needs to rerun the postupgrade steps
- `--rollback` reverts an Oracle Database to its original Oracle home
- `--upgradeOptions` use this option to pass DBUA-specific arguments to perform the Oracle Database upgrade. Refer to the corresponding Oracle documentation for the supported arguments and options.
  `--standby`
- `--resume` to resume the previous execution
- `--sessionID` to resume a specific session id.
- `--waitForCompletion` specify false to run the operation in background. Valid values : true| false.

Example 6-24 dbaascli database upgrade pre-upgrade requisite checks

```
dbaascli database upgrade --dbbname dbname --targetHome Target Oracle home
location --executePrereqs
```

dbaascli dataguard prepareStandbyBlob

To generate a blob file containing various files that are required on the standby site in case of a dataguard environment, use the `dbaascli dataguard prepareStandbyBlob` command.

Run the command as the `root` or `oracle` user.

Syntax

```
dbaascli dataguard prepareStandbyBlob --dbname <value> --blobLocation <value>
```

Where:

- `--dbname` specifies the Oracle Database name
- `--blobLocation` specifies the custom directory location where the standby blob file will be generated in a Data Guard environment

dbaascli dataguard updateDGConfigAttributes

To update Data Guard automation attributes across all the cluster nodes, use the `dbaascli dataguard updateDGConfigAttributes` command.

Run the command as the `root` or `oracle`user.

Syntax

```
dbaascli dataguard updateDGConfigAttributes --attributes <value>
```

Where:

- `--attributes` contains the Data Guard automation attributes that are to be modified. Accepts comma-delimited values in the format <*attribute*=*value*>. Attributes must be predefined in the Data Guard configuration file.

dbaascli dbhome create

To create an Oracle Database home of desired version, use the `dbaascli dbhome create` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli dbhome create --version <value>
[--oracleHome <value>]
[--oracleHomeName <value>]
[--enableUnifiedAuditing <value>]
[--imageTag <value>]
[--ImageLocation <value>
```

Where:

- `--version` specifies the version of Oracle Home specified as five numeric segments separated by periods, for example, 19.12.0.0.0
- `--oracleHome` specifies the location of Oracle home
- `--oracleHomeName` specifies user-defined Oracle home name. If not provided, then the default name will be used
- `--enableUnifiedAuditing` specifies `true` or `false` to enable or disable unified auditing link option in Oracle home
- `--imageTag` specifies Oracle home image tag
- `--imageLocation` - path of the image to be used.
- `--waitForCompletion` specifies `false` to run the operation in background. Valid values: `true` or `false`.

Example 6-25 dbaascli dbhome create

```
dbaascli dbhome create --version 19.11.0.0.0
```

Alternatively, `dbaascli dbhome create --version 19.8.0.0.0 --imageTag 19.8.0.0.0` for cases where image tags are different from version.

dbaascli dbHome delete

To delete a given Oracle Database home, use the `dbaascli dbHome delete` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli dbHome delete
{ --oracleHome <value>
| --oracleHomeName <value> } [--resume [--sessionID <value>]]
```

Where:

- `--oracleHome` specifies the location of the Oracle home
- `--oracleHomeName` specifies the name of the Oracle home
- `--resume` resumes the previous execution
  - `--sessionID` specifies to resume a specific session ID

dbaascli dbhome getDatabases

To view information about all Oracle Databases running from a given database Oracle home, use the `dbaascli dbHome getDatabases` command. Specify either the Oracle home location or Oracle home name.

Run the command as the `root` user.

Syntax

```
dbaascli dbHome getDatabases
{ --oracleHomeName value | --oracleHome value }
```

Where:

- `--oracleHomeName` specifies user-defined Oracle home name
- `--oracleHome` specifies the location (path) of Oracle home

Example 6-26 dbaascli dbHome getDatabases --oracleHome

```
dbaascli dbHome getDatabases --oracleHome /u02/app/mar_home/
```

dbaascli dbHome getDetails

To view information about a specific Oracle home, use the `dbaascli dbHome getDetails` command. Specify either the Oracle home location or Oracle home name.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli dbHome getDetails
{ --oracleHomeName value | --oracleHome value }
```

Where:

- `--oracleHomeName` specifies user-defined Oracle home name
- `--oracleHome` specifies the location of Oracle home

Example 6-27 dbaascli dbHome getDetails - using Oracle home location

```
dbaascli dbHome getDetails --oracleHome /u02/app/home_db19c/
```

Example 6-28 dbaascli dbHome getDetails - using Oracle home name

```
dbaascli dbHome getDetails --oracleHomeName home_db19c
```

dbaascli dbHome patch

To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli dbHome patch --oracleHome | --oracleHomeName
--targetVersion
[--resume]
    [--sessionID]
[--continueWithDbDowntime]
[--skipUnreachableNodes]
[--nodes]
[--executePrereqs]
[--skipDatapatch]
[--imageLocation]
[--skipPDBs]
[--skipClosedPDBs]
[--rollback]
```

Where:

- `--oracleHome` specifies the path of Oracle home
- `--oracleHomeName` specifies the name of Oracle home
- `--targetVersion` specifies the target version of Oracle Home specified as five numeric segments separated by periods (e.g. 19.12.0.0.0).
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--continueWithDbDowntime` continues patching with database downtime. This option can be used in environments wherein there is only one active instance up and the patching operation can be continued even with a downtime.
- `--skipUnreachableNodes` skips operation on unreachable nodes
- `--nodes` specifies a comma-delimited list of nodes if patching has to be performed on a subset of nodes
- `--executePrereqs` runs prereqs
- `--skipDatapatch` skips running `datapatch` on the databases
- `--imageLocation` specifies custom location for database image
- `--skipPDBs` skips running the datapatch on a specified comma-delimited list of PDBs. For example: *cdb1:pdb1*,*cdb2:pdb2*, and so on
- `--skipClosedPdbs` skips running `datapatch` on closed PDBs
- `--rollback` rolls back patched Oracle home.

Example 6-29 dbaascli dbhome patch

```
dbaascli dbhome patch --targetVersion 19.10.0.0.0 --oracleHome /u02/app/
oracle/product/19.0.0.0/dbhome_2
```

dbaascli dbimage purge

The `dbimage purge` command removes the specified software image from your Exadata Cloud Infrastructure environment.

Connect to the compute node as the `opc` user and execute this command as the `root` user.

```
# dbaascli dbimage purge --version software_version --bp software_bp [--cdb ( yes | no )]
```

In the preceding command:

- *`software_version`* — specifies the Oracle Database software version. For example, `11204`, `12102`, `12201`, `18000`, `19000`.
- *`software_bp`* — identifies the bundle patch release. For example, `APR2018`, `JAN2019`, `OCT2019`, and so on.
- `--cdb` — optionally specifies whether to remove the software image that supports the Oracle multitenant architecture. Default is `yes`. If you specify `--cdb no`, then the software image that contains binaries to support non-container databases (non-CDB) is removed.

If the command will remove a software image that is not currently available in the software image library, and therefore cannot be downloaded again, then the command pauses and prompts for confirmation.

You cannot remove the current default software image for any software version. To avoid this restriction, you must make another software image the current default.

dbaascli diag collect

To collect diagnostics, use the `dbaascli diag collect` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli diag collect [--components <value>] [--startTime <value>] [--endTime
<value>] [--nodes <value>] [--dbNames <value>]
        {
            [--objectStoreBucketUri <value>]
            | [--destLocation <value>]
        }
        [--waitForCompletion <value>]
```

Where:

- `--components` specifies a list of components for log collection.
  Valid values:
  - `db`
  - `gi`
  - `os`
  - `dbaastools`
  - `all`
- `--startTime` specifies the start time for log collection. Valid date and time format: `YYYY-MM-DDTHH24:MM:SS`
- `--endTime` specifies the end time for log collection. Valid date and time format: `YYYY-MM-DDTHH24:MM:SS`
- `--nodes` specifies a comma-delimited list of nodes to collect logs
- `--dbNames` specifies the database name for which to collect logs. You can specify only one database name.
- `--objectStoreBucketURI` specifies an Object Storage service pre-authenticated request (PAR) URL used to upload collected logs. Logs are collected from Guest VM. For more information, see *Using Pre-Authenticated Requests*.
- `--destLocation` specifies the location on Guest VM to collect logs. Default: `/var/opt/oracle/dbaas_acfs`
- `--waitForCompletion` Values: `true`|`false`. Default `true`. Specify `false` to run in the background.

Related Topics

- Using Pre-Authenticated Requests

- Collecting Tooling Log Data Examples
  The dbaascli dbaascli diag collect command uses the syntax shown below to collect tooling log data:

dbaascli diag healthCheck

To run diagnostic health checks, use the `dbaascli diag healthCheck` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli diag healthCheck
[--destLocation]
[--nodes]
[--objectStoreBucketURI]
```

Where:

- `--destLocation` specifies the location on Guest VM to collect logs. Default: `/var/opt/oracle/dbaas_acfs`
- `--nodes` specifies a comma-delimited list of nodes to collect logs
- `--objectStoreBucketURI` specifies an Object Storage service pre-authenticated request (PAR) URL used to upload collected logs. Logs are collected from Guest VM. For more information, see *Using Pre-Authenticated Requests*.

Related Topics

- Using Pre-Authenticated Requests

dbaascli gridHome create

To configure Grid Infrastructure home, use the `dbaascli gridHome create` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli gridHome create --version value [--resume [--sessionID value]] [--
waitForCompletion value]
```

Where:

- `--version` specifies the Grid home version
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true|false`

dbaascli grid configureTCPS

To configure TCPS for the existing cluster, use the `dbaascli grid configureTCPS` command.

Prerequisite

Run the command as the `root` user.

Syntax

> **Note:**
>
> By default, TCPS is enabled for databases on Oracle Exadata Database Service on Dedicated Infrastructure systems.

> **Note:**
>
> TCPS is not enabled for databases on Exadata Database Service on Cloud@Customer systems. To enable TCPS for a given database, update the database specific `sqlnet.ora` file with `WALLET_LOCATION = (SOURCE=(METHOD=FILE) (METHOD_DATA=(DIRECTORY=/var/opt/oracle/dbaas_acfs/grid/tcps_wallets)))` on all database nodes and then bounce the database. This will enable TCPS usage for the database. However, enabling TCPS will cause ZDLRA connection to fail. On Exadata Database Service on Cloud@Customer systems, you can enable either ZDLRA or TCPS configuration. Enabling both ZDLRA and TCPS simultaneously will not work.

```
dbaascli grid configureTCPS
[--pkcs12WalletPath]
[--caCertChain]
[--precheckOnly]
[--serverCert]
[--privateKey]
[--certType]
[--privateKeyPasswordProtected]
```

Where:

- `--pkcs12WalletPath` specifies the path of the certificate, which is in `pkcs12` wallet format
- `--caCertChain` concatenated list of certs, containing intermediate CA's and root CA certs
- `--precheckOnly` specifies `yes` to run only the prechecks for this operation. Valid values: `yes` or `no`.
- `--serverCert` specifies the path of PEM certificate to use or rotate for TCPS configuration.
- `--privateKey` specifies the path of the private key file of the certificate.
- `--certType` type of the cert to be added to the Grid Infrastructure wallet. Accepted values are: `SELF_SIGNED_CERT`, `CA_SIGNED_CERT`, or `PKCS12_CERT`. Default: `SELF_SIGNED_CERT`

- `--privateKeyPasswordProtected` specifies if the private key is password protected or not. Valid values: `true` or `false`. Default: `true`.

Example 6-30 dbaascli grid configureTCPS

To configure grid using self-signed certificate:

```
dbaascli grid configureTCPS
```

To configure grid using CA-signed certificate:

```
dbaascli grid configureTCPS --cert_type CA_SIGNED_CERT --server_cert /tmp/
certs/server_cert.pem --ca_cert_chain /tmp/certs/ca.pem --private_key /tmp/
certs/encrypted_private.key --private_key_password_protected false
```

dbaascli grid patch

To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.

Prerequisites

Run the command as the `root` user.

Syntax

```
dbaascli grid patch --targetVersion <value>
[--containerURL <value>]
[--executePrereqs]
[--nodeList <value>]
[--rollback]
[--resume [--sessionID <value>]]
[--continueWithDbDowntime]
        {
            [--createImage [--createImageDir <value>]]
            | [--imageFile <value>]
        }
[--waitForCompletion <value>]
```

Where:

- `--targetVersion` specifies the target version of Oracle Home specified as five numeric segments separated by periods (e.g. 19.12.0.0.0)
- `--containerURL` specifies custom URL for fetching Grid Infrastructure image
- `--executePrereqs` option to run prereqs
- `--nodeList` specifies a comma-delimited list of nodes if patching has to be performed on a subset of nodes
- `--rollback` specifies to roll back patched Oracle home
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID

- `--continueWithDbDowntime` continues patching with database downtime. This option can be used in environments wherein there is only 1 active instance up and the patching operation can be continued even with a downtime.
- `--createImage` creates an image from a copy of the active Grid home, patched to the specified target version
  - `--createImageDir` specifies fully qualified path of the directory where the image is to be created
- `--imageFile` specifies fully qualified path of the image to be used
- `--waitForCompletion` specifies `false` to run the operation in background. Valid values: `true|false`

Example 6-31 dbaascli grid patch

```
dbaascli grid patch --targetVersion 19.12.0.0.0
```

dbaascli grid removeTCPSCert

To remove existing TCPS certificates from Grid Infrastructure wallet, use the `dbaascli grid removeTCPSCert` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli grid removeTCPSCert --subject <value>
 {
   --userCert | --trustedCert | --requestedCert
 }
 [--serialNumber <value>] [--executePrereqs] [--resume [--sessionID <value>]]
[--bounceListeners]
```

Where:

- `--subject` specifies subject of the certificate
- `--userCert` flag to indicate user certificate
- `--trustedCert` flag to indicate trusted certificate
- `--requestedCert` flag to indicate requested certificate
- `--serialNumber` specifies the serial number of the certificate
- `--executePrereqs` runs the prerequisite checks and reports the results
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--bounceListeners` flag to bounce the Grid Infrastructure listener and scan listener

dbaascli grid rotateTCPSCert

To rotate TCPS certificates, use the dbaascli grid rotateTCPSCert command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli grid rotateTCPSCert
[--pkcs12WalletPath]
[--caCertChain]
[--precheckOnly]
[--serverCert]
[--privateKey]
[--certType]
[--privateKeyPasswordProtected]
```

Where:

- `--pkcs12WalletPath` specifies the path of the certificate, which is in `pkcs12` wallet format
- `--caCertChain` concatenated list of certs, containing intermediate CA's and root CA certs
- `--precheckOnly` specifies `yes` to run only the prechecks for this operation. Valid values: `yes` or `no`.
- `--serverCert` specifies the path of PEM certificate to use or rotate for TCPS configuration.
- `--privateKey` specifies the path of the private key file of the certificate.
- `--certType` type of the cert to be added to the Grid Infrastructure wallet. Accepted values are: `SELF_SIGNED_CERT`, `CA_SIGNED_CERT`, or `PKCS12_CERT`. Default: `SELF_SIGNED_CERT`
- `--privateKeyPasswordProtected` specifies if the private key is password protected or not. Valid values: `true` or `false`. Default: `true`.

Example 6-32 dbaascli grid rotateTCPSCert

To rotate cert using self-signed certificate (default option):

```
dbaascli grid rotateTCPSCert
```

To rotate cert using CA-signed certificate:

```
dbaascli grid rotateTCPSCert --cert_type CA_SIGNED_CERT --server_cert /tmp/
certs/server_cert.pem --ca_cert_chain /tmp/certs/ca.pem --private_key /tmp/
certs/encrypted_private.key --privateKeyPasswordProtected true
```

dbaascli grid upgrade

To upgrade Oracle Grid Infrastrucure from one major version to another, use the `dbaascli grid upgrade` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli grid upgrade --version
[--resume]
[--executePrereqs]
[--containerURL]
[--softwareOnly]
[--targetHome]
[--revert]
```

Where:

- `--version` specifies the target version
- `--resume` resumes the previous run
- `--executePrereqs` runs prereqs for Grid Infrastrucure upgrade
- `--containerUrl` specifies the custom URL for fetching Grid Infrastrucure image
- `--softwareOnly` installs only the Grid Infrastructure software
- `--targetHome` specifies the path of existing target Grid home
- `--revert` reverts failed run

Example 6-33 dbaascli grid upgrade

```
daascli grid upgrade --version 19.11.0.0.0 --executePrereqs
DBAAS CLI version MAIN
Executing command grid upgrade --version 19.11.0.0.0 --executePrereqs
```

dbaascli job getStatus

To view the status of a specified job, use the `dbaascli job getStatus` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli job getStatus --jobID
```

Where:

- `--jodID` specifies the job ID

Example 6-34 dbaascli job getStatus

```
dbaascli job getStatus --jobID 13c82031-f202-41b7-9aef-f4a71df0f551
DBAAS CLI version MAIN
Executing command job getStatus --jobID 13c82031-f202-41b7-9aef-f4a71df0f551
{
  "jobId" : "13c82031-f202-41b7-9aef-f4a71df0f551",
  "status" : "Success",
  "message" : "database create job: Success",
  "createTimestamp" : 1628095442431,
  "updatedTime" : 1628095633660,
  "description" : "Service job report for operation database create",
  "appMessages" : {
    "schema" : [ ],
    "errorAction" : "SUCCEED_AND_SHOW"
  },
  "resourceList" : [ ],
  "pct_complete" : "100"
}
```

dbaascli patch db apply

> **Note:**
>
> `dbaascli patch db prereq` and `dbaascli patch db apply` commands have been deprecated in `dbaascli` release 21.2.1.2.0, and replaced with `dbaascli grid patch`, `dbaascli dbhome patch`, and `dbaascli database move` commands.
> For more information, see:
>
> - `dbaascli grid patch`
> - `dbaascli dbhome patch`
> - `dbaascli database move`
> - *Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli*

Related Topics

- dbaascli grid patch
  To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.
- dbaascli dbHome patch
  To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.
- dbaascli database move
  To move the database from one home to another, use the `dbaascli database move` command.
- Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli
  Learn to use the `dbaascli` utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata Cloud Infrastructure system.

dbaascli patch db prereq

> **Note:**
>
> `dbaascli patch db prereq` and `dbaascli patch db apply` commands have been deprecated in `dbaascli` release 21.2.1.2.0, and replaced with `dbaascli grid patch`, `dbaascli dbhome patch`, and `dbaascli database move` commands.
> For more information, see:
>
> - `dbaascli grid patch`
> - `dbaascli dbhome patch`
> - `dbaascli database move`
> - *Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli*

Related Topics

- dbaascli grid patch
  To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.
- dbaascli dbHome patch
  To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.
- dbaascli database move
  To move the database from one home to another, use the `dbaascli database move` command.
- Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli
  Learn to use the `dbaascli` utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata Cloud Infrastructure system.

dbaascli pdb backup

To backup a pluggable database (PDB), query PDB backups, and delete a PDB backup, use the `dbaascli pdb backup` command.

Prerequisite

- Run the command as the `root` user.

Syntax

```
dbaascli pdb backup --pdbName <value> --dbname <value>
        {
            --start
                {
                    [--level1]
                    | [--archival --tag <value>]
                }
            | --delete --backupTag <value>
```

```
            | --status --uuid <value>
            | --getBackupReport --json <value> --tag <value>
            | --list [--json <value>]
        }
```

Where:

```
--pdbName: PDB name.
--dbname: Oracle Database name.
--start | --delete | --status | --getBackupReport | --list
--start: Begins PDB backup.
      [--level1 | --archival]
      [--level1: Creates a Level-1 (incremental) backup.]
      [--archival: Creates an archival full backup.]
          --tag: Specify backup tag.
--delete: Deletes archival backup.
          --backupTag: Specify backup tag to delete.
--status
          --uuid <value>
--getBackupReport: Returns backup report.
         --json: Specify the file name for JSON output.
         --tag: Specify backup tag.
--list: Returns PDB backup information.
         [--json: Specify the file name for JSON output.]
```

Example 6-35 Examples

- To take level1 backup for a PDB *pdb1* in a CDB *myTestDb*:

  ```
  dbaascli pdb backup --dbname myTestDb --pdbName pdb1 --start --level1
  ```

- To query the status of PDB backup request submitted with `uuid` *eef16b26361411ecb13800163e8e4fac*:

  ```
  dbaascli pdb backup --dbname myTestDb --pdbName pdb1 --status --uuid
  eef16b26361411ecb13800163e8e4fac
  ```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

dbaascli pdb bounce

To bounce a pluggable database (PDB), use the `dbaascli pdb bounce` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb bounce --dbname --pdbName | --pdbUID
[-openMode]
```

Where:

- `--dbname` specifies the name of the container database that hosts the PDB
- `--pdbName` specifies the name of the PDB
- `--pdbUID` specifies the identifier of the PDB
- `--openMode` specifies the target `OPEN MODE` of PDB

Example 6-36 dbaascli pdb bounce

```
dbaascli pdb bounce --dbname cdb_name --pdbName pdb name associated with the
CDB


dbaascli pdb bounce --dbname cdb_name --pdbUID con_uid of that pdb
```

Optional:

- `--openMode READ_WRITE`
- `--openMode READ_ONLY`

dbaascli pdb close

To close a pluggable database (PDB), use the `dbaascli pdb close` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb close --dbname --pdbName | --pdbUID
```

Where:

- `--dbname` specifies the name of the container database that hosts the PDB.
- `--pdbname` specifies the name of the PDB that you want to close.
- `--pdbUID` specifies the identifier of the PDB

Upon successful completion of running this command, the PDB is closed on all of the container database instances.

Example 6-37 dbaascli pdb close

```
dbaascli pdb close --dbname cdb name --pdbName pdb name associated with the
CDB


dbaascli pdb close --dbname cdb name --pdbUID con_uid of that pdb
```

dbaascli pdb getConnectString

To display Oracle Net connect string information for a pluggable database (PDB) run the `dbaascli pdb getConnectString` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb getConnectString --dbname --pdbName | --pdbUID
```

Where:

- `--dbname` specifies the name of the container database that hosts the PDB
- `--pdbname` specifies the name of the PDB for which you want to display connect string information
- `--pdbUID` specifies the identifier of the PDB

Example 6-38 dbaascli pdb getConnectString

```
dbaascli pdb getConnectString --dbname dbname --pdbName pdbName
```

dbaascli pdb create

To create a new pluggable database (PDB), use the `dbaascli pdb create` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb create --pdbName <value> --dbName <value>
[--maxCPU <value>]
[--maxSize <value>]
[--pdbAdminUserName <value>]
[--lockPDBAdminAccount <value>]
[--resume [--sessionID <value>]]
[--executePrereqs <value>]
[--waitForCompletion <value>]
[--blobLocation |--standbyBlobFromPrimary <value>]
```

Where:

- `--pdbName` specifies the name of the new PDB that you want to create
- `--dbName` specifies the name of the container database that hosts the new PDB
- `--maxCPU` optionally specifies the maximum number of CPUs that are available to the PDB. Setting this option is effectively the same as setting the `CPU_COUNT` parameter in the PDB

- `--maxSize` optionally specifies the maximum total size of data files and temporary files for tablespaces belonging to the PDB. Setting this option is effectively the same as setting the `MAXSIZE PDB` storage clause in the `CREATE PLUGGABLE DATABASE` SQL command. You can impose a limit by specifying an integer followed by a size unit (`K`, `M`, `G`, or `T`), or you can specify `UNLIMITED` to explicitly enforce no limit
- `--pdbAdminUserName` specifies the new PDB admin user name
- `--lockPDBAdminAccount` specifies `true` or `false` to lock the PDB admin user account. Default value is `true`.
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes` or `no`
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`
- `--blobLocation` custom directory location where the standby blob file will be generated in a DG environment.
- `--standbyBlobFromPrimary` specifies the location of the standby blob file, which is prepared from the primary database. This is required only for standby database PDB operations.

> **Note:**
>
> the parameters`blobLocation` and `standbyBlobFromPrimary` are mutually exclusive.

During the PDB creation process, you are prompted to specify the administration password for the new PDB.

Example 6-39 dbaascli pdb create

To create a PDB from seed in a standard database in a non-Data Guard environment:

```
dbaascli pdb create --dbName db721 --pdbName new_pdb1 --maxsize 5G --maxcpu 2
```

To create PDB in Data Guard environment:

```
dbaascli pdb create --dbName db721 --pdbName new_pdb1
```

```
dbaascli pdb create --dbName db721 --pdbName new_pdb1 --
standbyBlobFromPrimary /tmp/send_db721.tar
```

dbaascli pdb delete

To delete a pluggable database (PDB) run the `dbaascli pdb delete` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb delete --dbName value
{ --pdbName value | --pdbUID value }
[--executePrereqs value]
[--waitForCompletion value]
[--resume [--sessionID value]]
[--allStandbyPrepared]
[--cleanupRelocatedPDB]
```

Where:

- `--dbName` specifies the name of the container database that hosts the PDB
- `--pdbName` specifies the name of the PDB that you want to delete
- `--pdbUID` specifies the UID of the PDB that you want to delete
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes` or `no`
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`
- `--resume` specifies to resume the previous execution
  - `--sessionID` specifies to resume a specific session ID
- `--allStandbyPrepared` specifies to confirm that the operation has been successfully run on all the standby databases
- `--cleanupRelocatedPDB` - option to cleanup source database after a PDB has been relocated.

Example: dbaascli pdb delete
To delete a PDB from a standard database in a non-Data Guard environment or from Standby database in Data Guard environment.

```
dbaascli pdb delete --dbName db721 --pdbName pdb1
```

To create PDB from Primary database in Data Guard environment:

```
dbaascli pdb create --dbName db721 --pdbName pdb1 --allStandbyPrepared
```

dbaascli pdb getDetails

To view details of a pluggable database (PDB), use the `dbaascli pdb getDetails` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb getDetails --dbname --pdbName | --pdbUID
```

Where:

- `--dbname` specifies the name of the container database that hosts the PDB
- `--pdbname` specifies the name of the PDB that you want to delete
- `--pdbUID` specifies the identifier of the PDB

Example 6-40 dbaascli pdb getDetails

```
dbaascli pdb getDetails--dbname cdb name --pdbName pdb name associated with
the CDB


dbaascli pdb getDetails--dbname cdb name --pdbUID con_uid of that pdb
```

dbaascli pdb list

To view the list of pluggable databases (PDB) in a container database, use the `dbaascli pdb list` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb list --dbname
```

Where:

- `--dbname` specifies the name of the container database that hosts the PDB

Example 6-41 dbaascli pdb list

```
dbaascli pdb list --dbname cdb name
```

dbaascli pdb localClone

To create a new pluggable database (PDB) as a clone of an existing PDB in the same container database (CDB), use the `dbaascli pdb localClone` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb localClone --pdbName <value> --dbName <value>
[--targetPDBName <value>]
[--powerLimit <value>]
[--maxCPU <value]
[--maxSize <value>]
[--resume [--sessionID <value>]]
[--executePrereqs]
[--waitForCompletion <value>]
{[--blobLocation <value>]|[--standbyBlobFromPrimary <value>]}
[--excludeUserTablespaces <value>]
```

```
[--excludePDBData <value>]
[--pdbAdminUserName <value>]
[--lockPDBAdminAccount <value>]
[--sourcePDBServiceConvertList <value>]
```

Where:

- `--pdbName` specifies the name of the new PDB that you want to clone
- `--dbName` specifies the name of the database
- `--targetPDBName` specifies the name for the target PDB (new cloned PDB)
- `--powerLimit` specifies the degree of parallelism to be used for the clone operation. Valid value is between 1 and 128
- `--maxCPU` specifies the maximum number of CPUs to be allocated for the PDB
- `--maxSize` specifies the maximum storage size in GB for the new PDB
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes` or `no`
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`
- `--blobLocation` custom directory location where the standby blob file will be generated in a DG environment.
- `--standbyBlobFromPrimary` specifies the location of the standby blob file which is prepared from the primary database. This is required only for standby database PDB operations.

> **Note:**
>
> The parameters `--blobLocation` and `--standbyBlobFromPrimary` are mutually exclusive.

- `--excludeUserTablespaces` option to skip user table spaces, example t1,t2,t3.
- `--excludePDBData` specify true/yes to skip user data from source pdb.
- `--pdbAdminUserName` specify new PDB admin user name.
- `--lockPDBAdminAccount` specify true or false to lock the PDB admin user account. Default value is true.
- `--sourcePDBServiceConvertList` specify comma separated list of source to target service names which need to be converted. Syntax is source_srv1:new_srv1,source_srv2:new_srv2.

The newly cloned PDB inherits administration passwords from the source PDB.

Example 6-42 dbaascli pdb localClone

```
dbaascli pdb localClone --dbName db35 --pdbName PDB35 --targetPDBName
local_clone1 --maxCPU 2 --maxSize 15
```

dbaascli pdb open

To open a pluggable database (PDB), use the `dbaascli pdb open` command.

Run the command as the `root` or `oracle` user.

Syntax

```
dbaascli pdb open
 {
   --pdbName <value> | --pdbUID <value>
 }
--dbname <value> [--openMode <value>] [--startServices <value>] [--
waitForCompletion <value>] [--setPDBRefreshModeNone [--skipPDBRefresh] [--
pdbAdminUserName <value>]]
```

Where:

- `--pdbName` specifies the name of the PDB that you want to open
- `--pdbUID` specifies the identifier of the PDB
- `--dbname` specifies the name of the container database that hosts the PDB.
- `--openMode` specifies the target OPEN MODE of PDB
- `--startServices`: specifies to start all or list all services corresponding to a PDB. Accepted values are `all` or a comma-delimited list of PDB services.
- `--waitForCompletion`: specify `false` to run the operation in the background. Valid values: `true|false`
- `--setPDBRefreshModeNone`: specifies to convert a refreshable PDB to non-refreshable PDB
  - `--skipPDBRefresh`: specifies to skip refreshable PDB refresh
  - `--pdbAdminUserName`: specifies new PDB admin user name

Upon successful completion, the PDB is opened on all of the container database instances.

Example 6-43 dbaascli pdb open

```
dbaascli pdb open --dbname cdb name --pdbName pdb name associated with the CDB


dbaascli pdb open --dbname cdb name --pdbUID con_uid of that pdb
```

Optional: `--openMode READ_WRITE/READ_ONLY`

dbaascli pdb recover

To recover a pluggable database (PDB), use the `dbaascli pdb recover` command.

Prerequisite

- Run the command as the `root` user.
- Database must be configured with backup storage destination details where backups are stored.

Syntax

```
dbaascli pdb recover --pdbName <value> --dbname <value>
        {
            --start
                {
                    --untilTime <value>
                    | --untilSCN <value>
                    | --latest
                    | --tag <value>
                }
            | --status --uuid <value>
        }
```

Where:

```
--pdbName: PDB name.
--dbname: Oracle Database name.
--start | --status
--start
      --untilTime | --untilSCN | --latest | --tag
      --untilTime: Recovers PDB until time. Input format: DD-MON-YYYY HH24:MI:SS.
      --untilSCN: Recovers PDB until SCN.
      --latest: Recovers PDB to last known state.
      --tag: Recovers PDB to archival tag.
--status
      --uuid <value>
```

Example 6-44 Examples

- To recover a PDB *pdb1* in a CDB *myTestDb* to latest:

  ```
  dbaascli pdb recover --dbname myTestDb --pdbName pdb1 --start --latest
  ```

- To query the status of PDB recovery request submitted with `uuid` *81a17352362011ecbc3000163e8e4fac*:

  ```
  dbaascli pdb recover --dbname myTestDb --pdbName pdb1 --status --uuid
  81a17352362011ecbc3000163e8e4fac
  ```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

dbaascli pdb refresh

To refresh a specified pluggable database (PDB), use the `dbaascli pdb refresh` command.

Run the command as the `root` or `oracle` user.

Syntax

```
dbaascli pdb refresh --dbname <value>
   {
      --pdbName <value> | --pdbUID <value>
    }
    [--waitForCompletion <value>]
```

Where:

- `--dbname`: specifies the name of the Oracle Database
- `--pdbName`: specifies the name of the pluggable database
- `--pdbUID`: specifies the identifier of the pluggable database
- `--waitForCompletion`: specify `false` to run the operation in the background. Valid values: `true|false`

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

dbaascli pdb relocate

To relocate the specified PDB from the remote database into local database, use the `dbaascli pdb relocate` command.

Prerequisite

Run the command as the `oracle` user. When prompted, you must supply the SYS user password for the source database.

Syntax

```
dbaascli pdb relocate --pdbName <value> --dbName <value> --
sourceDBConnectionString <value>
[--targetPDBName <value>]
[--powerLimit <value>]
[--maxCpu <value>]
[--maxSize <value>]
[--resume [--sessionID <value>]]
[--executePrereqs <value>]
[--sourcePDBServices <value>]
[--sourcePDBReadOnlyServices <value>]
[--waitForCompletion <value>]
{
    [--blobLocation <value>] | [--standbyBlobFromPrimary <value>]
}
[--upgradePDB <value>]
[--updateDBBlockCacheSize]
{
    [skipOpenPDB] | [--completePDBRelocate]
}
```

Where:

- `--pdbName` specifies the source PDB name to relocate
- `--dbName` specifies the target database name
- `--sourceDBConnectionString` specifies the source database connection string in the format *<scan_name>:<scan_port>/<database_service_name>*
- `--targetPDBName` specifies a name for the target PDB (new relocated PDB)
- `--powerLimit` specifies the degree of parallelism to be used for the relocate operation
- `--maxCpu` specifies the maximum number of CPUs to be allocated for the PDB
- `--maxSize` specifies the maximum storage size in GB for the new PDB
- `--resume` specifies to resume the previous execution
  - `--sessionID` specifies to resume a specific session ID
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes|no`
- `--sourcePDBServices` specifies a list of comma-delimited source PDB services
- `--sourcePDBReadOnlyServices` specifies a comma-delimited list of source PDB read-only services
- `--waitForCompletion` specifies false to run the operation in the background. Valid values: `true|false`
- `--blobLocation` custom directory location where the standby blob file will be generated in a DG environment.
- `--standbyBlobFromPrimary` specify the location of the standby blob file which is prepared from the primary database. This is required only for standby operations.

> **Note:**
>
> The parameters `--blobLocation` and mutually exclusive.

- `--upgradePDB` specify `true` to upgrade the PDB as part of this operation. Valid values : `true | false`.
- `--updateDBBlockCachesize` option to enable application to set db block cache size initialization parameters in order to support data copy with different block size.
- `--skipOpenPDB` - to indicate that the PDB should not be opened at the end of the current operation.
- `--completePDBRelocate` - complete the PDB relocation if done as a two-step operation.

Example 6-45 dbaascli pdb relocate

```
dbaascli pdb relocate --sourceDBConnectionString test-
scan.dbaastoolslrgsu.dbaastoolslrgvc.oraclevcn.com:1521/
source_cdb_service_name --pdbName source_pdb --dbName target_db
```

dbaascli pdb remoteClone

To create a new pluggable database (PDB) as a clone of an existing PDB in another container database (CDB), use the `dbaascli pdb remoteClone` command.

Run the command as the `root` or `oracle` user.

Syntax

```
dbaascli pdb remoteClone --pdbName <value> --dbName <value> --
sourceDBConnectionString <value> [--targetPDBName <value>] [--powerLimit
<value>] [--maxCPU <value>] [--maxSize <value>] [--resume [--sessionID
<value>]] [--executePrereqs] [--waitForCompletion <value>] [--
sourcePDBExportedTDEKeyFile <value>]
        {
            [--blobLocation <value>]
            | [--standbyBlobFromPrimary <value>]
        }
[--excludeUserTablespaces <value>]
[--excludePDBData <value>]
[--pdbAdminUserName <value>]
[--lockPDBAdminAccount <value>]
[--sourcePDBServiceConvertList <value>]
[--refreshablePDB --refreshMode <value> [--refreshIntervalInMinutes <value>]
--dblinkUsername <value> [--honorCaseSensitiveUserName]]
[--updateDBBlockCacheSize]
```

Where:

- `--pdbName` specifies the name of the source PDB that you want to clone
- `--dbname` specifies the name (`DB_NAME`) of the CDB that hosts the newly cloned PDB
- `--sourceDBConnectionString` specifies the source database connection string in the format *`scan_name`*`:`*`scan_port`*`/`*`database_service_name`*
- `--targetPDBName` specifies the name for the target PDB (new cloned PDB)
- `--powerLimit` specifies the degree of parallelism to be used for the clone operation. Valid value is between 1 and 128
- `--maxCPU` specifies the maximum number of CPUs to be allocated for the PDB
- `--maxSize` specifies the maximum storage size in GB for the new PDB
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes` or `no`
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`
- `--sourcePDBExportedTDEKeyFile` specifies the source PDB exported key file. This variable is applicable to only 12.1 database.
- `--blobLocation` specifies the custom path where the standby blob file will be generated in a Data Guard environment

- `--standbyBlobFromPrimary` specify the location of the standby blob file, which is prepared from the primary database. This is required only for standby database PDB operations

> **Note:**
>
> The parameters `--blobLocation` and `--standbyBlobFromPrimary` are mutually exclusive.

- `--excludeUserTablespaces` option to skip user table spaces, example *t1*,*t2*,*t3*.
- `--excludePDBData` specify `true`/`yes` to skip user data from source PDB.
- `--pdbAdminUserName` specifies new PDB admin user name
- `--lockPDBAdminAccount` specify `true` or `false` to lock the PDB admin user account. Default value is `true`.
- `--sourcePDBServiceConvertList` specify a comma-delimited list of source to target service names, which need to be converted. Syntax is `source_srv1:new_srv1, source_srv2:new_srv2`.
- `--refreshablePDB` specifies to create refreshable PDB
  - `--refreshMode` specifies refresh mode for refreshable PDB. Valid values: `AUTO`|`MANUAL`
    - `--refreshIntervalInMinutes` specifies refresh interval for `refreshablePDB` in minutes
  - `--dblinkUsername` specifies common user of a remote database used for database link to connect to the remote database
    - `--honorCaseSensitiveUserName` indicates specified username is case sensitive
- `--updateDBBlockCacheSize`: specifies to enable application to set `db block cache size` initialization parameters to support data copy with a different block size

When promoted, you must supply the SYS user password for the source PDB. The newly cloned PDB inherits administration passwords from the source PDB. The cloned PDB is named using the following format: `dbname_sourcepdbname`. This command is supported only for databases that are not in a Data Guard configuration and use Oracle Database version 12.2.0.1, or later.

Example 6-46 dbaascli pdb remoteClone

```
dbaascli pdb remoteClone --sourceDBConnectionString test-
can.dbaastoolslrgsu.dbaastoolslrgvc.oraclevcn.com:1521 --pdbName source_pdb1
--dbName db9944 --targetPDBName new_pdb1 --maxsize 5 --maxcpu 2


dbaascli pdb remoteClone --sourceDBConnectionString
orcla.dbaastoolslrgsu.dbaastoolslrgvc.oraclevcn.com --pdbName source_pdb1 --
dbName db9944 --targetPDBName new_pdb1 --maxsize 5 --maxcpu 2
```

dbaascli system getDBHomes

To view information about all the Oracle homes, use the `dbaascli system getDBHomes` command.

Prerequisite

Run the command as the `root` or `oracle` user.

Syntax

```
dbaascli system getDBHomes
```

Example 6-47 dbaascli system getDBHomes

```
dbaascli system getDBHomes
```

dbaascli system getGridHomes

To list the details of all Grid homes, use the `dbaascli system getGridHomes` command.

Prerequisite

Run the command as the `root` or `oracle` user.

Syntax

```
dbaascli system getGridHomes
```

dbaascli tde changePassword

To change TDE keystore password as well as DB wallet password for the alias `tde_ks_passwd`, use the `dbaascli tde changePassword` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde changePassword [--dbname <value>]
  {            [--prepareStandbyBlob <value> [--blobLocation <value>]]
               | [--standbyBlobFromPrimary <value>]
  }
  [--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the name of the database
- `--prepareStandbyBlob` - specify true to generate a blob file containing the artifacts needed to perform the operation in a DG environment.

- `--blobLocation` - custom path where the standby blob file will be generated in a DG environment.
- `--standbyBlobFromPrimary` - specify the location of the standby blob file which is prepared from the primary database. This is required only for standby operations.
- `--resume` - to resume the previous execution
- `--sessionID` - to resume a specific session id.

```
dbaascli tde changepassword --dbname
      <dbname>
```

1. Change the TDE password in primary database.

   ```
   dbaascli tde changepassword --dbname
         <dbname> --prepareStandbyBlob true --blobLocation
         <Location where blob file has to be generated>
   ```

2. Copy the created standby blob to standby database environment.
3. Change the TDE password in standby database

   ```
   dbaascli tde changepassword --dbname
        <dbname> --standbyBlobFromPrimary <Location of blob generated from
       primary>
   ```

dbaascli tde addSecondaryHsmKey

To add a secondary HSM (KMS) key to the existing HSM (KMS) configuration, use the `dbaascli tde addSecondaryHsmKey` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde addSecondaryHsmKey --dbname <value> --secondaryKmsKeyOCID <value>
[--executePrereqs]
```

Where:

- `--secondaryKmsKeyOCID` specifies the secondary KMS key to add to the existing HSM (KMS) configuration
- `--dbname` specifies the name of the database
- `--executePrereqs` sexecute the prerequisites checks and report the results.

Example 6-48 dbaascli tde addSecondaryHsmKey

```
dbaascli tde addSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID
ocid1.key.oc1.eu-
```

```
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5ux
ainvvxza


dbaascli tde addSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5ux
ainvvxza --precheckOnly yes
```

dbaascli tde enableWalletRoot

To enable `wallet_root` spfile parameter for the existing database, use the `dbaascli tde enableWalletRoot` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde enableWalletRoot
[--dbRestart]
[--dbname]
[--precheckOnly]
```

Where:

- `--dbrestart` specifies the database restart option. Valid values are: `rolling` or `full`. Default value: `rolling`
  If you do not pass the `dbrestart` argument, then the database restarts in a `rolling` manner.
- `--dbname` specifies the name of the Oracle Database.
- `--precheckOnly` runs only the precheck for this operation. Valid values are: `yes` or `no`

Example 6-49 dbaascli tde enableWalletRoot

```
dbaascli tde enableWalletRoot --dbname db name --dbrestart rolling|full


dbaascli tde enableWalletRoot --dbname orcl


dbaascli tde enableWalletRoot --dbname orcl--dbrestart full
```

dbaascli tde encryptTablespacesInPDB

To encrypt all the tablespaces in the specified PDB, use the `dbaascli tde encryptTablespacesInPDB` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde encryptTablespacesInPDB --pdbName
[--dbname]
[--precheckOnly]
[--useSysdbaCredential]
```

Where:

- `--pdbName` specifies the name of the PDB to encrypt all the tablespaces.
- `--dbname` specifies the name of the Oracle Database.
- `--precheckOnly` runs only the precheck for this operation. Valid values: `yes` or `no`
- `--useSysdbaCredential` uses SYSDBA credentials for this operation if passed value is `true`. Valid values: `true` or `false`

Example 6-50 dbaascli tde encryptTablespacesInPDB

```
dbaascli tde encryptTablespacesInPDB --dbname dbname --pdbName pdb
```

```
dbaascli tde encryptTablespacesInPDB --dbname dbname --pdbName pdb --
precheckOnly yes
```

```
dbaascli tde encryptTablespacesInPDB --dbname dbname --pdbName pdb --
useSysdbaCredential true
```

dbaascli tde fileToHsm

To convert FILE based TDE to HSM (KMS/OKV) based TDE, use the `dbaascli tde fileToHsm` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde fileToHsm --kmsKeyOCID <value> --dbname <value>
[--skipPatchCheck <value>]
[--executePrereqs ]
[--primarySuc <value>]
{
    [--resume [--sessionID <value>]] | [--revert [--sessionID <value>]]
}
[--waitForCompletion <value>]
```

Where:

- `--kmsKeyOCID` specifies the KMS key OCID to use for TDE. This is applicable only if KMS is selected for TDE
- `--dbname` specifies the name of the database

- `--skipPatchCheck` skips validation check for required patches if the value passed for this argument is `true`. Valid values: `true` or `false`
- `--executePrereqs` sexecute the prerequisites checks and report the results.
- `--primarySuc` specify this property in the standby database of the Data Guard environment once the command is successfully run on the primary database
- `--resume` specifies to resume the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--revert` specifies to rollback the previous run
  - `--sessionID` specifies to rollback a specific session ID
- `--waitForCompletion` specify false to run the operation in background. Valid values : true| false.

Example 6-51 dbaascli tde fileToHsm --kmsKeyOCID

```
dbaascli tde fileToHSM --dbname dbname --kmsKeyOCID ocid1.key.oc1.eu-
frankfurt-.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5uxa
invvxza
```

```
dbaascli tde fileToHSM --dbname dbname --kmsKeyOCID ocid1.key.oc1.eu-
frankfurt-.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5uxa
invvxza --executePrereqs
```

```
dbaascli tde fileToHSM --dbname dbname --kmsKeyOCID ocid1.key.oc1.eu-
frankfurt-.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5uxa
invvxza --resume
```

dbaascli tde getHsmKeys

To get TDE active key details, use the `dbaascli tde getHsmKeys` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde getHsmKeys
[--dbname]
[--infoFile]
```

Where:

- `--dbname` specifies the name of the database
- `--infoFile` specifies the file path where the list of OCIDs will be saved. The output is in JSON format

Example 6-52 dbaascli tde getHsmKeys

```
dbaascli tde getHsmkeys --dbname dbname


dbaascli tde getHsmkeys --dbname dbname --infoFile infoFilePath
```

dbaascli tde getMkidForKeyVersionOCID

To get Master Key ID associated with the KMS key version OCID, use the `dbaascli tde getMkidForKeyVersionOCID` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde getMkidForKeyVersionOCID --kmsKeyVersionOCID <value>
[--dbname <value>]
[--waitForCompletion <value>]
```

Where:

- `--kmsKeyVersionOCID` specifies the KMS key version OCID to set
- `--dbname` specifies the name of the database
- `--waitForCompletion` specify `false` to run the operation in background. Valid values : `true|false`.

Example 6-53 dbaascli tde getMkidForKeyVersionOCID

```
dbaascli tde getMkidForKeyVersionOCID --dbname dbname --kmsKeyVersionOCID
ocid1.keyversion.oc1.eu-
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcwvylkd2l
x56lu2s6iwnxwgigu23nha
```

dbaascli tde getPrimaryHsmKey

To get primary HSM (KMS) key from the existing HSM (KMS) configuration, use the `dbaascli tde getPrimaryHsmKey` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde getPrimaryHsmKey
[--dbname]
```

Where:

- `--dbname` specifies the name of the database

Example 6-54 dbaascli tde getPrimaryHsmKey

```
dbaascli tde getPrimaryHsmKey --dbname dbname
```

dbaascli tde hsmToFile

To convert HSM (KMS/OKV) based TDE to FILE based TDE, use the `dbaascli tde hsmToFile` command.

Run the command as the `root` user.

Syntax

```
dbaascli tde hsmToFile
[--dbname <value>]
{
    [--prepareStandbyBlob <value> [--blobLocation <value>]
   | [--standbyBlobFromPrimary <value>]
}
]
[--skipPatchCheck <value>]
[--executePrereqs ]
[--primarySuc <value>]
{
     [--resume [--sessionID <value>]] |
     [--revert [--sessionID <value>]]
}
[--waitForCompletion <value>]
```

Where:

- `--dbname` specifies the name of the database
- `--prepareStandbyBlob` specify `true` to generate a blob file containing the artifacts needed to perform the operation in a DG environment.
- `--blobLocation` custom directory location where the standby blob file will be generated in a DG environment.
- `--standbyBlobFromPrimary` specify the location of the standby blob file which is prepared from the primary database. This is required only for standby operations.]
- `--skipPatchCheck` skips validation check for required patches if the value passed for this argument is `true`. Valid values: `true` or `false`
- `--executePrereqs` execute the prerequisites checks and report the results.
- `--primarySuc` specify this property in the standby database of the Data Guard environment once the command is successfully run on the primary database
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--revert` specifies to roll back the previous run
  - `--sessionID` specifies to rollback a specific session ID
- `--waitForCompletion` specifies `false` to run the operation in background. Valid values: `true|false`

Example 6-55 dbaascli tde hsmToFile

```
dbaascli tde hsmToFile --dbname dbname
```

```
dbaascli tde hsmToFile --dbname dbname --executePrereqs
```

```
dbaascli tde hsmToFile --dbname dbname --resume
```

dbaascli tde listKeys

To list TDE master keys, use the `dbaascli tde listKeys` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde listKeys
[--dbname <value>]
[--infoFilePath <value>]
```

Where:

- `--dbname` specifies the name of the database
- `--infoFilePath` specify the absolute path of the file where the results will be saved.

Example 6-56 dbaascli tde listKeys

```
dbaascli tde listKeys --dbname dbname
```

```
dbaascli tde listKeys --dbname dbname --infoFilePath infoFilePath
```

dbaascli tde removeSecondaryHsmKey

To remove secondary HSM (KMS) key from the existing HSM (KMS) configuration, use the `dbaascli tde removeSecondaryHsmKey` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde removeSecondaryHsmKey --dbname <value>
[--confirmDeletion]
[--secondaryKmsKeyOCID]
[--executePrereqs]
```

Where:

- `--dbname` specifies the name of the database
- `--confirmDeletion` if not specified the user will be prompted while deleting all existing HSM(KMS) keys.
- `--secondaryKmsKeyOCID` secondary KMS key to be removed from existing HSM(KMS) configuration. If not specified all secondary KMS keys will be removed.
- `--executePrereqs` execute the prerequisites checks and report the results.

Example 6-57 dbaascli tde removeSecondaryHsmKey

```
dbaascli tde removeSecondaryHsmKey --dbname dbname


dbaascli tde removeSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5ux
ainvvxza


dbaascli tde removeSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5ux
ainvvxza --executePrereqs
```

dbaascli tde rotateMasterKey

Rotate the master key for database encryption.

Prerequisites:

Run the command as the `root` user.

Syntax

(Optional) <Enter syntax information here.>

```
dbaascli tde rotateMasterKey --dbname <value>
[--rotateMasterKeyOnAllPDBs]
[--pdbName <value>]
[--executePrereqs]
[--resume [--sessionID <value>]]
{
     [--prepareStandbyBlob <value> [--blobLocation <value>]]
     | [--standbyBlobFromPrimary <value>]
 }
```

Where:

- `--dbname` - Oracle database name.
- `--rotateMasterKeyOnAllPDBs` - specify true to rotate master key of all PDBs in CDB. Valid values: `true|false`
- `--pdbName` - specify PDB name.
- `--executePrereqs` - execute the prerequisites checks and report the results.

- `--resume` - to resume the previous execution
- `--sessionID` - to resume a specific session id.
- `--prepareStandbyBlob | --standbyBlobFromPrimary]`
- `--prepareStandbyBlob` - specify true to generate a blob file containing the artifacts needed to perform the operation in a DG environment.
- `--blobLocation` - custom directory location where the standby blob file will be generated in a DG environment.
- `--standbyBlobFromPrimary` - specify the location of the standby blob file which is prepared from the primary database. This is required only for standby operations

dbaascli tde setKeyVersion

To set the version of the primary key to be used in DB/CDB or PDB, use the `dbaascli tde setKeyVersion` command.

Run the command as the `root` user.

Syntax

```
dbaascli tde setKeyVersion --kmsKeyVersionOCID <value> --dbname <value>
[--pdbName <value>]
[--masterKeyID <value>]
[--standbySuc]
[--executePrereqs]
[--waitForCompletion <value>]
```

Where:

- `--kmsKeyVersionOCID` specifies the KMS key version OCID to set.
- `--dbname` specifies the name of the database.
- `--pdbName` name of the PDB to use the key version OCID.
- `--masterKeyID` specifies the master key ID of the given key version OCID. This is applicable to the Data Guard environment.
- `--standbySuc` specify this property in the primary database of the Data Guard environment once the command is successfully run on the standby database
- `--executePrereqs` execute the prerequisites checks and report the results.
- `--waitForCompletion` specify `false` to run the operation in background. Valid values: `true|false`

Example 6-58 dbaascli tde setKeyVersion

```
dbaascli tde setKeyVersion --dbname dbname --kmsKeyVersionOCID
ocid1.keyversion.oc1.eu-
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcwvylkd2l
x56lu2s6iwnxwgigu23nha


dbaascli tde setKeyVersion --dbname dbname --kmsKeyVersionOCID
ocid1.keyversion.oc1.eu-
```

```
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcwvylkd2l
x56lu2s6iwnxwgigu23nha --executePrereqs
```

```
dbaascli tde setKeyVersion --dbname dbname --pdbName pdb --kmsKeyVersionOCID
ocid1.keyversion.oc1.eu-
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcwvylkd2l
x56lu2s6iwnxwgigu23nha
```

dbaascli tde setPrimaryHsmKey

To change the primary HSM (KMS) key for the existing HSM (KMS) configuration, use the `dbaascli tde setPrimaryHsmKey` command.

Run the command as the `root` user.

Syntax

```
dbaascli tde setPrimaryHsmKey --primaryKmsKeyOCID <value> --dbname <value>
[--allStandbyPrepared]
[--bounceDatabase]
[--executePrereqs]
[--resume [--sessionID <value>]]
```

Where:

- `--primaryKmsKeyOCID` specifies the primary KMS key to set
- `--dbname` specifies the name of the database
- `--allStandbyPrepared` specify to confirm that the operation has been successfully run on all the standby databases.
- `--bounceDatabase` specify this flag to do rolling database bounce for this operation
- `--executePrereqs` execute the prerequisites checks and report the results.
- `--resume` to resume the previous execution
- `--sessionID` to resume a specific session id.

Example 6-59 dbaascli tde setPrimaryHsmKey

```
dbaascli tde setPrimaryHsmKey --dbname dbname --primaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5ux
ainvvxza
```

```
dbaascli tde setPrimaryHsmKey --dbname dbname --primaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5ux
ainvvxza --executePrereqs
```

dbaascli tde status

To display information about the keystore for the specified database, use the `dbaascli tde status` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli tde status --dbname dbname
```

Where:

- `--dbname` specifies the name of the database that you want to check.

Output from the command includes the type of keystore, and the status of the keystore.

Example 6-60 dbaascli tde status

```
dbaascli tde status --dbname dbname
```

Monitoring and Managing Exadata Storage Servers with ExaCLI

The ExaCLI command line utility allows you to perform monitoring and management functions on Exadata storage servers in an Exadata Cloud Infrastructure instance.

- About the ExaCLI Command
  The ExaCLI command provides a subset of the commands found in the on-premises Exadata command line utility.
- Exadata Storage Server Username and Password
  You need a username and password to connect to the Exadata Storage Server.
- ExaCLI Command Syntax
  For Exadata Storage Server targets, construct your commands using the syntax that follows.
- Connecting to a Storage Server with ExaCLI
  To use ExaCLI on storage servers, you will need to know your target storage server's IP address.
- ExaCLI Command Reference
  You can execute various ExaCLI commands to monitor and manage Exadata Storage Servers associated with your Oracle Cloud Infrastructure Exadata DB system. ExaCLI allows you to get up-to-date, real-time information about your Exadata Cloud Service.

About the ExaCLI Command

The ExaCLI command provides a subset of the commands found in the on-premises Exadata command line utility.

ExaCLI offers a subset of the commands found in the on-premises Exadata command line utility *CellCLI utility*. The utility runs on the database virtual machines in the Exadata Cloud Service.

See the *ExaCLI Command* list in this topic to learn what commands are available.

Related Topics

- Using the CellCLI Utility
- ExaCLI Command Syntax
  For Exadata Storage Server targets, construct your commands using the syntax that follows.

Exadata Storage Server Username and Password

You need a username and password to connect to the Exadata Storage Server.

On Exadata Cloud Infrastructure, the preconfigured user for Exadata Storage Server is `cloud_user_clustername`, where *`clustername`* is the name of the virtual machine (VM) cluster that is being used.

You can determine the name of the VM cluster by running the following `crsctl` command as the `grid` user on any cluster node:

```
crsctl get cluster name
```

IThis command returns `CRS-6724: Current cluster name is <cluster_name>`

The password for `cloud_user_clustername` is initially set to a random value, which you can view by running the following command as the `root` user on any cluster node:

```
/opt/exacloud/get_cs_data.py
```

This returns a password <pwd>

Then test with ExaCLI as root:

```
exacli -l cloud_user_<clusternmae> -c 192.168.136.14
Password: **************************
exacli cloud_user_<cluster_name>@192.168.136.14>
```

ExaCLI Command Syntax

For Exadata Storage Server targets, construct your commands using the syntax that follows.

Note that the syntax example assumes you are the opc user on a compute node.

```
exacli -c [username@]remotehost[:port] [-l username] [--xml] [--cookie-jar
filename] [-e {command | 'command; command' | @batchfile}]
```

NOT_SUPPORTED

This example shows the user on an Exadata compute node issuing the command to log in to ExaCLI start an interactive ExaCLI session on a storage server:

```
[opc@exacs-node1 ~]$ exacli -l cloud_user_clustername -c 192.168.136.7
```

See *Connecting to a Storage Server with ExaCLI* for information on determining your storage server's IP address.

Once logged in, run additional commands as follows:

```
exacli cloud_user_clustername@192.168.136.7> LIST DATABASE
ASM
HRCDB
```

NOT_SUPPORTED

Example 2
This example shows a single command issued on a compute node that does the following:

- Connects to a storage server
- Performs a LIST action
- Exits the session (specified with the "-e" flag)

```
[opc@exacs-node1 ~]$ exacli -l cloud_user_clustername -c 192.168.136.7 --xml
--cookie-jar -e list griddisk detail
```

NOT_SUPPORTED

Option	Description
`-c [username@]remotehost` or `--connect [username@]remotehost[:port]`	Specifies the remote node to which you want to connect. ExaCLI prompts for the user name if not specified.
`-l username` or `--login-name username`	Specifies the user name to log into the remote node. The preconfigured user is `cloud_user_clustername`.
`--xml`	Displays the output in XML format.
`--cookie-jar [filename]`	Specifies the filename of the cookie jar to use. If filename is not specified, the cookie is stored in a default cookie jar located at `HOME/.exacli/cookiejar`, where HOME is the home directory of the OS user running the ExaCLI command. The presence of a valid cookie allows the ExaCLI user to execute commands without requiring to login in subsequent ExaCLI sessions.
`-e command` or `-e 'command[; command]'` or `-e @batchFile`	Specifies either the ExaCLI commands to run or a batch file. ExaCLI exits after running the commands. If specifying multiple commands to run, enclose the commands in single quotes to prevent the shell from interpreting the semi-colon. Omit this option to start an interactive ExaCLI session.
`--cert-proxy proxy[:port]`	Specifies the proxy server to use when downloading certificates. If `port` is omitted, port 80 is used by default.
`-n` or `--no-prompt`	Suppresses prompting for user input.

NOT_SUPPORTED

- Notes for the --cookie-jar option:
  - The user name and password are sent to the remote node for authentication. On successful authentication, the remote node issues a cookie (the login credentials) that is stored in the specified `filename` on the database node. If `filename` is not specified, the cookie is stored in a default cookie jar located at `HOME/.exacli/cookiejar`, where HOME is the home directory of the operating system user running the ExaCLI command. For the opc user, the home is `/home/opc`.
  - The operating system user running the ExaCLI command is the owner of the cookie-jar file.
  - A cookie jar can contain multiple cookies from multiple users on multiple nodes in parallel sessions.
  - Cookies are invalidated after 24 hours.
  - If the cookie is not found or is no longer valid, ExaCLI prompts for the password. The new cookie is stored in the cookie jar identified by `filename`, or the default cookie jar if `filename` is not specified.
  - Even without the `--cookie-jar` option, ExaCLI still checks for cookies from the default cookie jar. However, if the cookie does not exist or is no longer valid, the new cookie will not be stored in the default cookie jar if the `--cookie-jar` option is not specified.
- Notes for the -e option:
  - ExaCLI exits after running the commands.
  - If specifying multiple commands to run, be sure to enclose the commands in single quotes to prevent the shell from interpreting the semi-colon.
  - The batch file is a text file that contains one or more ExaCLI commands to run.
- Notes for the `-n` (`--no-prompt`) option:
  - If ExaCLI needs additional information from the user, for example, if ExaCLI needs to prompt the user for a password (possibly because there were no valid cookies in the cookie-jar) or to prompt the user to confirm the remote node's identity, then ExaCLI prints an error message and exits.

Related Topics

- Connecting to a Storage Server with ExaCLI
  To use ExaCLI on storage servers, you will need to know your target storage server's IP address.

Connecting to a Storage Server with ExaCLI

To use ExaCLI on storage servers, you will need to know your target storage server's IP address.

If you do not know the IP address of the node you want to connect to, you can find it by viewing the contents of the `cellip.ora` file.

The following example illustrates how to do so on the UNIX command line for a quarter rack system. (Note that a quarter rack has three storage cells, and each cell has two connections, so a total of six IP addresses are shown.)

```
cat /etc/oracle/cell/
network-config/cellip.oracle
cell="192.168.136.5;cell="192.168.136.6"
cell="192.168.136.7;cell="192.168.136.8"
cell="192.168.136.9;cell="192.168.136.10"
```

If you are connecting to a storage cell for the first time using ExaCLI, you may be prompted to accept an SSL certificate. The ExaCLI output in this case will look like the following:

```
exacli -l cloud_user_clustername -c 192.168.136.7 --cookie-jar
No cookies found for cloud_user_clustername@192.168.136.7
Password: *********
EXA-30016: This connection is not secure. You have asked ExaCLI to connect to
cell 192.168.136.7 securely. The identity of 192.168.136.7 cannot be verified.
Got certificate from server:
C=US,ST=California,L=Redwood City,O=Oracle Corporation,OU=Oracle
Exadata,CN=ed1cl03clu01-priv2.usdc2.oraclecloud.com
Do you want to accept and store this certificate? (Press y/n)
```

Accept the self-signed Oracle certificate by pressing "y" to continue using ExaCLI.

ExaCLI Command Reference

You can execute various ExaCLI commands to monitor and manage Exadata Storage Servers associated with your Oracle Cloud Infrastructure Exadata DB system. ExaCLI allows you to get up-to-date, real-time information about your Exadata Cloud Service.

Use the LIST command with the following services and objects:

- **ACTIVEREQUEST**- Lists all active requests that are currently being served by the storage servers.
- **ALERTDEFINITION** - Lists all possible alerts and their sources for storage servers.
- **ALERTHISTORY** - Lists all alerts that have been issues for the storage servers.
- **CELL** - Used to list the details of a specific attribute of the storage servers or storage cells. The syntax is as follows:

  ```
  LIST CELL ATTRIBUTES A,B,C
  ```

  , with *A*, *B*, and *C* being attributes. To see all cell attributes, use the

  ```
  LIST CELL ATTRIBUTES ALL
  ```

  command.

- **CELLDISK** - Lists the attributes of the cell disks in the storage servers. Use the following syntax to list the cell disk details:

  ```
  LIST CELLDISK cell_disk_name
            DETAIL
  ```

  .

- **DATABASE** - Lists details of the databases. Uses the regular LIST command syntax:

  ```
  LIST DATABASE
  ```

  and

  ```
  LIST DATABASE DETAIL
  ```

  . You can also use this command to show an individual attribute with the following syntax:

  ```
  LIST DATABASE
            ATTRIBUTES NAME
  ```

  .

- **FLASHCACHE** - Lists the details of the Exadata system's flash cache. For this object, you can use the following syntax patterns:

  ```
  LIST FLASHCACHE DETAIL
  ```

  or

  ```
  LIST FLASHCACHE
            ATTRIBUTES attribute_name
  ```

  .

- **FLASHCACHECONTENT** - Lists the details of all objects in the flash cache, or the details of a specified object ID. To list all the details of all objects, use

  ```
  LIST FLASHCACHECONTENT DETAIL
  ```

  . To list details for a specific object, use a where clause as follows:

  ```
  LIST FLASHCACHECONTENT WHERE objectNumber=12345 DETAIL
  ```

  Note: To find the object ID of a specific object, you can query

  ```
  user_objects
  ```

  using the object's name to get the

  ```
  data_object_id
  ```

of a partition or table.

- **FLASHLOG** - Lists the attributes for the Oracle Exadata Smart Flash Log.
- **GRIDDISK** - Lists the details of a particular grid disk. The syntax is similar to the CELLDISK command syntax. To view all attributes:

```
LIST GRIDDISK grid_disk_name
          DETAIL
```

. To view specified attributes of the grid disk:

```
LIST
          GRIDDISK grid_disk_name ATTRIBUTES size,
        name
```

.

- **IBPORT** - Lists details of the InfiniBand ports. Syntax is

```
LIST IBPORT DETAIL
```

.

- **IORMPROFILE** - Lists any IORM profiles that have been set on the storage servers. You can also refer back to the profile attribute on the DATABASE object if a database has an IORM profile on it. Syntax is

```
LIST
          IORMPROFILE
```

.

- **LUN** - The LUN (logical unit number) object returns the number and the detail of the physical disks in the storage servers. List the LUNs of the disks with

```
LIST LUN
```

. List the details of each LUN with

```
LIST LUN lun_number
          DETAIL
```

.

- **METRICCURRRENT** - Lists the current metrics for a particular object type. Syntax is

```
LIST METRICCURRENT WHERE
            objectType = 'CELLDISK'
```

. This command also allows for sorting and results limits as seen in the following example:

```
LIST METRICCURRENT attributes name, metricObjectName
            ORDER BY metricObjectName asc, name desc LIMIT 5
```

- **METRICDEFINITION** - Lists metric definitions for the object that you can then get details for. With the command

```
LIST
          metricDefinition WHERE objectType=cell
```

, you can get all the metrics for that object type. You can then use the metric definition object again to get details for one of those specific metrics just listed:

```
LIST metricDefinition WHERE
          name= IORM_MODE DETAIL
```

.

- **METRICHISTORY** - List metrics over a specified period of time. For example, with the command

```
LIST METRICHISTORY WHERE ageInMinutes <
          30
```

, you can list all the metrics collected over the past 30 minutes. You can also use the predicate collectionTime to set a range from a specific time. Use collectionTime as shown in the follow example:

```
LIST METRICHISTORY WHERE collectionTime
          > '2018-04-01T21:12:00-10:00'
```

. The metric history object can also be used to see a specific metric using the object's name (for example,

```
LIST
          METRICHISTORY CT_FD_IO_RQ_SM
```

) or with a "where" clause to get objects with similar attributes like name (for example,

```
LIST METRICHISTORY WHERE
          name like 'CT_.*'
```

).

- **OFFLOADGROUP** - Lists the attributes for the offload group that are running on your storage servers. You can list all details for all groups with

```
LIST OFFLOADGROUP DETAIL
```

, or list the attributes for a specific group, as shown in the following example:

```
LIST
          OFFLOADGROUP offloadgroup4
```

. List specific attributes with

```
LIST OFFLOADGROUP ATTRIBUTES name
```

.

- **PHYSICALDISK** - Lists all physical disks. Use the results of

```
LIST PHYSICALDISK
```

to identify a specific disk for further investigation, then list the details of that disk using the command as follows:

```
LIST PHYSICALDISK 20:10 DETAIL
```

. To list the details of flash disks, use the command as follows:

```
LIST PHYSICALDISK
          FLASH_1_0 DETAIL
```

).

- **PLUGGABLEDATABASE** - Lists all PDBs. View the details of a specific PDB with

```
LIST PLUGGABLEDATABASE pdb_name
```

.

- **QUARANTINE** - Lists all SQL statements that you prevented from using Smart Scans. The syntax is

```
LIST QUARANTINE
          DETAIL
```

. You can also use a "where" clause on any of the available attributes.

Use the ExaCLI CREATE, ALTER, DROP, and LIST commands to act on the following Exadata Storage Server objects:

- **DIAGPACK** - Lists the diagnostic packages and their status in your Exadata system. The syntax is

```
LIST DIAGPACK
            [DETAIL]
```

, with DETAIL being an optional attribute. Use

```
CREATE DIAGPACK
```

with the

```
packStartTime
```

attribute to gather logs and trace files into a single compressed file for downloading, as in the following example:

```
CREATE DIAGPACK
            packStartTime=2019_12_15T00_00_00
```

. You can also use the value "now" with packStartTime:

```
CREATE DIAGPACK packStartTime=now
```

To download a diagnostic package, use

```
DOWNLOAD DIAGPACK
              package_namelocal_directory
```

. For example, the following command downloads a diagnostic package to the /tmp directory:

```
DOWNLOAD DIAGPACK  cfclcx2647_diag_2018_06_03T00_44_24_1 /tmp
```

- **IORMPLAN** - You can List, create, alter, and drop IORM plans using ExaCLI. To see the details of all IORM plans, use

```
LIST
          IORMPLAN DETAIL
```

. You can also use the command to create and alter IORM plans, and to apply plans to storage servers.

NOT_SUPPORTED

```
select object_name, data_object_id from user_objects where object_name =
'BIG_CENSUS';
OBJECT_NAME               DATA_OBJECT_ID
----------------------------------------
BIG_CENSUS                29152
```

Monitor Metrics for VM Cluster Resources

You can monitor the health, capacity, and performance of your VM clusters and databases with metrics, alarms, and notifications. You can use Oracle Cloud Infrastructure Console, Monitoring APIs, or Database Management APIs to view metrics.

Note: To view metrics you must have the required access as specified in an Oracle Cloud Infrastructure policy (whether you're using the Console, the REST API, or another tool). See Getting Started with Policies for information on policies.

> **WARNING:**
>
> Metrics, events, and audit events will not be sent if Cluster Ready Services (CRS) is not running before Autonomous Health Framework (AHF) starts.

Prerequisites for Using Metrics

The following prerequisites are required for the metrics to flow out of the VM cluster.

1. Metrics on the VM clusters depends on Oracle Trace File Analyzer (TFA) agent. Ensure that these components are up and running. AHF version **22.2.4** or higher is required for capturing metrics from the VM clusters. To start, stop, or check the status of TFA, see *Manage Oracle Trace File Analyzer*.
2. To view the metrics on the Oracle Cloud Infrastructure Console, the TFA flag `defaultocimonitoring` must be set to `ON`. This flag is set to `ON` by default and you need not perform any action to set this. If you are not seeing metrics on the Console, then as `root` user on the guest VM, check if the flag is set to `ON`.

   ```
   tfactl get defaultocimonitoring
   .-------------------------------------------------------------------.
   |                            <host name>                            |
   +-----------------------------------------------------------+-------+
   | Configuration Parameter                                   | Value |
   +-----------------------------------------------------------+-------+
   | Send CEF metrics to OCI Monitoring ( defaultOciMonitoring ) | ON    |
   '-----------------------------------------------------------+-------'
   ```

   If the `defaultocimonitoring` flag is set to `OFF`, then run the `tfactl set defaultocimonitoring=on` or `tfactl set defaultocimonitoring=ON` command to turn it on:

   ```
   tfactl set defaultocimonitoring=on
   Successfully set defaultOciMonitoring=ON
   .-------------------------------------------------------------------.
   |                            <host name>                            |
   +-----------------------------------------------------------+-------+
   | Configuration Parameter                                   | Value |
   +-----------------------------------------------------------+-------+
   | Send CEF metrics to OCI Monitoring ( defaultOciMonitoring ) | ON    |
   '-----------------------------------------------------------+-------'
   ```

3. The following network configurations are required.
   a. **Egress rules for outgoing traffic**: The default egress rules are sufficient to enable the required network path : For more information, see Default Security List .If you have blocked the outgoing traffic by modifying the default egress rules on your Virtual Cloud Network(VCN), you will need to revert the settings to allow outgoing traffic. The default egress rule allowing outgoing traffic (as shown in the *Rules Required for both Client and Backup Networks*) is as follows:
      - Stateless: No (all rules must be stateful)
      - Destination Type: CIDR
      - Destination CIDR: All <region> Services in Oracle Services Network
      - IP Protocol: 443 (HTTPS)
   b. **Public IP or Service Gateway**: The compute instance must have either a public IP address or a service gateway to be able to send compute instance metrics to the Monitoring service.
      If the instance does not have a public IP address, set up a service gateway on the virtual cloud network (VCN). The service gateway lets the instance send compute instance metrics to the Monitoring service without the traffic going over the internet. Here are special notes for setting up the service gateway to access the Monitoring service:
      i. When creating the service gateway, enable the service label called **All <region> Services in Oracle Services Network**. It includes the Monitoring service.
      ii. When setting up routing for the subnet that contains the instance, set up a route rule with **Target Type** set to **Service Gateway**, and the **Destination Service** set to **All <region> Services in Oracle Services Network**.

         For detailed instructions, see Access to Oracle Services: Service Gateway.

Related Topics

- Manage Oracle Trace File Analyzer
- Rules Required for Both the Client Network and Backup Network
  This topic has several general rules that enable essential connectivity for hosts in the VCN.
- Rules Required for Monitoring Service
  The compute instance must have either a public IP address or a service gateway to be able to send compute instance metrics to the Monitoring service.

View Metrics for VM Cluster

Perform the following steps to view the metrics for Guest VMs using the console.

> **Note:**
>
> When there is a network problem and Oracle Trace File Analyzer (TFA) is unable to post metrics, TFA will wait for one hour before attempting to retry posting the metrics. This is required to avoid creating a backlog of metrics processing on TFA.
>
> Potentially one hour of metrics will be lost between network restore and the first metric posted.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**. A list of VM clusters is displayed.
3. In the list of VM clusters, click the VM cluster for which you want to view the metrics. Details of the VM cluster you selected are displayed.
4. In the **Resources** section, click **Metrics**.
   A chart for each metrics is displayed. By default, the metrics for the last one hour are displayed.

   You can only select the `oci_database_cluster` namespace from the **Metric namespace** drop-down.
5. If you want to change the interval, select the required start time and end time. Alternatively, you can select the interval from the Quick Selects drop down menu. The metrics are refreshed immediately for the selected interval.
6. For each metric, you can choose the interval and statistic independently.
   - Interval - The time period for which the metric is calculated.
   - Statistic - The mathematical method by which the metric is calculated.
7. For each metric, you can choose the following options from the 'Options' drop down menu.
   - View Query in Metrics Explorer
   - Copy Chart URL
   - Copy Query (MQL)
   - Create an Alarm on this Query
   - Table View

For Detailed information on various options for viewing the metrics chart, see Viewing Default Metric Charts.

View Metrics for a Database

Perform the following steps to view the metrics for a database using the console.

> **Note:**
>
> When there is a network problem and Oracle Trace File Analyzer (TFA) is unable to post metrics, TFA will wait for one hour before attempting to retry posting the metrics. This is required to avoid creating a backlog of metrics processing on TFA.
>
> Potentially one hour of metrics will be lost between network restore and the first metric posted.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata on Oracle Public Cloud**.
2. Choose your **Compartment**. A list of VM clusters is displayed.
3. In the list of VM clusters, click the VM cluster that contains the database for which you want to view the metrics. Details of the VM cluster you selected are displayed.
4. In the list of databases, click the database for which you want to view the metrics.

5. In the **Resources** section, click **Metrics**.
   A chart for each metrics is displayed. By default, the metrics for the last one hour are displayed.
6. Select a namespace from the **Metric namespace** from where you wish to view metrics.

> **Note:**
>
> - When Database Management is enabled, you will have an option to choose from `oci_database` or `oracle_oci_database` namespace.
> - When Database Management is disabled, then you can view metrics only from the `oci_database` namespace.

7. If you want to change the interval, select the required start time and end time. Alternatively, you can select the interval from the Quick Selects drop down menu. The metrics are refreshed immediately for the selected interval.
8. For each metric, you can choose the interval and statistic independently.
   - Interval - The time period for which the metric is calculated.
   - Statistic - The mathematical method by which the metric is calculated.
9. For each metric, you can choose the following options from the 'Options' drop down menu.
   - View Query in Metrics Explorer
   - Copy Chart URL
   - Copy Query (MQL)
   - Create an Alarm on this Query
   - Table View

For Detailed information on various options for viewing the metrics chart, see Viewing Default Metric Charts.

View Metrics for a PDB

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata on Oracle Public Cloud**.
2. Choose your **Compartment**. A list of VM clusters is displayed.
3. In the list of VM clusters, click the VM cluster that contains the database for which you want to view the metrics. Details of the VM cluster you selected are displayed.
4. In the list of databases, click the database that contains the PBD for which you want to view the metrics.
5. Under **Resources**, click **Pluggable Databases**.
6. In the list of VM clusters, click the PDB that you wish to view metrics.
7. Select a namespace from the **Metric namespace** from where you wish to view metrics.

> **Note:**
>
> - When Database Management is enabled, you will have an option to choose from `oracle_oci_database` namespace.
> - When Database Management is disabled, then the system will display a banner asking you to enable Database Management to provide metrics.

View Metrics for VM Clusters in a Compartment

Perform the following steps to view the metrics for databases in a compartment using the console.

> **Note:**
>
> When there is a network problem and Oracle Trace File Analyzer (TFA) is unable to post metrics, TFA will wait for one hour before attempting to retry posting the metrics. This is required to avoid creating a backlog of metrics processing on TFA.
>
> Potentially one hour of metrics will be lost between network restore and the first metric posted.

1. Open the Oracle Cloud Infrastructure **Console** by clicking the menu icon next to **Oracle Cloud**.
2. From the left navigation list click **Observability & Management**.
3. Under **Monitoring**, click **Service Metrics**.
4. On the Service Metrics page, under **Compartment** select your compartment.
5. On the Service Metrics page, under **Metric Namespace** select `oci_database_cluster`.
6. If there are multiple VM clusters in the compartment you can show metrics aggregated across the clusters by selecting **Aggregate Metric Streams**.
7. If you want to limit the metrics you see, next to **Dimensions** click **Add** (click **Edit** if you have already added dimensions).
8. In the **Dimension Name** field select a dimension.
9. In the **Dimension Value** field select a value.
10. Click **Done**.
11. In the **Edit dimensions** dialog click **+Additional Dimension** to add an additional dimension. Click **X** to remove a dimension.
12. To create an alarm on a specific metric, click **Options** and select **Create an Alarm on this Query**. See *Managing Alarms* for information on setting and using alarms.

> **Note:**
>
> If you don't see any metrics, check the network settings and AHF version listed in the prerequisites section.

Related Topics

- Managing Alarms

View Metrics for Databases in a Compartment

Perform the following steps to view the metrics for databases in a compartment using the console.

> **Note:**
>
> When there is a network problem and Oracle Trace File Analyzer (TFA) is unable to post metrics, TFA will wait for one hour before attempting to retry posting the metrics. This is required to avoid creating a backlog of metrics processing on TFA.
>
> Potentially one hour of metrics will be lost between network restore and the first metric posted.

1. Open the Oracle Cloud Infrastructure **Console** by clicking the menu icon next to **Oracle Cloud**.
2. From the left navigation list click **Observability & Management**.
3. Under **Monitoring**, click **Service Metrics**.
4. On the Service Metrics page, under **Compartment** select your compartment.
5. On the Service Metrics page, under **Metric Namespace** select `oci_database`.
6. If there are multiple databases in the compartment you can show metrics aggregated across the databases by selecting **Aggregate Metric Streams**.
7. If you want to limit the metrics you see, next to **Dimensions** click **Add** (click **Edit** if you have already added dimensions).
8. In the **Dimension Name** field select a dimension.
9. In the **Dimension Value** field select a value.
10. Click **Done**.
11. In the **Edit dimensions** dialog click **+Additional Dimension** to add an additional dimension. Click **X** to remove a dimension.
12. To create an alarm on a specific metric, click **Options** and select **Create an Alarm on this Query**. See Managing Alarms for information on setting and using alarms.

Manage Oracle Trace File Analyzer

The deployment of the cloud-certified Autonomous Health Framework (AHF), which includes Oracle Trace File Analyzer, is managed by Oracle. You shouldn't install this manually on the guest VMs.

- To check the run status of Oracle Trace File Analyzer, run the `tfactl status` command as `root` or a non-root user:

```
# tfactl status
.-----------------------------------------------------------------------------
----------------------.
| Host           | Status of TFA | PID    | Port | Version    | Build
ID             | Inventory Status|
+----------------+---------------+--------+------+------------
+----------------------+------------+
| node1          | RUNNING       |  41312 | 5000 | 22.1.0.0.0 |
22100020220310214615| COMPLETE    |
| node2          | RUNNING       | 272300 | 5000 | 22.1.0.0.0 |
22100020220310214615| COMPLETE    |
'----------------+---------------+--------+------+------------
+----------------------+------------'
```

- To start the Oracle Trace File Analyzer daemon on the local node, run the `tfactl start` command as `root`:

```
# tfactl start
Starting TFA..
Waiting up to 100 seconds for TFA to be started..
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
Successfully started TFA Process..
. . . . .
TFA Started and listening for commands
```

- To stop the Oracle Trace File Analyzer daemon on the local node, run the `tfactl stop` command as `root`:

```
# tfactl stop
Stopping TFA from the Command Line
Nothing to do !
Please wait while TFA stops
Please wait while TFA stops
TFA-00002 Oracle Trace File Analyzer (TFA) is not running
TFA Stopped Successfully
Successfully stopped TFA..
```

Manage Database Service Agent

View the `/opt/oracle/dcs/log/dcs-agent.log` file to identify issues with the agent.

- To check the status of the Database Service Agent, run the `systemctl status` command:

```
# systemctl status dbcsagent.service
dbcsagent.service
Loaded: loaded (/usr/lib/systemd/system/dbcsagent.service; enabled; vendor
preset: disabled)
Active: active (running) since Fri 2022-04-0113:40:19UTC; 6min ago
Process: 9603ExecStopPost=/bin/bash -c kill `ps -fu opc |grep "java.*dbcs-
agent.*jar"|awk '{print $2}'` (code=exited, status=0/SUCCESS)
Main PID: 10055(sudo)
CGroup: /system.slice/dbcsagent.service
 10055sudo -u opc /bin/bash -c umask 077; /bin/java
```

- To start the agent if it is not running, run the `systemctl start` command as the `root` user:

```
systemctl start dbcsagent.service
```

Metrics for Oracle Exadata Database Service on Dedicated Infrastructure in the Monitoring Service

This article describes the metrics emitted by the Exadata Cloud Infrastructure Database service in the `oci_database_cluster` and `oci_database` namespaces for Oracle Databases.

Dimensions

All the metrics discussed in this topic include the following dimensions.

- RESOURCEID - The OCID of the VM Cluster.
- RESOURCENAME - The name of the VM Cluster.

NOT_SUPPORTED

The metrics listed in the following table are automatically available for the VM cluster.

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
`ASMDiskgroup Utilization`	**ASM Diskgroup Utilization**	percentage	Percentage of usable space used in a Disk Group. Usable space is the space available for growth. DATA disk group stores our Oracle database files. RECO disk group contains database files for recovery such as archives and flashback logs.	10 minutes	hostName deploymentType diskgroupName
`CpuUtilization`	**CPU Utilization**	percentage	Percent CPU utilization	1 minute	hostName deploymentType

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
`FilesystemUtilization`	**Filesystem Utilization**	percentage	Percent utilization of provisioned filesystem	1 minute	hostName deploymentType filesystemName
`LoadAverage`	**Load Average**	integer	System load average over 5 minutes	1 minute	hostName deploymentType
`MemoryUtilization`	**Memory Utilization**	percentage	Percentage of memory available for starting new applications, without swapping. The available memory can be obtained via the following command: `cat /proc/meminfo`	1 minute	hostName deploymentType
`NodeStatus`	**Node Status**	integer	Indicates whether the host is reachable.	1 minute	hostName deploymentType
`OcpusAllocated`	**OCPU Allocated**	integer	The number of OCPUs allocated	1 minute	deploymentType
`SwapUtilization`	**Swap Utilization**	percentage	Percent utilization of total swap space	1 minute	hostName deploymentType

NOT_SUPPORTED

The metrics listed in the following table are automatically available for the database.

Metric Name	Metric Display Name	Unit	Dsicription and Metric Chart Defaults	Collection Frequency	Dimensions
`CpuUtilization`	**CPU Utilization**	percentage	The CPU utilization expressed as a percentage, aggregated across all consumer groups. The utilization percentage is reported with respect to the number of CPUs the database is allowed to use, which is two times the number of OCPUs.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database|pdb} resourceName_{database|pdb}
`StorageUtilization`	**Storage Utilization**	percentage	The percentage of provisioned storage capacity currently in use. Represents the total allocated space for all tablespaces.	1 hour	deploymentType resourceId_{database|pdb} resourceName_{database|pdb}
`BlockChanges`	**DB Block Changes**	Changes per second	The Average number of blocks changed per second.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database|pdb} resourceName_{database|pdb}
`ExecuteCount`	**Execute Count**	Count	The number of user and recursive calls that executed SQL statements during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType

Metric Name	Metric Display Name	Unit	Dsicription and Metric Chart Defaults	Collection Frequency	Dimensions
`ExecuteCount`	**Execute Count**	Count	The number of user and recursive calls that executed SQL statements during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType
`CurrentLogons`	**Current Logons**	Count	The number of successful logons during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database|pdb} resourceName_{database|pdb}
`TransactionCount`	**Transaction Count**	Count	The combined number of user commits and user rollbacks during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database|pdb} resourceName_{database|pdb}
`UserCalls`	**User Calls**	Count	The combined number of logons, parses, and execute calls during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database|pdb} resourceName_{database|pdb}

Metric Name	Metric Display Name	Unit	Dsicription and Metric Chart Defaults	Collection Frequency	Dimensions
`ParseCount`	**Parse Count**	Count	The number of hard and soft parses during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database|pdb} resourceName_{database|pdb}
`StorageUsed`	**Storage Space Used**	GB	Total amount of storage space used by the database at the collection time.	1 hour	deploymentType resourceId_{database|pdb} resourceName_{database|pdb}
`StorageAllocated`	**Storage Space Allocated**	GB	Total amount of storage space allocated to the database at the collection time	1 hour	deploymentType resourceId_{database|pdb} resourceName_{database|pdb}
`StorageUsedByTablespace`	**Storage Space Used By Tablespace**	GB	Total amount of storage space used by tablespace at the collection time. In case of container database, this metric provides root container tablespaces.	1 hour	tablespaceName, tablespaceType deploymentType resourceId_{database|pdb} resourceName_{database|pdb}
StorageAllocatedByTablespace	**Allocated Storage Space By Tablespace**	GB	Total amount of storage space allocated to the tablespace at the collection time. In case of container database, this metric provides root container tablespaces.	1 hour	TablespaceName, tablespaceType, deploymentType, resourceId_{database|pdb} resourceName_{database|pdb}

Metric Name	Metric Display Name	Unit	Dsicription and Metric Chart Defaults	Collection Frequency	Dimensions
`StorageUtilizationByTablespace`	**Storage Space Utilization By Tablespace**	percentage	This indicates the percentage of storage space utilized by the tablespace at the collection time. In case of container database, this metric provides root container tablespaces..	1 hour	tablespaceName, tablespaceType deploymentType

Oracle Exadata Database Service on Dedicated Infrastructure Events

Exadata Cloud Infrastructure resources emit events, which are structured messages that indicate changes in resources.

- About Event Types on Exadata Cloud Infrastructure
  Learn about the event types available for Exadata Cloud Infrastructure resources.
- Prerequisites for Event Service
  The following prerequisites are required for the Events to flow out of the VM Cluster.
- Oracle Exadata Database Service on Dedicated Infrastructure Event Types
  The events in this section are emitted by the cloud Exadata infrastructure resource
- Oracle Exadata Database Service on Dedicated Infrastructure Maintenance Event Types
  The events in this section are emitted by the cloud Exadata infrastructure resource for Maintenance Events
- Exadata Cloud Infrastructure Critical and Information Event Types
  Exadata Cloud Infrastructure infrastructure resources emit "critical" and "information" data plane events that allow you to receive notifications when your infrastructure resource needs attention.
- Exadata Cloud Infrastructure VM Cluster Event Types
  Review the list of events that can be emitted by VM Cluster
- VM Node Subsetting Event Types
  Review the list of event types that VM Node Subsetting emits.
- Data Guard Association Event Types
  Review the list of event types that Data Guard associations emit.
- Oracle Database Home Event Types
  Review the list of events emitted by Oracle Database Homes.
- Database Event Types
  These are the event types that Oracle Databases in Exadata Cloud Service instances emit.
- Pluggable Database Event Types
  These are the event types that Oracle pluggable databases in Oracle Cloud Infrastructure emit.

- Database Service Events
  The Database Service emits events, which are structured messages that indicate changes in resources.
- Application VIP Event Types
  These are the event types that Application VIPs in Oracle Cloud Infrastructure emit.
- Interim Software Updates Event Types
  These are the event types that Interim Software Updates in Oracle Cloud Infrastructure emit.
- Serial Console Connection Event Types
  Review the list of event types that serial console connection emits.

About Event Types on Exadata Cloud Infrastructure

Learn about the event types available for Exadata Cloud Infrastructure resources.

Exadata Cloud Infrastructure resources emit events, which are structured messages that indicate changes in resources. For more information about Oracle Cloud Infrastructure Events, see *Overview of Events*. You may subscribe to events and be notified when they occur using the Oracle Notification service, see *Notifications Overview*.

Related Topics

- Overview of Events
- Notifications Overview

Prerequisites for Event Service

The following prerequisites are required for the Events to flow out of the VM Cluster.

The Event Service requires the following:

1. Events on the VM Cluster depends on Oracle Trace File Analyzer (TFA) agent. Ensure that these components are up and running. AHF version **22.2.2** or higher is required for capturing events from the VM Cluster. To start, stop, or check the status of TFA, see Incident Logs and Trace Files . To enable AHF Telemetry for the VM Cluster using the dbcli ulitilty, see AHF Telemetry Commands
2. The following network configurations are required.
   a. **Egress rules for outgoing traffic**: The default egress rules are sufficient to enable the required network path : For more information, see Default Security List .If you have blocked the outgoing traffic by modifying the default egress rules on your Virtual Cloud Network(VCN), you will need to revert the settings to allow outgoing traffic. The default egress rule allowing outgoing traffic (as shown in Security Rules for the Oracle Exadata Database Service on Dedicated Infrastructure) is as follows:
      - Stateless: No (all rules must be stateful)
      - Destination Type: CIDR
      - Destination CIDR: **All <region> Services in Oracle Services Network**
      - IP Protocol: TCP
      - Destination Port: 443 (HTTPS)
   b. **Public IP or Service Gateway**: The database server host must have either a public IP address or a service gateway to be able to send database server host metrics to the Monitoring service.

If the instance does not have a public IP address, set up a service gateway on the virtual cloud network (VCN). The service gateway lets the instance send database server host metrics to the Monitoring service without the traffic going over the internet. Here are special notes for setting up the service gateway to access the Monitoring service:

i. When creating the service gateway, enable the service label called **All <region> Services in Oracle Services Network**. It includes the Monitoring service.

ii. When setting up routing for the subnet that contains the instance, set up a route rule with **Target Type** set to **Service Gateway**, and the **Destination Service** set to **All <region> Services in Oracle Services Network**.

For detailed instructions, see Access to Oracle Services: Service Gateway.

Related Topics

- Service Gateway for the VCN
  Your VCN needs access to both Object Storage for backups and Oracle YUM repos for OS updates.

Oracle Exadata Database Service on Dedicated Infrastructure Event Types

The events in this section are emitted by the cloud Exadata infrastructure resource

> **Note:**
>
> Exadata systems that use the old DB system resource model are deprecated and will be desupported in a future release. The DB system event are not described.

Friendly Name	Event Type
Cloud Exadata Infrastructure - Create Begin	`com.oraclecloud.databaseservice.createcloudexadatainfrastructure.begin`
Cloud Exadata Infrastructure - Create End	`com.oraclecloud.databaseservice.createcloudexadatainfrastructure.end`
Cloud Exadata Infrastructure - Change Compartment Begin	`com.oraclecloud.databaseservice.changecloudexadatainfrastructurecompartment.begin`
Cloud Exadata Infrastructure - Change Compartment End	`com.oraclecloud.databaseservice.changecloudexadatainfrastructurecompartment.end`
Cloud Exadata Infrastructure - Critical See Exadata Cloud Service Infrastructure Critical and Information Event Types for details	`com.oraclecloud.databaseservice.cloudexadatainfrastructure.critical`
Cloud Exadata Infrastructure - Delete Begin	`com.oraclecloud.databaseservice.deletecloudexadatainfrastructure.begin`
Cloud Exadata Infrastructure - Delete End	`com.oraclecloud.databaseservice.deletecloudexadatainfrastructure.end`
Cloud Exadata Infrastructure - Information See Exadata Cloud Service Infrastructure Critical and Information Event Types for details	`com.oraclecloud.databaseservice.cloudexadatainfrastructure.information`
Cloud Exadata Infrastructure - Update Begin	`com.oraclecloud.databaseservice.updatecloudexadatainfrastructure.begin`

Friendly Name	Event Type
Cloud Exadata Infrastructure - Update End	com.oraclecloud.databaseservice.updatecloudexadatainfrastructure.end

This is a reference event for a Cloud Exadata Infrastructure resource:

```
{
  "cloudEventsVersion": "0.1",
  "eventId": "<unique_ID>",
  "eventType":
"com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenance.end",
  "source": "DatabaseService",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1.<unique_ID>"
  },
  "data": {
    "compartmentId": "ocid1.compartment.oc1.<unique_ID>",
    "compartmentName": "example_name",
    "resourceName": "my_exadata_infrastructure",
    "resourceId": "ocid1.dbsystem.oc1.eu-frankfurt-1.<unique_ID>", ,
    "availabilityDomain": "tXPJ:EU-FRANKFURT-1-AD-3",
    "freeFormTags": {
      "Department": "Finance"
    },
    "definedTags": {
      "Operations": {
        "CostCenter": "42"
      }
    },
    "additionalDetails" : {
"subnetId" : "ocid1.subnet.oc1.eu-frankfurt-1.<unique_ID>",
"lifecycleState" : "MAINTENANCE_IN_PROGRESS",
"sshPublicKeys" : "...",
"cpuCoreCount" : 32,
"version" : "19.2.8.0.0.191119",
"nsgIds" : "null",
"backupSubnetId" : "ocid1.subnet.oc1.eu-frankfurt-1.<unique_ID>",
"licenseType" : "BRING_YOUR_OWN_LICENSE",
"dataStoragePercentage" : 80,
"patchHistoryEntries" : "null",
"lifecycleMessage" : "The underlying infrastructure of this system (cell
storage) is being updated and this will not impact database
                      availability.",
"exadataIormConfig" : "ExadataIormConfigCache(lifecycleState=DISABLED,
lifeCycleDetails=null, objective=Auto,
                       dbPlans=[DbIormConfigCache(dbName=default, share=null,
flashCacheLimit=null), DbIormConfigCache(dbName=<my_database1>,
                       share=null, flashCacheLimit=null),
DbIormConfigCache(dbName=<my_database2>, share=null, flashCacheLimit=null),
                       DbIormConfigCache(dbName=<my_database3>, share=null,
flashCacheLimit=null), DbIormConfigCache(dbName=<my_database4>,
```

```
                        share=null, flashCacheLimit=null),
DbIormConfigCache(dbName=<my_database5>, share=null, flashCacheLimit=null),
                        DbIormConfigCache(dbName=<my_database6>, share=null,
flashCacheLimit=null), DbIormConfigCache(dbName=<my_database7>,
                        share=null, flashCacheLimit=null),
DbIormConfigCache(dbName=<my_database8>, share=null, flashCacheLimit=null),
                        DbIormConfigCache(dbName=<my_database9>, share=null,
flashCacheLimit=null), DbIormConfigCache(dbName=<my_database10>,
                        share=null, flashCacheLimit=null),
DbIormConfigCache(dbName=<my_database11>, share=null, flashCacheLimit=null)],
                        undoData=null)"
}
},
"eventID" : "<unique_ID>",
"extensions" : {
"compartmentId" : "ocid1.compartment.oc1.<unique_ID>"
}
}
```

This is a reference event for Cloud Exadata Infrastructure - Add Storage Capacity Begin:

```
{
  "id":
"ocid1.eventschema.oc1.phx.z1nzw5klc4r7ar1vkxunfvyfhtwmeaaylr0j5hjnu2j5uozwlie
xa53gwlk4",
  "serviceName": "Database",
  "displayName": "Cloud Exadata Infrastructure - Add Storage Capacity Begin",
  "eventType":
"com.oraclecloud.databaseservice.addstoragecapacitycloudexadatainfrastructure.
begin",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2023-01-06T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",
      "type": [
        "null",
        "string"
      ]
    },
    {
```

```
      "name": "shape",
      "type": "string"
    },
    {
      "name": "message",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "description",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "timeZone",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "maintenanceMode",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "maintenanceSubType",
      "type": [
        "null",
        "string"
      ]
    }
  ],
  "exampleEvent": {
    "eventType":
"com.oraclecloud.databaseservice.addstoragecapacitycloudexadatainfrastructure.
begin",
    "cloudEventsVersion": "0.1",
    "eventTypeVersion": "2.0",
    "source": "databaseservice",
    "eventID": "10274771-3706-4624-99d1-e036805a9ca7",
    "eventTime": "2023-01-06T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
      "eventGroupingId": "csida87218404b4291914305ec7a5a86/
d53ffb13f83244bbbfb8c7d0a8f0e2eb/FB95D76D5123C152C25DBF288489077F",
      "eventName": "AddStorageCapacityCloudExadataInfrastructure",
      "compartmentId": "ocid1.compartment.oc1.....unique_id",
      "compartmentName": null,
      "resourceName": "my_cloud_exadata_infrastructure",
```

```
"resourceId": "ocid1.cloudexadatainfrastructure.oc1.....unique_id",
"resourceVersion": null,
"availabilityDomain": "",
"tagSlug": "tag_slug",
"identity": {
  "principalName": null,
  "principalId": null,
  "authType": null,
  "callerName": null,
  "callerId": null,
  "tenantId": null,
  "ipAddress": null,
  "credentials": null,
  "authZPolicies": null,
  "userGroups": null,
  "userAgent": null,
  "consoleSessionId": null
},
"request": {
  "id": "7e83c538-28bf-453d-9fb7-125bf70546c4",
  "path": null,
  "action": null,
  "parameters": null,
  "headers": null
},
"response": {
  "status": null,
  "responseTime": null,
  "headers": null,
  "payload": null,
  "message": null
},
"stateChange": {
  "previous": null,
  "current": {
    "lifecycleState": "AVAILABLE",
    "shape": "Exadata.X9M",
    "displayName": "my_display_name",
    "freeTags": {},
    "definedTags": {}
  }
},
"additionalDetails": {
  "timeCreated": "2023-01-06T21:16:04.000Z",
  "timeUpdated": "2023-01-06T21:16:04.000Z",
  "lifecycleState": "AVAILABLE",
  "lifecycleDetails": null,
  "description": null,
  "message": null,
  "shape": "Exadata.X9M",
  "timeZone": null,
  "maintenanceMode": null,
  "maintenanceSubType": null
},
"internalDetails": {
  "attributes": null
```

```
      }
    }
  },
  "timeCreated": "2023-01-06T21:16:04.000Z"
}
```


This is a reference event for Cloud Exadata Infrastructure - Add Storage Capacity End:


```
{
  "id":
"ocid1.eventschema.oc1.phx.4aeklze2co1ynub2ojmu49shhduq9gh5qg6fvudm7h77w3og8sf
kau6a3not",
  "serviceName": "Database",
  "displayName": "Cloud Exadata Infrastructure - Add Storage Capacity End",
  "eventType":
"com.oraclecloud.databaseservice.addstoragecapacitycloudexadatainfrastructure.
end",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2023-01-06T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "shape",
      "type": "string"
    },
    {
      "name": "message",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "description",
      "type": [
```

```
            "null",
            "string"
          ]
        },
        {
          "name": "timeZone",
          "type": [
            "null",
            "string"
          ]
        },
        {
          "name": "maintenanceMode",
          "type": [
            "null",
            "string"
          ]
        },
        {
          "name": "maintenanceSubType",
          "type": [
            "null",
            "string"
          ]
        }
      ],
      "exampleEvent": {
        "eventType":
"com.oraclecloud.databaseservice.addstoragecapacitycloudexadatainfrastructure.
end",
        "cloudEventsVersion": "0.1",
        "eventTypeVersion": "2.0",
        "source": "databaseservice",
        "eventID": "b12abcc0-110a-9120-aab5-9a34bc799e72",
        "eventTime": "2023-01-06T21:16:04.000Z",
        "contentType": "application/json",
        "data": {
          "eventGroupingId":
"csida2cd1c8442f9b9fc16354a1f0912/95202d41125e4ce18e8dd52fa9f57f5e/
545A43343BC1D5020A85AA2919C06E25",
          "eventName": "AddStorageCapacityCloudExadataInfrastructure",
          "compartmentId": "ocid1.compartment.oc1.....unique_id",
          "compartmentName": null,
          "resourceName": "my_cloud_exadata_infrastructure",
          "resourceId": "ocid1.cloudexadatainfrastructure.oc1.....unique_id",
          "resourceVersion": null,
          "availabilityDomain": "",
          "tagSlug": "tag_slug",
          "identity": {
            "principalName": null,
            "principalId": null,
            "authType": null,
            "callerName": null,
            "callerId": null,
            "tenantId": null,
            "ipAddress": null,
```

```
        "credentials": null,
        "authZPolicies": null,
        "userGroups": null,
        "userAgent": null,
        "consoleSessionId": null
      },
      "request": {
        "id": "111b9da5-a7a7-4aca-bd05-a51558f7df55",
        "path": null,
        "action": null,
        "parameters": null,
        "headers": null
      },
      "response": {
        "status": null,
        "responseTime": null,
        "headers": null,
        "payload": null,
        "message": null
      },
      "stateChange": {
        "previous": null,
        "current": {
          "lifecycleState": "AVAILABLE",
          "shape": "Exadata.X9M",
          "displayName": "my_display_name",
          "freeTags": {},
          "definedTags": {}
        }
      },
      "additionalDetails": {
        "timeCreated": "2023-01-06T21:16:04.000Z",
        "timeUpdated": "2023-01-06T21:16:04.000Z",
        "lifecycleState": "AVAILABLE",
        "lifecycleDetails": null,
        "description": null,
        "message": null,
        "shape": "Exadata.X9M",
        "timeZone": null,
        "maintenanceMode": null,
        "maintenanceSubType": null
      },
      "internalDetails": {
        "attributes": null
      }
    }
  },
  "timeCreated": "2023-01-06T21:16:04.000Z"
}
```


This is a reference event for Cloud Exadata Infrastructure - Update Begin


```
{
  "id":
"ocid1.eventschema.oc1.phx.jlx9t3z6igwglicpbba6xs1uaewcb8txsegnuykc65n8rxl5tqd
```

```
26ect7i3f",
  "serviceName": "Database",
  "displayName": "Cloud Exadata Infrastructure - Update Begin",
  "eventType":
"com.oraclecloud.databaseservice.updatecloudexadatainfrastructure.begin",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "id",
      "type": "string"
    },
    {
      "name": "defineTags",
      "type": [
        "null",
        "Map<String, Map<String, Object>>"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "Map<String, String>"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "compartmentId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "availabilityDomain",
```

```
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "description",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "tenantId",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "message",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "shape",
        "type": [
          "null",
          "String"
        ]
      },
      {
        "name": "timeZone",
        "type": [
          "null",
          "string"
        ]
      }
    ],
    "exampleEvent": {
      "cloudEventsVersion": "0.1",
      "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
      "eventType":
"com.oraclecloud.databaseservice.updatecloudexadatainfrastructure.begin",
      "source": "databaseservice",
      "eventTypeVersion": "2.0",
      "eventTime": "2019-06-27T21:16:04.000Z",
      "contentType": "application/json",
      "data": {
        "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
        "eventName": "UpdateCloudExadataInfrastructure",
        "compartmentName": "example_compartment",
        "resourceName": "my_container_database",
        "resourceId": "ocid1.cloudexadatainfrastructure.oc1.....unique_id",
```

```
      "resourceVersion": null,
      "additionalDetails": {
        "availabilityDomain": "all",
        "compartmentId": "ocid1.compartment.oc1.......unique_id",
        "freeFormTags": {},
        "definedTags": {},
        "lifecycleState": "AVAILABLE"
      }
    }
  },
  "timeCreated": "2020-06-15T16:31:31.979Z"
}
```


This is a reference event for Cloud Exadata Infrastructure - Update End


```
{
  "id":
"ocid1.eventschema.oc1.phx.aq2fuvh1nh9h71bnyclhmsuj3bky7dr304xj7nejajjzwbnh2n4
0zy3tdand",
  "serviceName": "Database",
  "displayName": "Cloud Exadata Infrastructure - Update End",
  "eventType":
"com.oraclecloud.databaseservice.updatecloudexadatainfrastructure.end",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "id",
      "type": "string"
    },
    {
      "name": "defineTags",
      "type": [
        "null",
        "Map<String, Map<String, Object>>"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "Map<String, String>"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
```

```
    "name": "lifecycleState",
    "type": "string"
  },
  {
    "name": "lifecycleDetails",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "compartmentId",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "availabilityDomain",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "description",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "tenantId",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "message",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "shape",
    "type": [
      "null",
      "String"
    ]
  },
  {
    "name": "timeZone",
    "type": [
      "null",
```

```
        "string"
      ]
    }
  ],
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
    "eventType":
"com.oraclecloud.databaseservice.updatecloudexadatainfrastructure.end",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
      "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
      "eventName": "UpdateCloudExadataInfrastructure",
      "compartmentName": "example_compartment",
      "resourceName": "my_container_database",
      "resourceId": "ocid1.dbsystem-.....unique_id",
      "resourceVersion": null,
      "additionalDetails": {
        "availabilityDomain": "all",
        "compartmentId": "ocid1.compartment.oc1.......unique_id",
        "freeFormTags": {},
        "definedTags": {},
        "lifecycleState": "AVAILABLE"
      }
    }
  },
  "timeCreated": "2020-06-15T16:31:31.979Z"
}
```


Oracle Exadata Database Service on Dedicated Infrastructure Maintenance Event Types

The events in this section are emitted by the cloud Exadata infrastructure resource for Maintenance Events

> **Note:**
>
> Exadata systems that use the old DB system resource model are deprecated and will be desupported in a future release. The DB system event are not described.

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure – Maintenance Scheduled	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancescheduled`	• **Rolling:** Oracle Cloud Operations is announcing the availability of a new quarterly maintenance update for Cloud Exadata Infrastructure. Oracle has scheduled the installation of this new update on your service instance <*infra-name*>, `ocid` <*infra-ocid*> on <*time-scheduled*>. The maintenance method for this maintenance is <*maintenance-method*> as selected per the maintenance preferences. • **Non Rolling:** Oracle Cloud Operations is announcing the availability of a new quarterly maintenance update for Cloud Exadata Infrastructure. Oracle has scheduled the installation of this new update on your service instance <*infra-name*>, `ocid` <*infra-ocid*> on <*time-scheduled*>. The maintenance method for this maintenance is <*maintenance-method*> as selected per the maintenance preferences. Non-rolling maintenance minimizes maintenance time but will result in full system downtime.

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure – Maintenance Reminder	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancereminder`	• **Rolling:** This is an Oracle Cloud Operations reminder notice. Oracle has scheduled a quarterly maintenance update installation for Cloud Exadata Infrastructure <*infra-name*>, ocid <*ocid*> in approximately <*no-of-days*> days on <*time-scheduled*>. The maintenance method for this maintenance is <*maintenance-method*> as selected per the maintenance preferences. • **Non Rolling:** This is an Oracle Cloud Operations reminder notice. Oracle has scheduled a quarterly maintenance update installation for Cloud Exadata Infrastructure <*infra-name*>, ocid <*ocid*> in approximately <*no-of-days*> days on <*time-scheduled*>. The maintenance method for this maintenance is <*maintenance-method*> as selected per the maintenance preferences. Non-rolling maintenance minimizes maintenance time but will result in full system downtime.
Cloud Exadata Infrastructure - Maintenance Begin	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenance.begin`	This is an Oracle Cloud Operations notice regarding the quarterly maintenance update installation for your Cloud Exadata Infrastructure instance <*infra-name*>, ocid <*infra-ocid*>. The update installation for the service started at <*time scheduled*>. A follow-up notice will be sent when the maintenance update operation has completed.
Cloud Exadata Infrastructure - Maintenance End Success	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenance.end.success`	This is an Oracle Cloud Operations notice that your Cloud Exadata Infrastructure quarterly maintenance update installation for service instance <*infra-name*>, ocid <*infra-ocid*> which started at <*maintenance-start-time*> is now successfully complete.

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure - Maintenance End Failed	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenance.end.failed.`	This is an Oracle Cloud Operations notice that your Cloud Exadata Infrastructure quarterly maintenance update installation for service instance <*infra-name*>, ocid <*infra-ocid*> which started at <*maintenance-start-time*> has failed to complete due to technical reasons and operations team are currently looking into the issue. You will receive regular notifications to track progress of this maintenance.
Cloud Exadata Infrastructure - Maintenance VM Begin	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancevm.begin.`	This is an Oracle Cloud Operations notice regarding the quarterly maintenance update of Virtual Machines component of your Cloud Exadata Infrastructure instance <*infra-name*>, ocid <*infra-ocid*>, Database Server <*dbserver name*>, `ocid` <*dbserver ocid*> has started. A follow-up notice will be sent when Virtual Machines maintenance operation has completed.
Cloud Exadata Infrastructure - MaintenanceVM End	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancevm.end`	This is an Oracle Cloud Operations notice that quarterly maintenance update of the Database Server component of your Cloud Exadata Infrastructure instance <*infra-name*>, `ocid` <*infra-ocid*>; Database Server <*dbserver name*> `ocid` <*dbserver ocid*> has completed.
Cloud Exadata Infrastructure - Maintenance Storage Servers Start	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancestorageservers.start`	This is an Oracle Cloud Operations notice regarding the quarterly maintenance update of Storage servers component of your Cloud Exadata Infrastructure instance <*infra-name*>, `ocid` <*infra-ocid*> has started. A follow-up notice will be sent when storage servers maintenance operation has completed.
Cloud Exadata Infrastructure - Maintenance Storage Servers End	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancestorageservers.end`	This is an Oracle Cloud Operations notice that quarterly maintenance update of Storage servers component of your Cloud Exadata Infrastructure instance <*infra-name*>, `ocid` <*infra-ocid*> has completed.

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure - Maintenance Network Switches Begin	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancenetworkswitches.begin`	This is an Oracle Cloud Operations notice regarding the quarterly maintenance update of the network fabric switches component of your Cloud Exadata Infrastructure instance <*infra-name*>, ocid <*infra-ocid*> has started. A follow-up notice will be sent when the network fabric switches maintenance operation has completed.
Cloud Exadata Infrastructure - Maintenance Network Switches End	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancenetworkswitches.end`	This is an Oracle Cloud Operations notice that quarterly maintenance update of the network fabric switches component of your Cloud Exadata Infrastructure instance <*infra-name*>, ocid <*infra-ocid*> has completed.
Cloud Exadata Infrastructure - Maintenance Custom Action Time Begin	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancecustomactiontime.begin`	This is an Oracle Cloud Operations notice that the custom action timeout for your Cloud Exadata Infrastructure instance <*infra-name*>, ocid <*infra-ocid*>; Database Server <*dbserver name*>, ocid <*dbserver ocid*> has started. A follow-up notice will be sent when the custom action timeout has ended.
Cloud Exadata Infrastructure - Maintenance Custom Action Time End	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancecustomactiontime.end`	This is an Oracle Cloud Operations notice that the custom action timeout for your Cloud Exadata Infrastructure instance <*infra-name*>, ocid <*infra-ocid*>; Database Server <*dbserver name*>, ocid <*dbserver ocid*> has ended.
Cloud Exadata Infrastructure - Maintenance Rescheduled	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancerescheduled`	Oracle Cloud Operations is announcing reschedule of a quarterly maintenance update for Cloud Exadata Infrastructure. A maintenance run has been rescheduled on your service instance <*infra-name*>, ocid <*infra-ocid*> to <*new-schedule-time*>.

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure - Maintenance Method Change	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancemethodchange`	Oracle Cloud Operations is announcing a change related to quarterly maintenance update for Cloud Exadata Infrastructure. There's a change in maintenance method on your service instance <*infra-name*>, ocid <*infra-ocid*> to <*new-maintenance-method*>.

This is a reference event for a Cloud Exadata Infrastructure resource:

```
{
  "cloudEventsVersion": "0.1",
  "eventId": "<unique_ID>",
  "eventType":
"com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenance.end",
  "source": "DatabaseService",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1.<unique_ID>"
  },
  "data": {
    "compartmentId": "ocid1.compartment.oc1.<unique_ID>",
    "compartmentName": "example_name",
    "resourceName": "my_exadata_infrastructure",
    "resourceId": "ocid1.dbsystem.oc1.eu-frankfurt-1.<unique_ID>", ,
    "availabilityDomain": "tXPJ:EU-FRANKFURT-1-AD-3",
    "freeFormTags": {
      "Department": "Finance"
    },
    "definedTags": {
      "Operations": {
        "CostCenter": "42"
      }
    },
    "additionalDetails" : {
"subnetId" : "ocid1.subnet.oc1.eu-frankfurt-1.<unique_ID>",
"lifecycleState" : "MAINTENANCE_IN_PROGRESS",
"sshPublicKeys" : "...",
"cpuCoreCount" : 32,
"version" : "19.2.8.0.0.191119",
"nsgIds" : "null",
"backupSubnetId" : "ocid1.subnet.oc1.eu-frankfurt-1.<unique_ID>",
"licenseType" : "BRING_YOUR_OWN_LICENSE",
"dataStoragePercentage" : 80,
"patchHistoryEntries" : "null",
"lifecycleMessage" : "The underlying infrastructure of this system (cell
storage) is being updated and this will not impact database
                      availability.",
"exadataIormConfig" : "ExadataIormConfigCache(lifecycleState=DISABLED,
```

```
lifeCycleDetails=null, objective=Auto,
                       dbPlans=[DbIormConfigCache(dbName=default, share=null,
flashCacheLimit=null), DbIormConfigCache(dbName=<my_database1>,
                       share=null, flashCacheLimit=null),
DbIormConfigCache(dbName=<my_database2>, share=null, flashCacheLimit=null),
                       DbIormConfigCache(dbName=<my_database3>, share=null,
flashCacheLimit=null), DbIormConfigCache(dbName=<my_database4>,
                       share=null, flashCacheLimit=null),
DbIormConfigCache(dbName=<my_database5>, share=null, flashCacheLimit=null),
                       DbIormConfigCache(dbName=<my_database6>, share=null,
flashCacheLimit=null), DbIormConfigCache(dbName=<my_database7>,
                       share=null, flashCacheLimit=null),
DbIormConfigCache(dbName=<my_database8>, share=null, flashCacheLimit=null),
                       DbIormConfigCache(dbName=<my_database9>, share=null,
flashCacheLimit=null), DbIormConfigCache(dbName=<my_database10>,
                       share=null, flashCacheLimit=null),
DbIormConfigCache(dbName=<my_database11>, share=null, flashCacheLimit=null)],
                       undoData=null)"
}
},
"eventID" : "<unique_ID>",
"extensions" : {
"compartmentId" : "ocid1.compartment.oc1.<unique_ID>"
}
}
```

Exadata Cloud Infrastructure Critical and Information Event Types

Exadata Cloud Infrastructure infrastructure resources emit "critical" and "information" data plane events that allow you to receive notifications when your infrastructure resource needs attention.

Exadata Cloud Service infrastructure resources emit "critical" and "information" data plane events that allow you to receive notifications when your infrastructure resource needs urgent attention ("critical" events), or notifications for events that are not critical, but which you may want to monitor ("information" events). The eventType values for these events are the following:

- `com.oraclecloud.databaseservice.exadatainfrastructure.critical`
- `com.oraclecloud.databaseservice.exadatainfrastructure.information`

These events use the `additionalDetails` section of the event message to provide specific details about what is happening within the infrastructure resource emitting the event. In the `additionalDetails` section, the `eventName` field provides the name of the critical or information event. *(Note that some fields in the example that follows have been omitted for brevity.)*

```
 {
  "eventType" :
"com.oraclecloud.databaseservice.exadatainfrastructure.critical",
  ....
  "data" : {
   ....
     "additionalDetails" : {
      ....
```

```
      "description" : "SQL statement terminated by Oracle Database Resource
Manager due to excessive consumption of CPU and/or I/O.
                      The execution plan associated with the terminated SQL
stmt is quarantined. Please find the sql identifier in
                      sqlId field of this JSON payload. This feature protects
an Oracle database from performance degradation.
                      Please review the SQL statement. You can see the
statement using the following commands: \"set serveroutput off\",
                      \"select sql_id, sql_text from v$sqltext where sql_id
=<sqlId>\", \"set serveroutput on\"",
      "component" : "storage",
      "infrastructureType" : "exadata",
      "eventName" : "HEALTH.INFRASTRUCTURE.CELL.SQL_QUARANTINE",
      "quarantineMode" : "\"FULL Quarantine\""
       ....
    }
  },
  "eventID" : "<unique_ID>",
  ....
  }
}
```

In the tables below, you can read about the conditions and operations that trigger "critical" and "information" events. Each condition or operation is identified by a unique `eventName` value.

Critical events for Exadata Cloud Service infrastructure:

Critical Event - EventName	Description
`HEALTH.INFRASTRUCTURE.CELL.SQL_QUARANTINE`	SQL statement terminated by Oracle Database Resource Manager due to excessive consumption of CPU and/or I/O. The execution plan associated with the terminated SQL stmt is quarantined. Please find the sql identifier in sqlId field of this JSON payload. This feature protects an Oracle database from performance degradation. Please review the SQL statement. You can see the statement using the following commands: • `\"set serveroutput off\"` • `\"select sql_id, sql_text from v$sqltext where sql_id =<sqlId>\"` • `\"set serveroutput on\"`

Informational events for Exadata Cloud Service infrastructure:

Information Event - EventName	Description
`HEALTH.INFRASTRUCTURE.CELL.FLASH_DISK_FAILURE`	Flash Disk Failure has been detected. This is being investigated by Oracle Exadata team and the disk will be replaced if needed. No action needed from the customer.

NOT_SUPPORTED

In the following example of a "critical" event, you can see within the `additionalDetails` section of the event message that this particular message concerns an SQL statement that was terminated by Oracle Database Resource Manager because it was consuming excessive

CPU or I/O resources. The `eventName` and `description` fields within the `additionalDetails` section provide information regarding the critical situation:

```
 {
  "eventType" :
"com.oraclecloud.databaseservice.exadatainfrastructure.critical",
  "cloudEventsVersion" : "0.1",
  "eventTypeVersion" : "2.0",
  "source" : "Exadata Storage",
  "eventTime" : "2021-07-30T04:53:18Z",
  "contentType" : "application/json",
  "data" : {
    "compartmentId" : "ocid1.tenancy.oc1.<unique_ID>",
    "compartmentName" : "example_name",
    "resourceName" : "my_exadata_resource",
    "resourceId" : "ocid1.dbsystem.oc1.phx.<unique_ID>",
    "availabilityDomain" : "phx-ad-2",
     "additionalDetails" : {
      "serviceType" : "exacs",
      "sqlID" : "gnwfm1jgqcfuu",
      "systemId" : "ocid1.dbsystem.oc1.eu-frankfurt-1.<unique_ID>",
      "creationTime" : "2021-05-14T13:29:28+00:00",
      "dbUniqueID" : "1558836122",
      "quarantineType" : "SQLID",
      "dbUniqueName" : "AB0503_FRA1S6",
      "description" : "SQL statement terminated by Oracle Database Resource
Manager due to excessive consumption of CPU and/or I/O.
                       The execution plan associated with the terminated SQL
stmt is quarantined. Please find the sql identifier in sqlId
                       field of this JSON payload. This feature protects an
Oracle database from performance degradation.
                       Please review the SQL statement. You can see the
statement using the following commands: \"set serveroutput off\",
                       \"select sql_id, sql_text from v$sqltext where sql_id
=<sqlId>\", \"set serveroutput on\"",
      "quarantineReason" : "Manual",
      "asmClusterName" : "None",
      "component" : "storage",
      "infrastructureType" : "exadata",
      "name" : "143",
      "eventName" : "HEALTH.INFRASTRUCTURE.CELL.SQL_QUARANTINE",
      "comment" : "None",
      "quarantineMode" : "\"FULL Quarantine\"",
      "rpmVersion" : "OSS_20.1.8.0.0_LINUX.X64_210317",
      "cellsrvChecksum" : "14f73eb107dc1be0bde757267e931991",
      "quarantinePlan" : "SYSTEM"
    }
  },
  "eventID" : "<unique_ID>",
  "extensions" : {
    "compartmentId" : "ocid1.tenancy.oc1.<unique_ID>"
  }
}
```

NOT_SUPPORTED

In the following example of an "information" event, you can see within the `additionalDetails` section of the event message that this particular message concerns a flash disk failure that is being investigated by the Oracle Exadata operations team. The `eventName` and `description` fields within the `additionalDetails` section provide information regarding the event:

```
{
  "eventType" :
"com.oraclecloud.databaseservice.exadatainfrastructure.information",
  "cloudEventsVersion" : "0.1",
  "eventTypeVersion" : "2.0",
  "source" : "Exadata Storage",
  "eventTime" : "2021-12-17T19:14:42Z",
  "contentType" : "application/json",
  "data" : {
    "compartmentId" :
"ocid1.tenancy.oc1..aaaaaaaao3lj36x6lwxyvc4wausjouca7pwyjfwb5ebsq5emrpqlql2gj5
iq",
    "compartmentName" : "intexadatateam",
    "resourceId" :
"ocid1.dbsystem.oc1.phx.abyhqljt5y3taezn7ug445fzwlngjfszbedxlcbctw45ykkaxyzc5i
sxoula",
    "availabilityDomain" : "phx-ad-2",
    "additionalDetails" : {
      "serviceType" : "exacs",
      "component" : "storage",
      "systemId" :
"ocid1.dbsystem.oc1.phx.abyhqljt5y3taezn7ug445fzwlngjfszbedxlcbctw45ykkaxyzc5i
sxoula",
      "infrastructureType" : "exadata",
      "description" : "Flash Disk Failure has been detected. This is being
investigated by Oracle Exadata team and the disk will be
                       replaced if needed. No action needed from the
customer.",
      "eventName" : "HEALTH.INFRASTRUCTURE.CELL.FLASH_DISK_FAILURE",
      "FLASH_1_1" : "S2T7NA0HC01251  failed",
      "otto-ingestion-time" : "2021-12-17T19:14:43.205Z",
      "otto-send-EventService-time" : "2021-12-17T19:14:44.198Z"
    }
  },
  "eventID" : "30130ab4-42fa-4285-93a7-47e49522c698",
  "extensions" : {
    "compartmentId" :
"ocid1.tenancy.oc1..aaaaaaaao3lj36x6lwxyvc4wausjouca7pwyjfwb5ebsq5emrpqlql2gj5
iq"
  }
}
```

Exadata Cloud Infrastructure VM Cluster Event Types

Review the list of events that can be emitted by VM Cluster

Friendly Name	Event Type
Cloud VM Cluster - Change Compartment Begin	com.oraclecloud.databaseservice.changecloudvmclustercompartment.begin
Cloud VM Cluster - Change Compartment End	com.oraclecloud.databaseservice.changecloudvmclustercompartment.end
Cloud VM Cluster - Create Begin	com.oraclecloud.databaseservice.createcloudvmcluster.begin
Cloud VM Cluster - Create End	com.oraclecloud.databaseservice.createcloudvmcluster.end
Cloud VM Cluster - Delete Begin	com.oraclecloud.databaseservice.deletecloudvmcluster.begin
Cloud VM Cluster - Delete End	com.oraclecloud.databaseservice.deletecloudvmcluster.end
Cloud VM Cluster - Update Begin	com.oraclecloud.databaseservice.updatecloudvmcluster.begin
Cloud VM Cluster - Update End	com.oraclecloud.databaseservice.updatecloudvmcluster.end
Cloud VM Cluster - Update IORM Configuration Begin	com.oraclecloud.databaseservice.updatecloudvmclusteriormconfig.begin
Cloud VM Cluster - Update IORM Configuration End	com.oraclecloud.databaseservice.updatecloudvmclusteriormconfig.end
Cloud VM Cluster - Add Virtual Machine Begin	com.oraclecloud.databaseservice.cloudvmclusteraddvirtualmachine.begin
Cloud VM Cluster - Add Virtual Machine End	com.oraclecloud.databaseservice.cloudvmclusteraddvirtualmachine.end

NOT_SUPPORTED

This is a reference event for a cloud VM cluster resource:

```
{
    "cloudEventsVersion": "0.1",
    "eventID": "<unique_ID>",
    "eventType":
"com.oraclecloud.databaseservice.updatecloudvmclusteriormconfig.begin",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2022-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
      "eventGroupingId": "<unique_ID>",
      "eventName": "UpdateCloudVmClusterIormConfig",
      "compartmentName": "example_compartment",
      "resourceName": "my_container_database",
      "resourceId": "ocid1.cloudvmcluster.oc1.<unique_ID>",
      "resourceVersion": null,
      "additionalDetails": {
        "cloudExadataInfrastructureId":
```

```
"ocid1.cloudexadatainfrastructure.oc1.<unique_ID>",
        "freeFormTags": {},
        "definedTags": {},
        "licenseType": "BRING_YOUR_OWN_LICENSE",
        "lifecycleState": "AVAILABLE",
        "giVersion": "19.0.0.0.0",
        "cpuCoreCount": 16
      }
    }
  },
  "timeCreated": "2022-06-15T16:31:31.979Z"
}
```


This is a reference event for Add Virtual Machine Begin:


```
{
  "id":
"ocid1.eventschema.oc1.phx.n2p4ijm0jyuia5p6lzhps0axtqft2d2ueywaq4oxcr3ywlzt9jd
689kvxazo",
  "serviceName": "Database",
  "displayName": "Cloud VM Cluster - Add Virtual Machine Begin",
  "eventType":
"com.oraclecloud.databaseservice.cloudvmclusteraddvirtualmachine.begin",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2023-01-06T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "cloudExadataInfrastructureId",
      "type": [
        "null",
        "string"
      ]
    },
    {
```

```
      "name": "cpuCoreCount",
      "type": [
        "null",
        "Integer"
      ]
    },
    {
      "name": "ocpuCountFractional",
      "type": [
        "null",
        "Float"
      ]
    },
    {
      "name": "dataStorageSizeInTBs",
      "type": [
        "null",
        "Integer"
      ]
    },
    {
      "name": "dataStorageSizeInGBs",
      "type": [
        "null",
        "Integer"
      ]
    },
    {
      "name": "licenseType",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "giVersion",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "dbNodeIds",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "timeZone",
      "type": [
        "null",
        "string"
      ]
    }
  ],
```

```
  "exampleEvent": {
    "eventType":
"com.oraclecloud.databaseservice.cloudvmclusteraddvirtualmachine.begin",
    "cloudEventsVersion": "0.1",
    "eventTypeVersion": "2.0",
    "source": "databaseservice",
    "eventID": "bc78609a-783a-9034-ccd1-12ab908df913",
    "eventTime": "2023-01-06T23:18:04.000Z",
    "contentType": "application/json",
    "data": {
      "eventGroupingId": "csid201fe4f3443a853d76e9cec3ef4a/
3200918f142a44adb715d8aaf4f5ba99/DC62865A826A6E98699590E7F33C5064",
      "eventName": "CloudVmClusterAddVirtualMachine",
      "compartmentId": "ocid1.compartment.oc1.....unique_id",
      "compartmentName": null,
      "resourceName": "my_cloud_vm_cluster",
      "resourceId": "ocid1.cloudvmcluster.oc1.....unique_id",
      "resourceVersion": null,
      "availabilityDomain": "",
      "tagSlug": "tag_slug",
      "identity": {
        "principalName": null,
        "principalId": null,
        "authType": null,
        "callerName": null,
        "callerId": null,
        "tenantId": null,
        "ipAddress": null,
        "credentials": null,
        "authZPolicies": null,
        "userGroups": null,
        "userAgent": null,
        "consoleSessionId": null
      },
      "request": {
        "id": "01858321-0045-4bc5-b0d9-a917a6a40901",
        "path": null,
        "action": null,
        "parameters": null,
        "headers": null
      },
      "response": {
        "status": null,
        "responseTime": null,
        "headers": null,
        "payload": null,
        "message": null
      },
      "stateChange": {
        "previous": null,
        "current": {
          "licenseType": "BRING_YOUR_OWN_LICENSE",
          "dataStorageSizeGb": 60,
          "lifecycleState": "AVAILABLE",
          "sshPublicKeys": "...",
          "displayName": "my_cloud_vm_cluster",
```

```
          "cpuCoreCount": 16,
          "freeTags": {},
          "definedTags": {},
          "ocpuCountFractional": 16.0
        }
      },
      "additionalDetails": {
        "timeCreated": "2023-01-06T22:18:04.000Z",
        "timeUpdated": "2023-01-06T22:20:04.000Z",
        "lifecycleState": "AVAILABLE",
        "lifecycleDetails": null,
        "cloudExadataInfrastructureId":
"ocid1.cloudexadatainfrastructure.oc1.....unique_id",
        "cpuCoreCount": 16,
        "ocpuCountFractional": 16.0,
        "dataStorageSizeInTBs": 4,
        "dataStorageSizeInGBs": 60,
        "licenseType": "BRING_YOUR_OWN_LICENSE",
        "giVersion": "19.0.0.0.0",
        "dbNodeIds": "[ocid1.dbnode.oc1.....unique_id,...]",
        "timeZone": "UTC"
      },
      "internalDetails": {
        "attributes": null
      }
    }
  },
  "timeCreated": "2023-01-06T23:18:04.000Z"
}
```


This is a reference event for Add Virtual Machine End:


```
{
  "id":
"ocid1.eventschema.oc1.phx.v87pke1z9k9u6xaqo51taf6bunf0gc2wyhrbmjzbh3h1pjwakav
mf2borxgb",
  "serviceName": "Database",
  "displayName": "Cloud VM Cluster - Add Virtual Machine End",
  "eventType":
"com.oraclecloud.databaseservice.cloudvmclusteraddvirtualmachine.end",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2023-01-06T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
```

```
    "type": "string"
  },
  {
    "name": "lifecycleDetails",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "cloudExadataInfrastructureId",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "cpuCoreCount",
    "type": [
      "null",
      "Integer"
    ]
  },
  {
    "name": "ocpuCountFractional",
    "type": [
      "null",
      "Float"
    ]
  },
  {
    "name": "dataStorageSizeInTBs",
    "type": [
      "null",
      "Integer"
    ]
  },
  {
    "name": "dataStorageSizeInGBs",
    "type": [
      "null",
      "Integer"
    ]
  },
  {
    "name": "licenseType",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "giVersion",
    "type": [
      "null",
      "string"
```

```
        ]
      },
      {
        "name": "dbNodeIds",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "timeZone",
        "type": [
          "null",
          "string"
        ]
      }
    ],
    "exampleEvent": {
      "eventType":
"com.oraclecloud.databaseservice.cloudvmclusteraddvirtualmachine.end",
      "cloudEventsVersion": "0.1",
      "eventTypeVersion": "2.0",
      "source": "databaseservice",
      "eventID": "ced78bb7-3903-acd8-ff78-5567aa01a912",
      "eventTime": "2023-01-06T23:18:04.000Z",
      "contentType": "application/json",
      "data": {
        "eventGroupingId": "csid89a04ef74ccb8b48340f56e656cf/
729c99d3e5a34d548ddc31c054810454/634F086E8618E0A660946A6862C82A68",
        "eventName": "CloudVmClusterAddVirtualMachine",
        "compartmentId": "ocid1.compartment.oc1.....unique_id",
        "compartmentName": null,
        "resourceName": "my_cloud_vm_cluster",
        "resourceId": "ocid1.cloudvmcluster.oc1.....unique_id",
        "resourceVersion": null,
        "availabilityDomain": "",
        "tagSlug": "tag_slug",
        "identity": {
          "principalName": null,
          "principalId": null,
          "authType": null,
          "callerName": null,
          "callerId": null,
          "tenantId": null,
          "ipAddress": null,
          "credentials": null,
          "authZPolicies": null,
          "userGroups": null,
          "userAgent": null,
          "consoleSessionId": null
        },
        "request": {
          "id": "07197e12-b680-475e-851e-bb89fcd8376d",
          "path": null,
          "action": null,
          "parameters": null,
```

```
        "headers": null
      },
      "response": {
        "status": null,
        "responseTime": null,
        "headers": null,
        "payload": null,
        "message": null
      },
      "stateChange": {
        "previous": null,
        "current": {
          "licenseType": "BRING_YOUR_OWN_LICENSE",
          "dataStorageSizeGb": 60,
          "lifecycleState": "AVAILABLE",
          "sshPublicKeys": "...",
          "displayName": "my_cloud_vm_cluster",
          "cpuCoreCount": 16,
          "freeTags": {},
          "definedTags": {},
          "ocpuCountFractional": 16.0
        }
      },
      "additionalDetails": {
        "timeCreated": "2023-01-06T22:18:04.000Z",
        "timeUpdated": "2023-01-06T22:20:04.000Z",
        "lifecycleState": "AVAILABLE",
        "lifecycleDetails": null,
        "cloudExadataInfrastructureId":
"ocid1.cloudexadatainfrastructure.oc1.....unique_id",
        "cpuCoreCount": 16,
        "ocpuCountFractional": 16.0,
        "dataStorageSizeInTBs": 4,
        "dataStorageSizeInGBs": 60,
        "licenseType": "BRING_YOUR_OWN_LICENSE",
        "giVersion": "19.0.0.0.0",
        "dbNodeIds": "[ocid1.dbnode.oc1.....unique_id,...]",
        "timeZone": "UTC"
      },
      "internalDetails": {
        "attributes": null
      }
    }
  },
  "timeCreated": "2023-01-06T23:18:04.000Z"
}
```


This is a reference event for Cloud VM Cluster - Update Begin:


```
{
  "id":
"ocid1.eventschema.oc1.phx.ekmz1phzp4bl1k7m7tbygulbnakmjnrsi99eqjops3zvpt337pn
nfmj6r79j",
  "serviceName": "Database",
  "displayName": "Cloud VM Cluster - Update Begin",
```

```
"eventType": "com.oraclecloud.databaseservice.updatecloudvmcluster.begin",
"source": "databaseservice",
"eventTypeVersion": "2.0",
"eventTime": "2019-06-27T21:16:04.000Z",
"contentType": "application/json",
"additionalDetails": [
  {
    "name": "id",
    "type": "string"
  },
  {
    "name": "defineTags",
    "type": [
      "null",
      "Map<String, Map<String, Object>>"
    ]
  },
  {
    "name": "freeFormTags",
    "type": [
      "null",
      "Map<String, String>"
    ]
  },
  {
    "name": "timeCreated",
    "type": "string"
  },
  {
    "name": "timeUpdated",
    "type": "string"
  },
  {
    "name": "lifecycleState",
    "type": "string"
  },
  {
    "name": "lifecycleDetails",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "cloudExadataInfrastructureId",
    "type": "string"
  },
  {
    "name": "cpuCoreCount",
    "type": [
      "null",
      "Integer"
    ]
  },
  {
    "name": "dataStorageSizeInGBs",
```

```
        "type": [
          "null",
          "Integer"
        ]
      },
      {
        "name": "licenseType",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "giVersion",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "dbNodeIds",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "timeZone",
        "type": [
          "null",
          "string"
        ]
      }
    ],
    "exampleEvent": {
      "cloudEventsVersion": "0.1",
      "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
      "eventType": "com.oraclecloud.databaseservice.updatecloudvmcluster.begin",
      "source": "databaseservice",
      "eventTypeVersion": "2.0",
      "eventTime": "2019-06-27T21:16:04.000Z",
      "contentType": "application/json",
      "data": {
        "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
        "eventName": "UpdateCloudVmCluster",
        "compartmentName": "example_compartment",
        "resourceName": "my_container_database",
        "resourceId": "ocid1.cloudvmcluster.oc1.....unique_id",
        "resourceVersion": null,
        "additionalDetails": {
          "cloudExadataInfrastructureId":
"ocid1.cloudexadatainfrastructure.oc1.....unique_id",
        "freeFormTags": {},
        "definedTags": {},
        "licenseType": "BRING_YOUR_OWN_LICENSE",
        "lifecycleState": "AVAILABLE",
```

```
        "giVersion": "19.0.0.0.0",
        "cpuCoreCount": 16
      }
    }
  },
  "timeCreated": "2020-06-15T16:31:31.979Z"
}
```


This is a reference event for Cloud VM Cluster - Update End:


```
{
  "id":
"ocid1.eventschema.oc1.phx.svwkildsx63clp1q6phba7d6lns1rl92yc3uyc2ea5utjprqcwu
hbgvht4we",
  "serviceName": "Database",
  "displayName": "Cloud VM Cluster - Update End",
  "eventType": "com.oraclecloud.databaseservice.updatecloudvmcluster.end",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "id",
      "type": "string"
    },
    {
      "name": "defineTags",
      "type": [
        "null",
        "Map<String, Map<String, Object>>"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "Map<String, String>"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",
      "type": [
```

```
        "null",
        "string"
      ]
    },
    {
      "name": "cloudExadataInfrastructureId",
      "type": "string"
    },
    {
      "name": "cpuCoreCount",
      "type": [
        "null",
        "Integer"
      ]
    },
    {
      "name": "dataStorageSizeInGBs",
      "type": [
        "null",
        "Integer"
      ]
    },
    {
      "name": "licenseType",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "giVersion",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "dbNodeIds",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "timeZone",
      "type": [
        "null",
        "string"
      ]
    }
  ],
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
    "eventType": "com.oraclecloud.databaseservice.updatecloudvmcluster.end",
    "source": "databaseservice",
```

```
    "eventTypeVersion": "2.0",
    "eventTime": "2019-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
      "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
      "eventName": "UpdateCloudVmCluster",
      "compartmentName": "example_compartment",
      "resourceName": "my_container_database",
      "resourceId": "ocid1.cloudvmcluster.oc1.....unique_id",
      "resourceVersion": null,
      "additionalDetails": {
        "cloudExadataInfrastructureId":
"ocid1.cloudexadatainfrastructure.oc1.....unique_id",
        "freeFormTags": {},
        "definedTags": {},
        "licenseType": "BRING_YOUR_OWN_LICENSE",
        "lifecycleState": "AVAILABLE",
        "giVersion": "19.0.0.0.0",
        "cpuCoreCount": 16
      }
    }
  },
  "timeCreated": "2020-06-15T16:31:31.979Z"
}
```


VM Node Subsetting Event Types

Review the list of event types that VM Node Subsetting emits.

Table 6-1 VM Node Subsetting Events

Friendly Name	Event Type
VM Cluster - Add Virtual Machine Begin	com.oraclecloud.databaseservice.vmclusteraddvirtualmachine.begin
VM Cluster - Add Virtual Machine End	com.oraclecloud.databaseservice.vmclusteraddvirtualmachine.end
VM Cluster - Terminate Virtual Machine Begin	com.oraclecloud.databaseservice.vmclusterterminatevirtualmachine.begin
VM Cluster - Terminate Virtual Machine End	com.oraclecloud.databaseservice.vmclusterterminatevirtualmachine.end

Example 6-61 VM Node Subsetting Examples

This is a reference event for VM Cluster - Add Virtual Machine Begin:


```
"exampleEvent": {
"cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.vmclusteraddvirtualmachine.begin",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
```

```
  "contentType": "application/json",
  "extensions": {
"compartmentId": "ocid1.compartment.oc1..unique_ID"
  },
  "data": {
"compartmentId": "ocid1.compartment.oc1..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_database",
    "resourceId": "Vmcluster-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
"id": "ocid1.id..oc1...unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2019-09-03T12:00:00.000Z",
      "timeUpdated": "2019-09-03T12:30:00.000Z",
      "displayName": "testDisplayName",
      "lifecycleDetails": "detail message",
      "exadataInfrastructureId": "ExatraInfra-unique_ID",
      "vmClusterNetworkId": "VmCluster-unique_ID",
      "cpuCoreCount": 2,
      "dataStorageSizeInTBs": 4,
      "memorySizeInGBs": 30,
      "dbNodeStorageSizeInGBs": 60,
      "dbVersion": "19.0.0.0",
      "licenseType": "BRING_YOUR_OWN_LICENSE",
      "giVersion": "19.0.0.0",
      "dbNodeIds": "[ocid1.dbnode.1, ocid1.dbnode.2,...]",
      "dbServerIds": "[ocid1.dbserver.1, ocid1.dbserver.2,...]",
      "timeZone": "US/Pacific"
    }
  }
}
```


This is a reference event for VM Cluster - Add Virtual Machine End:


```
"exampleEvent": {
"cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.vmclusteraddvirtualmachine.end",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
"compartmentId": "ocid1.compartment.oc1..unique_ID"
  },
  "data": {
"compartmentId": "ocid1.compartment.oc1..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_database",
    "resourceId": "Vmcluster-unique_ID",
    "availabilityDomain": "all",
```

```
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
"id": "ocid1.id..oc1...unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2019-09-03T12:00:00.000Z",
      "timeUpdated": "2019-09-03T12:30:00.000Z",
      "displayName": "testDisplayName",
      "lifecycleDetails": "detail message",
      "exadataInfrastructureId": "ExatraInfra-unique_ID",
      "vmClusterNetworkId": "VmCluster-unique_ID",
      "cpuCoreCount": 2,
      "dataStorageSizeInTBs": 4,
      "memorySizeInGBs": 30,
      "dbNodeStorageSizeInGBs": 60,
      "dbVersion": "19.0.0.0",
      "licenseType": "BRING_YOUR_OWN_LICENSE",
      "giVersion": "19.0.0.0",
      "dbNodeIds": "[ocid1.dbnode.1, ocid1.dbnode.2,...]",
      "dbServerIds": "[ocid1.dbserver.1, ocid1.dbserver.2,...]",
      "timeZone": "US/Pacific"
    }
  }
}
```


This is a reference event for VM Cluster - Terminate Virtual Machine Begin:


```
"exampleEvent": {
"cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.vmclusterterminatevirtualmachine.begin",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
"compartmentId": "ocid1.compartment.oc1..unique_ID"
  },
  "data": {
"compartmentId": "ocid1.compartment.oc1..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_database",
    "resourceId": "Vmcluster-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
"id": "ocid1.id..oc1...unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2019-09-03T12:00:00.000Z",
      "timeUpdated": "2019-09-03T12:30:00.000Z",
      "displayName": "testDisplayName",
      "lifecycleDetails": "detail message",
      "exadataInfrastructureId": "ExatraInfra-unique_ID",
```

```
      "vmClusterNetworkId": "VmCluster-unique_ID",
      "cpuCoreCount": 2,
      "dataStorageSizeInTBs": 4,
      "memorySizeInGBs": 30,
      "dbNodeStorageSizeInGBs": 60,
      "dbVersion": "19.0.0.0",
      "licenseType": "BRING_YOUR_OWN_LICENSE",
      "giVersion": "19.0.0.0",
      "dbNodeIds": "[ocid1.dbnode.1, ocid1.dbnode.2,...]",
      "dbServerIds": "[ocid1.dbserver.1, ocid1.dbserver.2,...]",
      "timeZone": "US/Pacific"
    }
  }
}
```

This is a reference event for VM Cluster - Terminate Virtual Machine End:

```
"exampleEvent": {
"cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.vmclusterterminatevirtualmachine.end",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
"compartmentId": "ocid1.compartment.oc1..unique_ID"
  },
  "data": {
"compartmentId": "ocid1.compartment.oc1..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_database",
    "resourceId": "Vmcluster-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
"id": "ocid1.id..oc1...unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2019-09-03T12:00:00.000Z",
      "timeUpdated": "2019-09-03T12:30:00.000Z",
      "displayName": "testDisplayName",
      "lifecycleDetails": "detail message",
      "exadataInfrastructureId": "ExatraInfra-unique_ID",
      "vmClusterNetworkId": "VmCluster-unique_ID",
      "cpuCoreCount": 2,
      "dataStorageSizeInTBs": 4,
      "memorySizeInGBs": 30,
      "dbNodeStorageSizeInGBs": 60,
      "dbVersion": "19.0.0.0",
      "licenseType": "BRING_YOUR_OWN_LICENSE",
      "giVersion": "19.0.0.0",
      "dbNodeIds": "[ocid1.dbnode.1, ocid1.dbnode.2,...]",
      "dbServerIds": "[ocid1.dbserver.1, ocid1.dbserver.2,...]",
```

```
      "timeZone": "US/Pacific"
    }
  }
}
```

Data Guard Association Event Types

Review the list of event types that Data Guard associations emit.

Friendly Name	Event Type
Change Protection Mode Begin	com.oraclecloud.databaseservice.changeprotectionmode.begin
Change Protection Mode End	com.oraclecloud.databaseservice.changeprotectionmode.end
Data Guard Association - Create Begin	com.oraclecloud.databaseservice.createdataguardassociation.begin
Data Guard Association - Create End	com.oraclecloud.databaseservice.createdataguardassociation.end
Data Guard Association - Failover Begin	com.oraclecloud.databaseservice.failoverdataguardassociation.begin
Data Guard Association - Failover End	com.oraclecloud.databaseservice.failoverdataguardassociation.end
Data Guard Association - Reinstate Begin	com.oraclecloud.databaseservice.reinstatedataguardassociation.begin
Data Guard Association - Reinstate End	com.oraclecloud.databaseservice.reinstatedataguardassociation.end
Data Guard Association - Switchover Begin	com.oraclecloud.databaseservice.switchoverdataguardassociation.begin
Data Guard Association - Switchover End	com.oraclecloud.databaseservice.switchoverdataguardassociation.end

NOT_SUPPORTED

This is a reference event for Data Guard associations:

```
{
    "cloudEventsVersion": "0.1",
    "contentType": "application/json",
    "data": {
        "additionalDetails": {
            "ApplyLag": null,
            "DGConfigId": "7e8eff2b-a4cd-474a-abd5-940b05c0b1fd",
            "DGConfigState": "null",
            "DatabaseId": "ocid1.database.oc1.iad.<unique_ID>",
            "DbHomeId": "ocid1.dbhome.oc1.iad.<unique_ID>",
            "DbSystemId": "ocid1.dbsystem.oc1.iad.<unique_ID>",
            "LastSyncedTime": null,
            "SyncState": "null",
            "dcsDgUpdateTimestamp": null,
            "lastUpdatedIdentifier": null,
            "lifeCycleMessage": null,
            "lifecycleState": "PROVISIONING",
```

```
            "timeCreated": "2019-10-25T21:42:19.041Z",
            "timeUpdated": "2019-10-25T21:42:19.041Z"
        },
        "availabilityDomain": "XXIT:US-ASHBURN-AD-1",
        "compartmentId": "ocid1.compartment.oc1.<unique_ID>",
        "compartmentName": "example_compartment",
        "resourceId": "ocid1.dgassociation.oc1.iad.<unique_ID>"
    },
    "eventID": "5b8b7fbf-2e9a-4730-9761-e52715b7bc79",
    "eventTime": "2019-10-25T21:42:16.579Z",
    "eventType":
"com.oraclecloud.databaseservice.createdataguardassociation.begin",
    "eventTypeVersion": "2.0",
    "extensions": {
        "compartmentId": "ocid1.compartment.oc1.<unique_ID>"
    },
    "source": "DatabaseService"
}
```


Oracle Database Home Event Types

Review the list of events emitted by Oracle Database Homes.

Friendly Name	Event Type
DB Home - Create Begin	com.oraclecloud.databaseservice.createdbhome.begin
DB Home - Create End	com.oraclecloud.databaseservice.createdbhome.end
DB Home - Patch Begin	com.oraclecloud.databaseservice.patchdbhome.begin
DB Home - Patch End	com.oraclecloud.databaseservice.patchdbhome.end
DB Home - Terminate Begin	com.oraclecloud.databaseservice.deletedbhome.begin
DB Home - Terminate End	com.oraclecloud.databaseservice.deletedbhome.end
DB Home - Update Begin	com.oraclecloud.databaseservice.updatedbhome.begin
DB Home - Update End	com.oraclecloud.databaseservice.updatedbhome.end

NOT_SUPPORTED

This is a reference event for Database Homes:


```
{
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.createdbhome.begin",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-08-29T21:16:04Z",
    "contentType": "application/json",
```

```
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1.<unique_ID>"
  },
  "data": {
    "compartmentId": "ocid1.compartment.oc1.<unique_ID>",
    "compartmentName": "example_compartment",
    "resourceName": "my_dbhome",
    "resourceId": "DbHome-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.id.oc1.<unique_ID>",
      "lifecycleState": "PROVISIONING",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbSystemId": "DbSystem-unique_ID",
      "dbVersion": "19.0.0.0",
      "recordVersion": 4,
      "displayName": "example_display_name"
    }
  }
}
```

Database Event Types

These are the event types that Oracle Databases in Exadata Cloud Service instances emit.

Friendly Name	Event Type
Database - Automatic Backup Begin	`com.oraclecloud.databaseservice.automaticbackupdatabase.begin`
Database - Automatic Backup End	`com.oraclecloud.databaseservice.automaticbackupdatabase.end`
Database - Create Backup Begin	`com.oraclecloud.databaseservice.backupdatabase.begin`
Database - Create Backup End	`com.oraclecloud.databaseservice.backupdatabase.end`
Database - Critical (see Database Service Event Types for more information)	`com.oraclecloud.databaseservice.database.critical`
Database - Information	`com.oraclecloud.databaseservice.database.information`
Database - Delete Backup Begin	`com.oraclecloud.databaseservice.deletebackup.begin`
Database - Delete Backup End	`com.oraclecloud.databaseservice.deletebackup.end`
Database - Migrate to KMS Key Begin	`com.oraclecloud.databaseservice.migratedatabasekmskey.begin`
Database - Migrate to KMS Key End	`com.oraclecloud.databaseservice.migratedatabasekmskey.end`

Friendly Name	Event Type
Database - Move Begin	com.oraclecloud.databaseservice.movedatabase.begin
Database - Move End	com.oraclecloud.databaseservice.movedatabase.end
Database - Restore Begin	com.oraclecloud.databaseservice.restoredatabase.begin
Database - Restore End	com.oraclecloud.databaseservice.restoredatabase.end
Database - Rotate KMS Key Begin	com.oraclecloud.databaseservice.rotatedatabasekmskey.begin
Database - Rotate KMS Key End	com.oraclecloud.databaseservice.rotatedatabasekmskey.end
Database - Terminate Begin	com.oraclecloud.databaseservice.database.terminate.begin
Database - Terminate End	com.oraclecloud.databaseservice.database.terminate.end
Database - Update Begin	com.oraclecloud.databaseservice.updatedatabase.begin
Database - Update End	com.oraclecloud.databaseservice.updatedatabase.end
Database - Upgrade Begin	com.oraclecloud.databaseservice.upgradedatabase.begin
Database - Upgrade End	com.oraclecloud.databaseservice.upgradedatabase.end

NOT_SUPPORTED

This is a reference event for databases:

```
{
"eventType" : "com.oraclecloud.databaseservice.backupdatabase.begin",
udEventsVersion" : "0.1",
"eventTypeVersion" : "2.0",
"source" : "DatabaseService",
"eventTime" : "2020-01-08T17:31:43.666Z",
"contentType" : "application/json",
"data" : {
"compartmentId" : "ocid1.compartment.oc1.<unique_ID>",
"compartmentName": "example_compartment_name",
"resourceName": "my_backup",
"resourceId": "ocid1.dbbckup.oc1.<unique_ID>",
"availabilityDomain": "<availability_domain>",
"additionalDetails" : {
"timeCreated" : "2020-01-08T17:31:44Z",
"lifecycleState" : "CREATING",
"dbSystemId" : "ocid1.dbsystem.oc1.<unique_ID>",
"dbHomeId" : ocid1.dbhome.oc1.<unique_ID>",
"dbUniqueName" : DB1115_iad1dv",
"dbVersion" : "11.2.0.4.190716",
"databaseEdition" : "ENTERPRISE_EDITION_HIGH_PERFORMANCE",
"autoBackupsEnabled" : "false",
```

```
"backupType" : "FULL",
"databaseId" : "ocid1.database.oc1.<unique_ID>",
},
"definedTags" : {
      "My_example_tag_name" :
        { "Example_key" : "Example_value" }
      },
    "eventID": "<unique_ID>",
    "extensions" : {
      "compartmentId": "ocid1.compartment.oc1.<unique_ID>"
    }
}
```

Pluggable Database Event Types

These are the event types that Oracle pluggable databases in Oracle Cloud Infrastructure emit.

Friendly Name	Event Type
Pluggable Database - Create Begin	com.oraclecloud.databaseservice.createpluggabledatabase.begin
Pluggable Database - Create End	com.oraclecloud.databaseservice.createpluggabledatabase.end
Pluggable Database - Delete Begin	com.oraclecloud.databaseservice.deletepluggabledatabase.begin
Pluggable Database - Delete End	com.oraclecloud.databaseservice.deletepluggabledatabase.end
Pluggable Database - Local Clone Begin	com.oraclecloud.databaseservice.localclonepluggabledatabase.begin
Pluggable Database - Local Clone End	com.oraclecloud.databaseservice.localclonepluggabledatabase.end
Pluggable Database - Remote Clone Begin	com.oraclecloud.databaseservice.remoteclonepluggabledatabase.begin
Pluggable Database - Remote Clone End	com.oraclecloud.databaseservice.remoteclonepluggabledatabase.end
Start Pluggable Database - Begin	com.oraclecloud.databaseservice.startpluggabledatabase.begin
Start Pluggable Database - End	com.oraclecloud.databaseservice.startpluggabledatabase.end
Stop Pluggable Database - Begin	com.oraclecloud.databaseservice.stoppluggabledatabase.begin
Stop Pluggable Database - End	com.oraclecloud.databaseservice.stoppluggabledatabase.end
Pluggable Database - Convert to Regular Begin	com.oraclecloud.databaseservice.pluggabledatabase.converttoregular.begin
Pluggable Database - Convert to Regular End	com.oraclecloud.databaseservice.pluggabledatabase.converttoregular.end
Pluggable Database - Inplace Restore Begin	com.oraclecloud.databaseservice.pluggabledatabase.inplacerestore.begin
Pluggable Database - Inplace Restore End	com.oraclecloud.databaseservice.pluggabledatabase.inplacerestore.end

Friendly Name	Event Type
Pluggable Database - Refresh Begin	com.oraclecloud.databaseservice.pluggabledatabase.refresh.begin
Pluggable Database - Refresh End	com.oraclecloud.databaseservice.pluggabledatabase.refresh.end
Pluggable Database - Relocate Begin	com.oraclecloud.databaseservice.pluggabledatabase.relocate.begin
Pluggable Database - Relocate End	com.oraclecloud.databaseservice.pluggabledatabase.relocate.end

NOT_SUPPORTED

This is a reference event for pluggable databases (PDBs):

```
{
  "eventID": "unique_id",
  "eventTime": "2021-03-23T00:49:14.123Z",
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1.<unique_ID>"
  },
  "eventType":
"com.oraclecloud.databaseservice.remoteclonepluggabledatabase.begin",
  "eventTypeVersion": "2.0",
  "cloudEventsVersion": "0.1",
  "source": "databaseservice",
  "contentType": "application/json",
  "definedTags": {},
  "data": {
    "compartmentId": "ocid1.compartment.oc1.<unique_ID>",
    "compartmentName": "MyCompartment",
    "resourceName": "11092020_PKS_PDB1",
    "resourceId": "ocid1.pluggabledatabases.oc1.phx.<unique_ID>",
    "availabilityDomain": "XXIT:PHX-AD-1",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.pluggabledatabases.oc1.phx.<unique_ID>",
      "timeCreated": "2021-03-13T21:15:59.000Z",
      "timeUpdated": "2021-03-13T21:15:59.000Z",
      "databaseId": "ocid1.database.oc1.<unique_ID>",
      "lifecycleState": "AVAILABLE",
      "lifecycleDetails": "Pluggable Database is available",
      "displayName": "Pluggable Database - Remote Clone Begin"
    }
  }
}
```

This is a reference event for Pluggable Database - Convert to Regular Begin:

```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
```

```
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.converttoregular.begin",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020_PKS_PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique_id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique_id",
        "isRefreshableClone": true,
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "databaseId": "ocid1.database.oc1.....unique_id",
        "lifecycleState": "UPDATING",
        "displayName": "Pluggable Database - Convert to Regular Begin"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```


This is a reference event for Pluggable Database - Convert to Regular End:


```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.converttoregular.end",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020_PKS_PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique_id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
```

```
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique_id",
        "isRefreshableClone": false,
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "databaseId": "ocid1.database.oc1.....unique_id",
        "lifecycleState": "AVAILABLE",
        "displayName": "Pluggable Database - Convert to Regular End"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```

This is a reference event for Pluggable Database - Inplace Restore Begin:

```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.inplacerestore.begin",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020_PKS_PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique_id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique_id",
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "databaseId": "ocid1.database.oc1.....unique_id",
        "lifecycleState": "RESTORE_IN_PROGRESS",
        "isRefreshableClone": false,
        "displayName": "Pluggable Database - Inplace Restore Begin"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```

This is a reference event for Pluggable Database - Inplace Restore End:


```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.inplacerestore.end",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020_PKS_PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique_id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique_id",
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "databaseId": "ocid1.database.oc1.....unique_id",
        "lifecycleState": "AVAILABLE",
        "isRefreshableClone": false,
        "lifecycleDetails": "Pluggable Database is available",
        "displayName": "Pluggable Database - Inplace Restore End"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```


This is a reference event for Pluggable Database - Refresh Begin:


```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.refresh.begin",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
```

```
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020_PKS_PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique_id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique_id",
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "isRefreshableClone": true,
        "databaseId": "ocid1.database.oc1.....unique_id",
        "lifecycleState": "AVAILABLE",
        "lifecycleDetails": "Pluggable Database is available",
        "displayName": "Pluggable Database - Refresh Begin"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```


This is a reference event for Pluggable Database - Refresh End:


```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.refresh.end",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020_PKS_PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique_id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique_id",
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "databaseId": "ocid1.database.oc1.....unique_id",
        "lifecycleState": "AVAILABLE",
        "isRefreshableClone": true,
        "lifecycleDetails": "Pluggable Database is available",
        "displayName": "Pluggable Database - Refresh End"
```

```
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```

This is a reference event for Pluggable Database - Relocate Begin:

```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.relocate.begin",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020_PKS_PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique_id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique_id",
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "databaseId": "ocid1.database.oc1.....unique_id",
        "lifecycleState": "AVAILABLE",
        "isRefreshableClone": false,
        "lifecycleDetails": "Pluggable Database is available",
        "displayName": "Pluggable Database - Relocate Begin"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```

This is a reference event for Pluggable Database - Relocate End:

```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
```

```
"com.oraclecloud.databaseservice.pluggabledatabase.relocate.end",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020_PKS_PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique_id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique_id",
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "databaseId": "ocid1.database.oc1.....unique_id",
        "lifecycleState": "AVAILABLE",
        "lifecycleDetails": "Pluggable Database is available",
        "displayName": "Pluggable Database - Relocate End"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```


Database Service Events

The Database Service emits events, which are structured messages that indicate changes in resources.

- Overview of Database Service Events
- Receive Notifications about Database Service Events
  Subscribe to the Database Service Events and get notified.
- Database Service Event Types
  Review the list of event types that the Database Service emits.
- Temporarily Restrict Automatic Diagnostic Collections for Specific Events
  Use the `tfactl blackout` command to temporarily suppress automatic diagnostic collections.
- Remediation
  These topics cover some common issues you might run into and how to address them.

Overview of Database Service Events

Database Service Events feature implementation enables you to get notified about health issues with your Oracle Databases or other components on the Guest VM.

It is possible that Oracle Database or Clusterware may not be healthy or various system components may be running out of space in the Guest VM. You are not notified of this situation, unless you opt-in.

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

Database Service Events feature implementation generates events for Guest VM operations and conditions, as well as Notifications for customers by leveraging the existing OCI Events service and Notification mechanisms in their tenancy. Customers can then create topics and subscribe to these topics through email, functions, or streams.

> **Note:**
>
> Events flow on Exadata Cloud Infrastructure depends on the following components: Oracle Trace File Analyzer (TFA), sysLens, and Oracle Database Cloud Service (DBCS) agent. Ensure that these components are up and running.

Manage Oracle Trace File Analyzer

- To check the run status of Oracle Trace File Analyzer, run the `tfactl status` command as `root` or a non-root user:

```
# tfactl status
.-------------------------------------------------------------------------------
----------------------.
| Host    | Status of TFA | PID    | Port | Version    | Build ID
| Inventory Status|
+---------------+--------------+--------+------+------------
+----------------------+------------+
| node1      | RUNNING    | 41312  | 5000 | 22.1.0.0.0 |
22100020220310214615 | COMPLETE        |
| node2      | RUNNING    | 272300 | 5000 | 22.1.0.0.0 |
22100020220310214615 | COMPLETE        |
'---------------+--------------+--------+------+------------
+----------------------+------------'
```

- To start the Oracle Trace File Analyzer daemon on the local node, run the `tfactl start` command as `root`:

```
# tfactl start
Starting TFA..
Waiting up to 100 seconds for TFA to be started..
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
Successfully started TFA Process..
```

```
. . . . .
TFA Started and listening for commands
```

- To stop the Oracle Trace File Analyzer daemon on the local node, run the `tfactl stop` command as `root`:

```
# tfactl stop
Stopping TFA from the Command Line
Nothing to do !
Please wait while TFA stops
Please wait while TFA stops
TFA-00002 Oracle Trace File Analyzer (TFA) is not running
TFA Stopped Successfully
Successfully stopped TFA..
```

Manage sysLens

- If sysLens is running, then once every 15 minutes data is collected in the local domU to discover the events to be reported. To check if sysLens is running, run the `systemctl status syslens` command as `root` in the domU:

```
# systemctl status syslens
\u25cf syslens.service
Loaded: loaded (/etc/systemd/system/syslens.service; disabled; vendor
preset: disabled)
Active: active (running) since Wed 2022-03-16 18:08:59 UTC; 34s ago
Main PID: 358039 (python3)
Memory: 31.6M
CGroup: /system.slice/syslens.service
\u2514\u2500358039 /usr/bin/python3 /var/opt/oracle/syslens/bin/
syslens_main.py --archive /var/opt/oracle/log/...

Mar 16 18:08:59 node1 systemd[1]: Started syslens.service.
Mar 16 18:09:09 node1 su[360495]: (to oracle) root on none
Mar 16 18:09:09 node1 su[360539]: (to grid) root on none
Mar 16 18:09:10 node1 su[360611]: (to grid) root on none
Mar 16 18:09:11 node1 su[360653]: (to oracle) root on none
```

- If the sysLens is enabled, when there is a reboot of the domU, then sysLens starts automatically. To validate if sysLens is enabled to collect telemetry, run the `systemctl is-enabled syslens` command as `root` in the domU:

```
# systemctl is-enabled syslens
enabled
```

- To validate if sysLens is configured to notify events, run the `/usr/bin/syslens --config /var/opt/oracle/syslens/data/exacc.syslens.config --get-key enable_telemetry` command as `root` in the domU:

```
# /usr/bin/syslens --config /var/opt/oracle/syslens/data/
exacc.syslens.config --get-key enable_telemetry
syslens Collection 2.3.3
on
```

Manage Database Service Agent

View the `/opt/oracle/dcs/log/dcs-agent.log` file to identify issues with the agent.

- To check the status of the Database Service Agent, run the `systemctl status` command:

```
# systemctl status dbcsagent.service
dbcsagent.service
Loaded: loaded (/usr/lib/systemd/system/dbcsagent.service; enabled; vendor
preset: disabled)
Active: active (running) since Fri 2022-04-01 13:40:19 UTC; 6min ago
Process: 9603 ExecStopPost=/bin/bash -c kill `ps -fu opc |grep "java.*dbcs-
agent.*jar" |awk '{print $2}' ` (code=exited, status=0/SUCCESS)
Main PID: 10055 (sudo)
CGroup: /system.slice/dbcsagent.service
 10055 sudo -u opc /bin/bash -c umask 077; /bin/java -
Doracle.security.jps.config=/opt/oracle/...
```

- To start the agent if it is not running, run the `systemctl start` command as the `root` user:

```
systemctl start dbcsagent.service
```

Related Topics

- To create a cloud VM cluster resource
  Create a VM cluster in an Exadata Cloud Infrastructure instance.
- Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.
- Overview of Events
- Notifications Overview

Receive Notifications about Database Service Events

Subscribe to the Database Service Events and get notified.

To receive notifications, subscribe to Database Service Events and get notified using the Oracle Notification service, see *Notifications Overview*. For more information about Oracle Cloud Infrastructure Events, see *Overview of Events*.

Events Service - Event Types:

- Database - Critical
- DB Node - Critical
- DB Node - Error
- DB Node - Warning
- DB Node - Information
- DB System - Critical

Related Topics

- Overview of Events
- Notifications Overview

Database Service Event Types

Review the list of event types that the Database Service emits.

> **Note:**
>
> - Critical events are triggered due to several types of critical conditions and errors that cause disruption to the database and other critical components. For example, database hang errors, and availability errors for databases, database nodes, and database systems to let you know if a resource becomes unavailable.
> - Information events are triggered when the database and other critical components work as expected. For example, a clean shutdown of CRS, CDB, client, or scan listener, or a startup of these components will create an event with the severity of INFORMATION.
> - Threshold limits reduce the number of notifications customers will receive for similar incident events whilst at the same time ensuring they receive the incident events and are reminded in a timely fashion.

Table 6-2 Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
Resource Utilization - Disk Usage	`HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE` This event is reported when VM guest file system free space falls below 10% free, as determined by the operating system `df(1)` command, for the following file systems: • `/` • `/u01` • `/u02` • `/var` (X8M and later only) • `/tmp` (X8M and later only)	HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE	`com.oraclecloud.databaseservice.dbnode.critical`	Critical threshold: 90%

Table 6-2 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
CRS status Up/ Down	`AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN.` An event of type CRITICAL is created when the Cluster Ready Service (CRS) is detected to be down.	AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN	`com.oraclecloud.databaseservice.dbnode.critical` (if .DOWN and NOT "user_action")	N/A
	`AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN_CLEARED` An event of type INFORMATION is created once it is determined that the event for CRS down has cleared.	N/A	`com.oraclecloud.databaseservice.dbnode.information` (if .DOWN_CLEARED)	

Table 6-2 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
SCAN Listener Up/ Down	`AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN` A DOWN event is created when a SCAN listener goes down. The event is of type INFORMATION when a SCAN listener is shutdown due to user action, such as with the Server Control Utility (`srvctl`) or Listener Control (`lsnrctl`) commands, or any Oracle Cloud maintenance action that uses those commands, such as performing a grid infrastructure software update. The event is of type CRITICAL when a SCAN listener goes down unexpectedly. A corresponding DOWN_CLEARED event is created when a SCAN listener is started. There are three SCAN listeners per cluster called LISTENER_SCAN[1,2,3].	AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN	`com.oraclecloud.databaseservice.dbnode.critical` (if .DOWN and NOT "user_action")	N/A
	`AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN_CLEARED` An event of type INFORMATION is created once it is determined that the event for SCAN Listener down has cleared.	N/A	`com.oraclecloud.databaseservice.dbnode.information` (if .DOWN_CLEARED)	

Table 6-2 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
Net Listener Up/ Down	`AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN` A DOWN event is created when a client listener goes down. The event is of type INFORMATION when a client listener is shutdown due to user action, such as with the Server Control Utility (`srvctl`) or Listener Control (`lsnrctl`) commands, or any Oracle Cloud maintenance action that uses those commands, such as performing a grid infrastructure software update. The event is of type CRITICAL when a client listener goes down unexpectedly. A corresponding DOWN_CLEARED event is created when a client listener is started. There is one client listener per node, each called LISTENER.	AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN	`com.oraclecloud.databaseservice.database.critical` (if .DOWN and NOT "user_action")	N/A
	`AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN_CLEARED` An event of type INFORMATION is created once it is determined that the event for Client Listener down has cleared.	N/A	`com.oraclecloud.databaseservice.database.information` (if .DOWN_CLEARED)	

Table 6-2 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
CDB Up/Down	`AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN` A DOWN event is created when a database instance goes down. The event is of type INFORMATION when a database instance is shutdown due to user action, such as with the SQL*Plus (`sqlplus`) or Server Control Utility (`srvctl`) commands, or any Oracle Cloud maintenance action that uses those commands, such as performing a database home software update. The event is of type CRITICAL when a database instance goes down unexpectedly. A corresponding DOWN_CLEARED event is created when a database instance is started.	AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN	`com.oraclecloud.databaseservice.database.critical` (if .DOWN and NOT "user_action")	N/A
	`AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN_CLEARED` An event of type INFORMATION is created once it is determined that the event for the CDB down has cleared.	N/A	`com.oraclecloud.databaseservice.database.information` (if .DOWN_CLEARED)	

Table 6-2 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
CRS Eviction	AVAILABILITY.DB_GUEST.CRS_INSTANCE.EVICTION An event of type CRITICAL is created when the Cluster Ready Service (CRS) evicts a node from the cluster. The CRS alert.log is parsed for the CRS-1632 error indicating that a node is being removed from the cluster.	AVAILABILITY.DB_GUEST.CRS_INSTANCE.EVICTION	An event of type CRITICAL is created when the Cluster Ready Service (CRS) evicts a node from the cluster. The CRS alert.log is parsed for the CRS-1632 error indicating that a node is being removed from the cluster.	N/A
Critical DB Errors	HEALTH.DB_CLUSTER.CDB.CORRUPTION Database corruption has been detected on your primary or standby database. The database alert.log is parsed for any specific errors that are indicative of physical block corruptions, logical block corruptions, or logical block corruptions caused by lost writes.	HEALTH.DB_CLUSTER.CDB.CORRUPTION	com.oraclecloud.databaseservice.database.critical	N/A
Other DB Errors	HEALTH.DB_CLUSTER.CDB.ARCHIVER_HANG An event of type CRITICAL is created if a CDB is either unable to archive the active online redo log or unable to archive the active online redo log fast enough to the log archive destinations.	HEALTH.DB_CLUSTER.CDB.ARCHIVER_HANG	com.oraclecloud.databaseservice.database.critical	N/A

Table 6-2 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
	HEALTH.DB_CLUSTER.CDB.DATABASE_HANG An event of type CRITICAL is created when a process/session hang is detected in the CDB.	HEALTH.DB_CLUSTER.CDB.DATABASE_HANG		
Backup Failures	HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE An event of type CRITICAL is created if there is a CDB backup with a FAILED status reported in the v$rman_status view.	HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE	com.oraclecloud.databaseservice.database.critical	N/A
Disk Group Usage	HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE An event of type CRITICAL is created when an ASM disk group reaches space usage of 90% or higher. An event of type INFORMATION is created when the ASM disk group space usage drops below 90%.	HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE	com.oraclecloud.databaseservice.dbsystem.critical com.oraclecloud.databaseservice.dbsystem.information (if < 90%)	Critical threshold: 90%
Memory Usage	CONFIGURATION.DB_GUEST.MEMORY.HUGEPAGES_TOO_LARGE An event of type CRITICAL is created when the amount of memory in the VM configured for HugePages is 90% or more of the total VM memory.	CONFIGURATION.DB_GUEST.MEMORY.HUGEPAGES_TOO_LARGE	com.oraclecloud.databaseservice.dbnode.critical	90%

Table 6-2 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
sshd Configuration	`CONFIGURATION.DB_GUEST.SSHD.INVALID` An event of type CRITICAL is created if unexpected values are set in the `/etc/ssh/sshd_config` file.	CONFIGURATION.DB_GUEST.SSHD.INVALID	`com.oraclecloud.databaseservice.dbnode.critical`	N/A
Disk Issues	`HEALTH.DB_GUEST.FILESYSTEM.CORRUPTION` A Write-then-Read operation with a dummy file has failed for a file system, typically indicating the operating system had detected an I/O error or inconsistency (i.e. corruption) with the file system and changed the file system mount mode from read-write to read-only. The following file systems are tested: • `/` • `/u01` • `/u02`	HEALTH.DB_GUEST.FILESYSTEM.CORRUPTION	`com.oraclecloud.databaseservice.dbnode.critical`	N/A

Table 6-2 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
Oracle EXAchk Reported Issues	`HEALTH.DB_CLUSTER.EXACHK.CRITICAL_ALERT` Oracle EXAchk is Exadata database platform's holistic health check that includes software, infrastructure and database configuration checks. CRITICAL check alerts should be addressed in 24 hours to maintain the maximum stability and availability of your system. This database service event alerts every 24 hours whenever there are any CRITICAL checks that are flagged in the most recent Oracle EXAchk report. The event will point to the latest Oracle EXAchk zip report.	HEALTH.DB_CLUSTER.EXACHK.CRITICAL_ALERT	`com.oraclecloud.databaseservice.dbnode.critical`	N/A

Example 6-62 Database Service DB Node Critical Events Examples

DB node critical reference events:

```
{
 "eventType" : "com.oraclecloud.databaseservice.dbnode.critical",
 "cloudEventsVersion" : "0.1",
 "eventTypeVersion" : "2.0",
 "source" : "SYSLENS/host_Name/DomU",
 "eventTime" : "2022-03-04T18:19:42Z",
 "contentType" : "application/json",
 "data" : {
   "compartmentId" : "compartment_ID",
   "compartmentName" : "compartment_Name",
   "resourceName" : "resource_Name",
   "resourceId" : "resource_ID",
   "additionalDetails" : {
     "serviceType" : "EXACS",
     "hostName" : "host_Name",
     "description" : "The '/' filesystem is over 90% used.",
     "eventName" : "HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE",
     "status" : "online"
```

```
    }
  },
  "eventID" : "a9752630-9be7-11ec-a203-00163eb980bb",
  "extensions" : {
    "compartmentId" : "compartment_ID"
  }
}
```

Temporarily Restrict Automatic Diagnostic Collections for Specific Events

Use the `tfactl blackout` command to temporarily suppress automatic diagnostic collections.

If you set blackout for a target, then Oracle Trace File Analyzer stops automatic diagnostic collections if it finds events in the alert logs for that target while scanning. By default, blackout will be in effect for 24 hours.

You can also restrict automatic diagnostic collection at a granular level, for example, only for **ORA-00600** or even only **ORA-00600** with specific arguments.

Syntax

```
tfactl blackout add|remove|print
-targettype host|crs|asm|asmdg|database|dbbackup|db_dataguard|db_tablespace|
pdb_tablespace|pdb|listener|service|os
-target all|name
[-container name]
[-pdb pdb_name]
-event all|"event_str1,event_str2"|availability
[-timeout nm|nh|nd|none]
[-c|-local|-nodes "node1,node2"]
[-reason "reason for blackout"]
[-docollection]
```

Parameters

Table 6-3 tfactl blackout Command Parameters

Parameter	Description
`add|remove|print|`	Adds, removes, or prints blackout conditions.

Table 6-3 (Cont.) tfactl blackout Command Parameters

Parameter	Description
`targettype` *type* **Target type:** `host`|`crs`|`asm`|`asmdg`|`database`|`dbbackup`|`db_dataguard`|`db_tablespace` |`pdb_tablespace`|`pdb`|`listener`|`service`|`os`	Limits blackout only to the specified target type. `host`: The whole node is under blackout. If there is host blackout, then every blackout element that's shown true in the Telemetry JSON will have the reason for the blackout. `crs`: Blackout the availability of the Oracle Clusterware resource or events in the Oracle Clusterware logs. `asm`: Blackout the availability of Oracle Automatic Storage Management (Oracle ASM) on this machine or events in the Oracle ASM alert logs. `asmdg`: Blackout an Oracle ASM disk group. `database`: Blackout the availability of an Oracle Database, Oracle Database backup, tablespace, and so on, or events in the Oracle Database alert logs. `dbbackup`: Blackout Oracle Database backup events (such as CDB or archive backups). `db_dataguard`: Blackout Oracle Data Guard events. `db_tablespace`: Blackout Oracle Database tablespace events (container database). `pdb_tablespace`: Blackout Oracle Pluggable Database tablespace events (Pluggable database). `pdb`: Blackout Oracle Pluggable Database events. `listener`: Blackout the availability of a listener. `service`: Blackout the availability of a service. `os`: Blackout one or more operating system records.
`target all`|*name*	Specify the target for blackout. You can specify a comma-delimited list of targets. By default, the target is set to `all`.
`container` *name*	Specify the database container name (`db_unique_name`) where the blackout will take effect (for PDB, `DB_TABLESPACE`, and `PDB_TABLESPACE`).
`pdb` *pdb_name*	Specify the PDB where the blackout will take effect (for `PDB_TABLESPACE` only).
`events all`|"*str1*,*str2*"	Limits blackout only to the availability events, or event strings, which should not trigger auto collections, or be marked as blacked out in telemetry JSON. `all`: Blackout everything for the target specified. *string*: Blackout for incidents where any part of the line contains the strings specified. Specify a comma-delimited list of strings.
`timeout` *nh*|*nd*|`none`	Specify the duration for blackout in number of hours or days before timing out. By default, the timeout is set to 24 hours (24h).
`c`|`local`	Specify if blackout should be set to cluster-wide or local. By default, blackout is set to `local`.
`reason` *comment*	Specify a descriptive reason for the blackout.
`docollection`	Use this option to do an automatic diagnostic collection even if a blackout is set for this target.

Example 6-63 tfactl blackout

- To blackout **event:** ORA-00600 on **target type:** database, **target:** mydb

```
tfactl blackout add -targettype database -target mydb -event "ORA-00600"
```

- To blackout **event:** ORA-04031 on **target type:** database, **target:** all

```
tfactl blackout add -targettype database -target all -event "ORA-04031" -
timeout 1h
```

- To blackout **db backup events** on **target type:** dbbackup, **target:** mydb

```
tfactl blackout add -targettype dbbackup -target mydb
```

- To blackout **db dataguard events** on **target type:** db_dataguard, **target:** mydb

```
tfactl blackout add -targettype db_dataguard -target mydb -timeout 30m
```

- To blackout **db tablespace events** on **target type:** db_tablespace, **target:** system, **container:** mydb

```
tfactl blackout add -targettype db_tablespace -target system -container
mydb -timeout 30m
```

- To blackout **ALL events** on **target type:** host, **target:** all

```
tfactl blackout add -targettype host -event all -target all -timeout 1h -
reason "Disabling all events during patching"
```

- To print blackout details

```
tfactl blackout print
```

```
.-------------------------------------------------------------------------
-------------------------------------------------------------------------
-------------------------.
|

myhostname
                    |
+---------------+--------------------+-----------
+------------------------------+------------------------------+--------
+---------------+----------------------------------------+
| Target Type   | Target             | Events    | Start
Time                   | End Time                     | Status | Do
Collection | Reason                              |
+---------------+--------------------+-----------
+------------------------------+------------------------------+--------
+---------------+----------------------------------------+
| HOST          | ALL                | ALL       | Thu Mar 24 16:48:39
UTC 2022 | Thu Mar 24 17:48:39 UTC 2022 | ACTIVE | false         |
Disabling all events during patching |
```

```
| DATABASE      | MYDB                | ORA-00600 | Thu Mar 24 16:39:03
UTC 2022 | Fri Mar 25 16:39:03 UTC 2022 | ACTIVE | false         |
NA                                    |
| DATABASE      | ALL                 | ORA-04031 | Thu Mar 24 16:39:54
UTC 2022 | Thu Mar 24 17:39:54 UTC 2022 | ACTIVE | false         |
NA                                    |
| DB_DATAGUARD  | MYDB                | ALL       | Thu Mar 24 16:41:38
UTC 2022 | Thu Mar 24 17:11:38 UTC 2022 | ACTIVE | false         |
NA                                    |
| DBBACKUP      | MYDB                | ALL       | Thu Mar 24 16:40:47
UTC 2022 | Fri Mar 25 16:40:47 UTC 2022 | ACTIVE | false         |
NA                                    |
| DB_TABLESPACE | SYSTEM_CDBNAME_MYDB | ALL       | Thu Mar 24 16:45:56
UTC 2022 | Thu Mar 24 17:15:56 UTC 2022 | ACTIVE | false         |
NA                                    |
'---------------+---------------------+-----------
+-------------------------------+-------------------------------+--------
+---------------+-------------------------------------'
```

- To remove blackout for **event:** ORA-00600 on **target type:** database, **target:** mydb

  ```
  tfactl blackout remove -targettype database -event "ORA-00600" -target mydb
  ```

- To remove blackout for **db backup events** on **target type:** dbbackup, **target:** mydb

  ```
  tfactl blackout remove -targettype dbbackup -target mydb
  ```

- To remove blackout for **db tablespace events** on **target type:** db_tablespace, **target:** system, **container:** mydb

  ```
  tfactl blackout remove -targettype db_tablespace -target system -container
  mydb
  ```

- To remove blackout for **host events** on **target type:** host, **target:** all

  ```
  tfactl blackout remove -targettype host -event all -target all
  ```

Remediation

These topics cover some common issues you might run into and how to address them.

- HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE
- AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN
- AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN
- AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN
- AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN
- AVAILABILITY.DB_GUEST.CRS_INSTANCE.EVICTION
- HEALTH.DB_CLUSTER.CDB.CORRUPTION
- HEALTH.DB_CLUSTER.CDB.ARCHIVER_HANG
- HEALTH.DB_CLUSTER.CDB.DATABASE_HANG

- HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE
- HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE
- Managing the Log and Diagnostic Files on Oracle Exadata Database Service on Dedicated Infrastructure
- CONFIGURATION.DB_GUEST.MEMORY.HUGEPAGES_TOO_LARGE
- CONFIGURATION.DB_GUEST.SSHD.INVALID
- HEALTH.DB_GUEST.FILESYSTEM.CORRUPTION
- HEALTH.DB_CLUSTER.EXACHK.CRITICAL_ALERT

HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE

Problem Statement: One or more VM guest file systems has free space below 10% free.

Risk: Insufficient VM guest file system free space can cause disk space allocation failure, which can result in wide-ranging errors and failures in Oracle software (Database, Clusterware, Cloud, Exadata).

Action:

Oracle Cloud and Exadata utilities run automatically to purge old log files and trace files created by Oracle software to reclaim file system space.

If the automatic file system space reclamation utilities cannot sufficiently purge old files to clear this event, then perform the following actions:

1. Remove unneeded files and/or directories created manually or by customer-installed applications or utilities. Files created by customer-installed software are outside the scope of Oracle's automatic file system space reclamation utilities. The following operating system command, run as the `opc` user, is useful for identifying directories consuming excessive disk space:

   ```
   $ sudo du -hx file-system-mount-point | sort -hr
   ```

   Only remove files or directories you are certain can be safely removed.
2. Reclaim `/u02` file system disk space by removing Database Homes that have no databases. For more information about managing Database Homes, see *Manage Oracle Database Homes on Exadata Database Service on Exadata Cloud Infrastructure Instance*.
3. Open service request to receive additional guidance about reducing file system space use.

Related Topics

- Managing Oracle Database Homes on an Exadata Cloud Infrastructure Instance
  You can delete or view information about Oracle Database Homes (referred to as "Database Homes" in Oracle Cloud Infrastructure) by using the Oracle Cloud Infrastructure Console, the API, or the CLI.

AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN

Problem Statement: The Cluster Ready Stack is in an offline state or has failed.

Risk: If the Cluster Ready Service is offline on a node, then the node cannot provide database services for the application.

Action:

1. Check if CRS was stopped by your administrator, as part of a planned maintenance event, or a scale up or down of local storage.
   a. The following patching events will stop CRS:
      i. GRID Patching
      ii. Exadata VM patching of Guest
      iii. Exadata VM Patching of Host
2. If CRS has stopped unexpectedly, then the current status can be checked by issuing the `crsctl check crs` command.
   a. If the node is not responding, then the VM node may be rebooting. Wait for the node reboot to finish, CRS will normally be started through the `init` process.
3. If CRS is still down, then investigate the cause of the failure by referring to the `alert.log` found in `/u01/app/grid/diag/crs/<node_name>/crs/trace`.
   Review the log entries corresponding to the date/time of the down event. Act on any potential remediation.
4. Restart the CRS, by issuing the `crsctl start crs` command.
5. A successful restart of CRS will generate the clearing event: `AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN_CLEARED`.

AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN

Problem Statement: A SCAN listener is down and unable to accept application connections.

Risk: If all SCAN listeners are down, then application connections to the database through the SCAN listener will fail.

Action:

Start the SCAN listener to receive the `DOWN_CLEARED` event.

DOWN event of type INFORMATION

1. If the event was caused by an Oracle Cloud maintenance action, such as performing a Grid Infrastructure software update, then no action is required. The affected SCAN listener will automatically failover to an available instance.
2. If the event was caused by user action, then start the SCAN listener at the next opportunity.

DOWN event of type CRITICAL

Check SCAN status and restart the SCAN listener.

1. Login to the VM as `opc` user and `sudo` to the `grid` user:

```
sudo su - grid
```

2. Check the SCAN listener status on any node:

```
srvctl status scan_listener
```

3. Start the SCAN listener:

```
srvctl start scan_listener
```

4. Recheck the SCAN listeners status on any node:
If the `scan_listener` is still down, then investigate the cause of the scan listener failure:

   a. Collect both the CRS and operating system logs 30 minutes prior and 10 minutes for the <*hostName*>indicated in the log. Note the time in the event payload is always provided in UTC. For `tfactl` collection, adjust the time to the timezone of the VM Cluster. As the grid user:

```
tfactl diagcollect -crs -os -node <hostName> -from "<eventTime
adjusted for local vm timezone> - 30 minute " -to "<eventTime adjusted
for local vm timezone> + 10 minutes"
```

   b. Review the SCAN listener log located under `/u01/app/grid/diag/tnslsnr/<hostName>/<listenerName>/trace`

AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN

Problem Statement: A client listener is down and unable to accept application connections.

Risk:

- If the node's client listener is down, then the database instances on the node cannot provide services for the application.
- If the client listener is down on all nodes, then any application that connects to any database using the SCAN or VIP will fail.

Action:

Start the client listener to receive the `DOWN_CLEARED` event.

DOWN event of type INFORMATION

1. If the event was caused by an Oracle Cloud maintenance action, such as performing a Grid Infrastructure software update, then no action is required. The affected client listener will automatically restart when maintenance affecting the grid instance is complete.
2. If the event was caused by user action, then start the client listener at the next opportunity.

DOWN event of type CRITICAL

Check the client listener status and then restart the client listener.

1. Login to the VM as `opc` user and `sudo` to the `grid` user:

```
[opc@vm ~] sudo su - grid
```

2. Check the client listener status on any node:

```
[grid@vm ~] srvctl status listener
```

3. Start the client listener:

```
[grid@vm ~] srvctl start listener
```

4. Recheck the client listener status on any node:
   If the client listener is still down, then investigate the cause of the client listener failure:
   a. Use tfactl to collect both the CRS and operating system logs 30 minutes prior and 10 minutes for the <*hostName*> indicated in the log. Note the time in the event payload is always provided in UTC. For tfactl collection, adjust the time to the timezone of the VM Cluster.

      ```
      [grid@vm ~] tfactl diagcollect -crs -os -node <hostName> -from
      "<eventTime adjusted for local vm timezone> - 30 minute " -to
      "<eventTime adjusted for local vm timezone> + 10 minutes"
      ```

   b. Review the listener log located under `/u01/app/grid/diag/tnslsnr/<hostName>/<listenerName>/trace`

AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN

Problem Statement: A database instance has gone down.

Risk: A database instance has gone down, which may result in reduced performance if database instances are available on other nodes in the cluster, or complete downtime if database instances on all nodes are down.

Action:

Start the database instance to receive the `DOWN_CLEARED` event.

DOWN event of type INFORMATION

1. If the event was caused by an Oracle Cloud maintenance action, such as performing a Database Home software update, then no action is required. The affected database instance will automatically restart when maintenance affecting the instance is complete.
2. If the event was caused by user action, then start the affected database instance at the next opportunity.

DOWN event of type CRITICAL

1. Check database status and restart the down database instance.
   a. Login to the VM as `oracle` user:
   b. Set the environment:

      ```
      [oracle@vm ~] . <dbName>.env
      ```

   c. Check the database status:

      ```
      [oracle@vm ~] srvctl status database -db <dbName>
      ```

   d. Start the database instance:

      ```
      [oracle@vm ~] srvctl start instance -db <dbName> -instance
      <instanceName>
      ```

2. Investigate the cause of the database instance failure.

a. Review Trace File Analyzer (TFA) events for the database:

```
[oracle@vm ~] tfactl events -database <dbName> -instance <instanceName>
```

b. Review the database alert log located at `$ORACLE_BASE/diag/rdbms/<dbName>/<instanceName>/trace/alert_<instanceName>.log`

AVAILABILITY.DB_GUEST.CRS_INSTANCE.EVICTION

Problem Statement: The Oracle Clusterware is designed to perform a node eviction by removing one or more nodes from the cluster if some critical problem is detected. A critical problem could be a node not responding via a network heartbeat, a node not responding via a disk heartbeat, a hung or severely degraded machine, or a hung ocssd.bin process. The purpose of this node eviction is to maintain the overall health of the cluster by removing impaired members.

Risks: During the time it takes to restart the evicted node, the node cannot provide database services for the application.

Action: CRS node eviction could be caused by OCSSD (CSS daemon), CSSDAGENT, or CSSDMONITOR processes. This requires determining which process was responsible for the node eviction and reviewing the relevant log files. Common causes of OCSSD eviction are network failures/latencies, IO issues with CSS voting disks, a member kill escalation. CSSDAGENT or CSSDMONITOR evictions could be OS scheduler problem or a hung thread within CSS daemon.

Log files to review include:

- clusterware alert log
- cssdagent log
- cssdmonitor log
- ocssd log
- lastgasp log
- `/var/log/messages`
- CHM/OS Watcher data
- `opatch lsinventory` detail

For more information on collecting together most of the files, see *Autonomous Health Framework (AHF) - Including TFA and ORAchk/EXAchk (Doc ID 2550798.1)*.

For more information on troubleshooting CRS node eviction, see *Troubleshooting Clusterware Node Evictions (Reboots) (Doc ID 1050693.1)*.

Related Topics

- Autonomous Health Framework (AHF) - Including TFA and ORAchk/EXAchk (Doc ID 2550798.1)
- Troubleshooting Clusterware Node Evictions (Reboots) (Doc ID 1050693.1)

HEALTH.DB_CLUSTER.CDB.CORRUPTION

Problem Statement: Corruptions can lead to application or database errors and in worse case result in significant data loss if not addressed promptly.

A corrupt block is a block that was changed so that it differs from what Oracle Database expects to find. Block corruptions can be categorized as physical or logical:

- In a physical block corruption, which is also called a media corruption, the database does not recognize the block at all; the checksum is invalid or the block contains all zeros. An example of a more sophisticated block corruption is when the block header and footer do not match.
- In a logical block corruption, the contents of the block are physically sound and pass the physical block checks; however, the block can be logically inconsistent. Examples of logical block corruption include incorrect block type, incorrect data or redo block sequence number, corruption of a row piece or index entry, or data dictionary corruptions.

For more information, see *Physical and Logical Block Corruptions. All you wanted to know about it. (Doc ID 840978.1)*.

Block corruptions can also be divided into interblock corruption and intrablock corruption:

- In an intrablock corruption, the corruption occurs in the block itself and can be either a physical or a logical block corruption.
- In an interblock corruption, the corruption occurs between blocks and can only be a logical block corruption.

Oracle checks for the following errors in the `alert.log`:

- ORA-01578
- ORA-00752
- ORA-00753
- ORA-00600 [3020]
- ORA-00600 [kdsgrp1]
- ORA-00600 [kclchkblk_3]
- ORA-00600 [13013]
- ORA-00600 [5463]

Risk: A data corruption outage occurs when a hardware, software, or network component causes corrupt data to be read or written. The service-level impact of a data corruption outage may vary, from a small portion of the application or database (down to a single database block) to a large portion of the application or database (making it essentially unusable). If remediation action is not taken promptly, then potential downtime and data loss can increase.

Action:

The current event notification currently triggers on physical block corruptions (ORA-01578), lost writes (ORA-00752, ORA-00753 and ORA-00600 with first argument 3020), and logical corruptions (typical detected from ORA-00600 with first argument of kdsgrp1, kdsgrp1, kclchkblk_3, 13013 OR 5463).

Oracle recommends the following steps:

1. Confirm that these corruptions were reported in the alert.log trace file. Log a Service Request (SR) with latest EXAchk report, excerpt of the alert.log and trace file containing the corruption errors, any history of recent application, database or software changes and any system, clusterware and database logs for the same time period. For all these cases, a TFA collection should be available and should be attached to the SR.
2. For repair recommendations, refer to *Handling Oracle Database Corruption Issues (Doc ID 1088018.1)*.

For physical corruptions or ORA-1578 errors, the following notes will be helpful:

- Doc ID 1578.1 : OERR: ORA-1578 "ORACLE data block corrupted (file # %s, block # %s)" Primary Note
- Doc ID 472231.1 : How to identify all the Corrupted Objects in the Database reported by RMAN
- Doc ID 819533.1 : How to identify the corrupt Object reported by ORA-1578 / RMAN / DBVERIFY
- Depending on the object that has the corruption, follow the guidance in Doc ID 1088018.1. Note RMAN can be used to recover one or many data block that are physically corrupted. Also using Active Data Guard with real time apply, auto block repair of physical data corruptions would have occurred automatically.

For logical corruptions caused by lost writes (ORA-00752, ORA-00753 and ORA-00600 with first argument 3020) on the primary or standby databases, they will be detected on the primary or with standby's redo apply process. The following notes will be helpful:

- Follow the guidance, follow Doc ID 1088018.1.
- If you have a standby and lost write corruption on the primary or standby, refer to Resolving ORA-00752 or ORA-00600 [3020] During Standby Recovery (Doc ID 1265884.1)

For logical corruptions (typical detected from ORA-00600 with arguments of kdsgrp1, kclchkblk_3, 13013 OR 5463):

- Follow the guidance, follow Doc ID 1088018.1 for specific guidance on the error that was detected.
- If you have a standby and logical corruption on the primary, refer to Resolving Logical Block Corruption Errors in a Physical Standby Database (Doc ID 2821699.1)

Related Topics

- Physical and Logical Block Corruptions. All you wanted to know about it. (Doc ID 840978.1)
- OERR: ORA-1578 "ORACLE data block corrupted (file # %s, block # %s)" Primary Note (Doc ID 1578.1)
- How to identify all the Corrupted Objects in the Database with RMAN (Doc ID 472231.1)
- How to identify the corrupt Object reported by ORA-1578 / RMAN / DBVERIFY (Doc ID 819533.1)
- Resolving ORA-00752 or ORA-600 [3020] During Standby Recovery (Doc ID 1265884.1)
- Resolving Logical Block Corruption Errors in a Physical Standby Database (Doc ID 2821699.1)

HEALTH.DB_CLUSTER.CDB.ARCHIVER_HANG

Problem Statement: CDB RAC Instance may temporarily or permanently stall due to the log writer's (LGWR) inability to write the log buffers to an online redo log. This occurs because all online logs need archiving. Once the archiver (ARC) can archive at least one online redo log, LGWR will be able to resume writing the log buffers to online redo logs and the application impact will be alleviated.

Risk: If the archiver hang is temporary, then this can result in a small application brown out or stall for application processes attempting to commit their database changes. If the archiver is not unblocked, applications can experience extended delay in processing.

Action:

- See, *Script To Find Redo log Switch History And Find Archivelog Size For Each instance In RAC (Doc ID 2373477.1)* to determine the hourly frequency for each thread/instance.
- If any hourly bucket is greater than 12, then consider resizing the online redo logs. See item 2 below for resizing steps.
- If the database hangs are temporary, then the archiver may be unable to keep up with the redo log generated. Check the `alert.log`, `$ORACLE_BASE/diag/rdbms/<dbName>/<instanceName>/trace/alert_<instanceName>.log`, for "All online logs need archiving", multiple events in a short period can indicate 2 possible solutions.
  - If the number of redo logs groups per thread is less than 4, then consider adding additional logs groups to reach 4, see item 1 below for add redo log steps.
  - The other possible solution is to resize the redo logs, see item 2 below for resizing steps.
- For Data Guard and Non Data Guard review the *Configure Online Redo Logs Appropriately* of section Oracle Database High Availability Overview and Best Practices for sizing guidelines.

1. Add a redo log group for each thread. The additional redo log should equal the current log size.
   a. Use the following query:

      ```
      select max(group#) Ending_group_number, thread#, count(*)
      number_of_groups_per_thread, bytes redo_size_in_bytes from v$log group
      by thread#,bytes
      ```

   b. Add one new group per thread using the same size as the current redo logs.

      ```
      alter database add logfile thread <thread_number> Group <max group + 1>
      ('<DATA_DISKGROUP>') size <redo_size_in_bytes>
      ```

2. Resize the online redo logs by adding larger redo logs and dropping the current smaller redo logs.
   a. Use the following query:

      ```
      select max(group#) Ending_group_number, thread#, count(*)
      number_of_groups_per_thread, bytes redo_size_in_bytes from v$log group
      by thread#,bytes
      ```

   b. Add the same number of redo logs for each thread <*number_of_groups_per_thread*> that currently exist. The <*new_redo_size_in_bytes*> should be based on *Configure Online Redo Logs Appropriately* of section Oracle Database High Availability Overview and Best Practices.

      i. `alter database add logfile thread <thread_number> Group <max group + 1> ('<DATA_DISKGROUP>') size <new_redo_size_in_bytes>`

      ii. The original smaller redo logs should be deleted. A redo log can only be deleted if its status is inactive.

To determine the status of a redo logs issue:

```
select group#, thread#, status, bytes from v$log order by bytes,
group#, thread#;
```

To delete the original smaller redo logs:

```
alter database drop logfile <group#>
```

- If the database is hung, the primary log archive destination and alternate may be full. Review the *HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE* for details on freeing space in RECO and DATA disk groups.

Related Topics

- Script To Find Redolog Switch History And Find Archivelog Size For Each Instances In RAC (Doc ID 2373477.1)
- Configure Online Redo Logs Appropriately
- HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE

HEALTH.DB_CLUSTER.CDB.DATABASE_HANG

Problem Statement: Hang management detected a process hang and generated a ORA-32701 error message. Additionally, this event may be raised if Diagnostic Process (DIA0) process detects a hang in a critical database process.

Risk: A hang can indicate resource, operating system, or application coding related issues.

Action:

Investigate the cause of the session hang.

1. Review TFA events for the database for the following message patterns corresponding to the date/time of the event: ORA-32701, "DIA0 Critical Database Process Blocked" or "DIA0 Critical Database Process As Root".

   ```
   [oracle@vm ~] tfactl events -database <dbName> -instance <instanceName>
   ```

2. Review the `alert.log` file.

   ```
   $ORACLE_BASE/diag/rdbms/<dbName>/<instanceName>/trace/
   alert_<instanceName>.log
   ```

3. **For ora-32701:** An overloaded system can cause slow progress, which can be interpreted as a hang.
   The hang manager may attempt to resolve the hang by terminating the final blocker process.
4. **For DIA0 Critical Database Process messages:** Review the related diagnostic lines indicating the process and the reason for the hang.

HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE

Problem Statement: A daily incremental BACKUP of the CDB failed.

Risk: A failure of the backup can compromise the ability to use the backups for restore/recoverability of the database. Recoverability Point Object (RPO) and the Recoverability Time Object (RTO) can be impacted.

Action:

Review the RMAN logs corresponding to the date/time of the event. Note the event time stamp <*eventTime*> is in UTC, adjust as necessary for the VM's timezone.

- For Exadata Cloud Infrastructure Oracle Managed Backups or User Configured Backups under `dbaascli`:
  - RMAN output can be found at `/var/opt/oracle/log/<DB_NAME>/obkup`. Daily incremental logs have the following format `obkup_yyyy-mm-dd_24hh:mm:ss.zzzzzzzzzzz.log` within the `obkup` directory. The logs are located on the lowest active node/instance of the database when the backup was initiated.
  - Review the log for any failures:
    - If the failure is due to an external event outside of RMAN, for example the backup location was full or a networking issue, resolve the external issue.
    - For other RMAN script errors, collect the diagnostic logs and open a Service Request. See *DBAAS Tooling: Using dbaascli to Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check*.
  - If the issue is transient or is resolved, take a new incremental backup: See *dbaascli database backup*.
- For Customer owned and managed backup taken through RMAN:
  - Review the RMAN logs for the backup.

Related Topics

- DBAAS Tooling: Using dbaascli to Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check
- dbaascli database backup
  To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE

Problem Statement: ASM disk group space usage is at or exceeds 90%.

Risk: Insufficient ASM disk group space can cause database creation failure, tablespace and data file creation failure, automatic data file extension failure, or ASM rebalance failure.

Action:

ASM disk group used space is determined by the running the following query while connected to the ASM instance.

```
[opc@node ~] sudo su - grid
[grid@node ~] sqlplus / as sysasm

SQL> select 'ora.'||name||'.dg', total_mb, free_mb, round ((1-(free_mb/
total_mb))*100,2) pct_used from v$asm_diskgroup;

NAME                           TOTAL_MB    FREE_MB   PCT_USED
------------------------------ ---------- ---------- ----------
```

```
ora.DATAC1.dg                75497472    7408292      90.19
ora.RECOC1.dg                18874368   17720208       6.11
```

ASM disk group capacity can be increased in the following ways:

1. Scale Exadata VM Cluster storage to add more ASM disk group capacity. See *Scaling an Exadata Cloud Infrastructure Instance*.
2. Scale Exadata Infrastructure storage to add more ASM disk group capacity. See *Scaling Exadata X8M and X9M Compute and Storage*.

DATA disk group space use can be reduced in the following ways:

1. Drop unused data files and temp files from databases. See *Dropping Data Files*.
2. Terminate unused databases (e.g. test databases). See *Using the Console to Terminate a Database*.

RECO disk group space use can be reduced in the following ways:

1. Drop unnecessary Guaranteed Restore Points. See *Using Normal and Guaranteed Restore Points*.
2. Delete archived redo logs or database backups already backed up outside the Flash Recovery Area (FRA). See *Maintaining the Fast Recovery Area*.

SPARSE disk group space use can be reduced in the following ways:

1. Move full copy test master databases to another disk group (e.g. DATA).
2. Drop unused snapshot databases or test master databases. See *Managing Exadata Snapshots*.

For more information about managing the log and diagnostic files, see *Managing the Log and Diagnostic Files on Oracle Exadata Database Service on Dedicated Infrastructure*.

Related Topics

- Scaling Exadata X8M and X9M Compute and Storage
  The flexible X8M and X9M system model is designed to be easily scaled in place, with no need to migrate the database using a backup or Data Guard
- Dropping Data Files
- To terminate a database
- Using Normal and Guaranteed Restore Points
- Maintaining the Fast Recovery Area
- Managing Exadata Snapshots
- Managing the Log and Diagnostic Files on Oracle Exadata Database Service on Dedicated Infrastructure

Managing the Log and Diagnostic Files on Oracle Exadata Database Service on Dedicated Infrastructure

The software components in Oracle Exadata Database Service on Dedicated Infrastructure generate a variety of log and diagnostic files, and not all these files are automatically archived and purged. Thus, managing the identification and removal of these files to avoid running out of file storage space is an important administrative task.

Database deployments on ExaDB-D include the `cleandblogs` script to simplify this administrative task. The script runs daily as a `cron` job on each compute node to archive key files and remove old log and diagnostic files.

The `cleandblogs` script operates by using the `adrci` (Automatic Diagnostic Repository [ADR] Command Interpreter) tool to identify and purge target diagnostic folders and files for each Oracle Home listed in `/etc/oratab`. It also targets Oracle Net Listener logs, audit files, and core dumps.

On ExaDB-D, the script is run separately as the `oracle` user to clean log and diagnostic files that are associated with Oracle Database, and as the `grid` user to clean log and diagnostic files that are associated with Oracle Grid Infrastructure.

The `cleandblogs` script uses a configuration file to determine how long to retain each type of log or diagnostic file. You can edit the file to change the default retention periods. The file is located at `/var/opt/oracle/cleandb/cleandblogs.cfg` on each compute node.

> **Note:**
>
> Configure an optimal retention period for each type of log or diagnostic file. An insufficient retention period will hinder root cause analysis and problem investigation.

Parameter	Description and Default Value
`AlertRetention`	Alert log (`alert_instance.log`) retention value in days. Default value: 14
`ListenerRetention`	Listener log (`listener.log`) retention value in days. Default value: 14
`AuditRetentionDB`	Database audit (`*.aud`) retention value in days. Default value: 1
`CoreRetention`	Core dump/files (`*.cmdp*`) retention value in days. Default value: 7
`TraceRetention`	Trace file (`*.tr*` and `*.prf`) retention value in days. Default value: 7
`longpRetention`	Data designated in the Automatic Diagnostic Repository (ADR) as having a long life (the `LONGP_POLICY` attribute). For information about ADR, see *Automatic Diagnostic Repository (ADR)* in the *Oracle Database Administrator's Guide*. Default value: 14
`shortpRetention`	Data designated in the Automatic Diagnostic Repository (ADR) as having a short life (the `SHORTP_POLICY` attribute). For information about ADR, see *Automatic Diagnostic Repository (ADR)* in the *Oracle Database Administrator's Guide*. Default value: 7

Parameter	Description and Default Value
LogRetention	Log file retention in days for files under /var/opt/oracle/log and log files in ACFS under /var/opt/oracle/dbaas_acfs/log. Default value: 14
LogDirRetention	cleandblogs logfile retention in days. Default value: 14
ScratchRetention	Temporary file retention in days for files under /scratch. Default value: 7

Archiving Alert Logs and Listener Logs

When cleaning up alert and listener logs, cleandblogs first archives and compresses the logs, operating as follows:

1. The current log file is copied to an archive file that ends with a date stamp.
2. The current log file is emptied.
3. The archive file is compressed using gzip.
4. Any existing compressed archive files older than the retention period are deleted.

Running the cleandblogs Script Manually

The cleandblogs script automatically runs daily on each compute node, but you can also run the script manually if the need arises.

1. Connect to the compute node as the oracle user to clean log and diagnostic files that are associated with Oracle Database, or connect as the grid user to clean log and diagnostic files that are associated with Oracle Grid Infrastructure.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.

   Change to the directory containing the cleandblogs script:

   ```
   $ cd /var/opt/oracle/cleandb
   ```

2. Run the cleandblogs script:

   ```
   $ ./cleandblogs.pl
   ```

   When running the script manually, you can specify an alternate configuration file to use instead of cleandblogs.cfg by using the --pfile option:

   ```
   $ ./cleandblogs.pl --pfile config-file-name
   ```

3. Close your connection to the compute node:

   ```
   $ exit
   ```

Related Topics

- Automatic Diagnostic Repository (ADR)

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

CONFIGURATION.DB_GUEST.MEMORY.HUGEPAGES_TOO_LARGE

Problem Statement: Too much VM memory is allocated for HugePages use.

Risk: Excessive memory allocated to HugePages may result in poor database performance, or the system running out of memory, experiencing excessive swapping, or having crucial system services fail, causing system crash or node eviction.

Action:

1. Reduce HugePages memory use. To determine the proper setting for operating system parameter vm.nr_hugepages, see My Oracle Support document 361323.1.
2. Scale up VM memory. For more information about scaling VM memory, see *Introduction to Scale Up or Scale Down Operations*.

Related Topics

- https://support.oracle.com/rs?type=doc&id=361323.1
- Introduction to Scale Up or Scale Down Operations
  With the Multiple VMs per Exadata system (MultiVM) feature release, you can scale up or scale down your VM cluster resources.

CONFIGURATION.DB_GUEST.SSHD.INVALID

Problem Statement: SSHD configuration is unexpected.

SSHD Configuration Setting	Expected Value
`PubkeyAuthentication`	yes
`AuthorizedKeysFile`	`.ssh/authorized_keys` This file must exist in `root` user home directory.
`HostbasedAuthentication`	no
`IgnoreUserKnownHosts`	yes
`IgnoreRhosts`	yes
`PermitEmptyPasswords`	no
`PasswordAuthentication`	no
`ChallengeResponseAuthentication`	no
`GSSAPIAuthentication`	no
`UsePAM`	yes
`PrintMotd`	no
`UsePrivilegeSeparation`	yes
`PermitUserEnvironment`	no
`Compression`	delayed
`MaxStartups`	100

SSHD Configuration Setting	Expected Value
`AcceptEnv`	Must contain one of the following: • `LANG` • `LC_CTYPE` • `LC_NUMERIC` • `LC_TIME` • `LC_COLLATE` • `LC_MONETARY` • `LC_MESSAGES` • `LC_PAPER` • `LC_NAME` • `LC_ADDRESS` • `LC_TELEPHONE` • `LC_MEASUREMENT` • `LC_IDENTIFICATION` • `LC_ALL` • `LANGUAGE` • `XMODIFIERS`
`Subsystem`	**sftp** `/usr/libexec/openssh/sftp-server`
`Protocol`	2
`AddressFamily`	`inet`

Risk: SSHD configuration is unexpected which may cause Oracle Cloud automation failure or prevent customer SSH access to the VM.

Action: Change SSHD to match expected configuration.

1. Verify SSHD service is active.

   ```
   $ sudo systemctl is-active sshd.service
   active
   ```

   If SSHD service is inactive, then start it.

   ```
   $ sudo systemctl start sshd.service
   ```

2. Verify SSHD service is enabled.

   ```
   $ sudo /opt/oracle.cellos/host_access_control ssh-service --status
   [INFO] [IMG-SEC-1201] Service sshd is enabled {1}
   ```

   If SSHD service is disabled, then enable it.

   ```
   $ sudo /opt/oracle.cellos/host_access_control ssh-service --enable
   ```

3. Change SSHD configuration to match the expected values according to the table shown in the Problem Statement section above.

SSHD Configuration Setting	How to Change Current setting
`Ciphers`	`/opt/oracle.cellos/host_access_control sshciphers --help`
`MACs`	`/opt/oracle.cellos/host_access_control ssh-macs --help`
`PermitRootLogin`	`/opt/oracle.cellos/host_access_control rootssh --help`
`ClientAliveInterval`	`/opt/oracle.cellos/host_access_control idle-timeout --help`
`ClientAliveCountMax`	`/opt/oracle.cellos/host_access_control idle-timeout --help`
`ListenAddress`	`/opt/oracle.cellos/host_access_control ssh-listen --help`
ALL OTHER PARAMETERS	• Edit `/etc/ssh/sshd_config`. • Restart `sshd.service`.`$ sudo systemctl restart sshd.service`

HEALTH.DB_GUEST.FILESYSTEM.CORRUPTION

Problem Statement: A file system that is expected to be read-write can no longer be written to.

Risk: Oracle software (Linux, Database, Clusterware, Cloud, Exadata) requires write access to file systems to operate correctly.

Action:

`/u01` and `/u02` file systems:

1. Stop running services, if any, that are using the file system, such as Oracle Clusterware, Trace File Analyzer (TFA), and Enterprise Manager (EM) agent.
2. Unmount the file system.
3. Run file system check and repair.
   - **ext4:** Refer to *Checking and Repairing a File System*.
   - **xfs:** Refer to *Checking and Repairing an XFS File System*.
   - If the file system cannot be repaired then open a service request with Oracle Support for assistance.
4. Mount the file system.
5. Start the services.

/ (root) file system:

Open a service request with Oracle Support for assistance.

- If there is VM access, then collect full `dmesg(1)` command output and provide it to Oracle Support.
- Note that / (root) file system repair is possible only with console access.

Related Topics

- Checking and Repairing a File System

- Checking and Repairing an XFS File System

HEALTH.DB_CLUSTER.EXACHK.CRITICAL_ALERT

Problem Statement: A CRITICAL Exachk check failed and should be reviewed and addressed as soon as possible.

Risk: A CRITICAL check is expected to impact a large number of customers AND should be addressed immediately (for example, within 24 hours) AND meets one or more of the following criteria:

1. On-disk corruption or data loss
2. Intermittent wrong results with Exadata feature usage (e.g. smart scan)
3. System wide availability impact
4. Severe system wide performance impact seriously affecting application service Service Level Agreements (SLAs)
5. Compromised redundancy and inability to restore redundancy
6. Inability to update software in a rolling manner
7. Configuration error that could lead to an unexpected or unknown impact

Action:

Recommend that you bring up the EXAchk HTML report from the latest EXAchk zip file and click "**view**" on each CRITICAL check and follow the recommendation guidance that contains: Benefit/Impact, Risk, and Action/Repair guidance. Once the CRITICAL check is addressed, the next EXAchk run will pass that check. For more information about Oracle EXAchk, see *Oracle Exadata Database Machine Exachk (Doc ID 1070954.1)*.

As the `root` user, you can re-run EXAchk command by issuing:

```
/usr/bin/exachk -profile exatier1 -noupgrade -dball
```

If the check results are returning false data, then log a Service Request.

If there is a CRITICAL check that needs to be temporarily excluded, then follow the "**Skipping Specific Best Practice Checks in Exadata Cloud**" section of *Oracle Exadata Database Machine Exachk (Doc ID 1070954.1)*.

Related Topics

- Oracle Exadata Database Machine Exachk (Doc ID 1070954.1)

Application VIP Event Types

These are the event types that Application VIPs in Oracle Cloud Infrastructure emit.

Friendly Name	Event Type
Application Virtual IP (VIP) - Create Begin	`com.oraclecloud.databaseservice.createapplicationvip.begin`
Application Virtual IP (VIP) - Create End	`com.oraclecloud.databaseservice.createapplicationvip.end`
Application Virtual IP (VIP) - Delete Begin	`com.oraclecloud.databaseservice.deleteapplicationvip.begin`

Friendly Name	Event Type
Application Virtual IP (VIP) - Delete End	com.oraclecloud.databaseservice.deleteapplicationvip.end

Application VIP Event Types Examples:

This is a reference event for Application Virtual IP (VIP) - Create Begin:


```
{
  "id":
"ocid1.eventschema.oc1.phx.5ur5er8bddumnu9r84rtt2c3282s5no31vsthibyqvvsisotnwp
csg9idv6q",
  "serviceName": "Database",
  "displayName": "Application Virtual IP (VIP) - Create Begin",
  "eventType": "com.oraclecloud.databaseservice.createapplicationvip.begin",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2022-12-15T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "id",
      "type": "string"
    },
    {
      "name": "definedTags",
      "type": [
        "null",
        "Map<String, Map<String, Object>>"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "Map<String, String>"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",
      "type": [
        "null",
```

```
        "string"
      ]
    },
    {
      "name": "hostnameLabel",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "cloudVmClusterId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "compartmentId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "vcnIpId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "ipAddress",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "subnetId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "networkType",
      "type": [
        "null",
        "string"
      ]
    }
  ],
  "exampleEvent": {
    "eventType": "com.oraclecloud.databaseservice.createapplicationvip.begin",
    "cloudEventsVersion": "0.1",
```

```
    "eventTypeVersion": "2.0",
    "source": "databaseservice",
    "contentType": "application/json",
    "eventID": "ab2ac219-b435-1045-aaf3-13cd909ec106",
    "eventTime": "2022-12-16T21:16:04.000Z",
    "data": {
      "resourceId": "ocid1.applicationvip.oc1.....unique_id",
      "resourceName": "my_application_vip",
      "tagSlug": null,
      "compartmentId": "ocid1.compartment.oc1.....unique_id",
      "request": {
        "id": "4260c9fd-d36b-4bc8-866e-c2dd53f34b2f",
        "path": null,
        "action": null,
        "parameters": null,
        "headers": null
      },
      "response": {
        "status": null,
        "responseTime": null,
        "headers": null,
        "payload": null,
        "message": ""
      },
      "stateChange": {
        "previous": null,
        "current": {
          "lifecycleState": "PROVISIONING",
          "hostnameLabel": "my_application_vip",
          "freeTags": {},
          "definedTags": {}
        }
      },
      "eventGroupingId": "csid74237ee84398b60cf1b834c81602/
f43a881dc99542318d46fa9285bdf2c5/6AC9F7641E1A5AD5C27D1650CB17E822",
      "eventName": "CreateApplicationVip",
      "availabilityDomain": "",
      "resourceVersion": null,
      "additionalDetails": {
        "id": "ocid1.applicationvip.oc1.....unique_id",
        "freeformTags": {},
        "definedTags": {},
        "timeCreated": "2022-12-15T21:17:59.000Z",
        "timeUpdated": "2022-12-15T21:18:04.389Z",
        "lifecycleState": "PROVISIONING",
        "lifecycleDetails": "",
        "hostnameLabel": "my_application_vip",
        "cloudVmClusterId": "ocid1.cloudvmcluster.oc1.....unique_id",
        "compartmentId": "ocid1.compartment.oc1.....unique_id",
        "vcnIpId": "ocid1.privateip.oc1.....unique_id",
        "ipAddress": "10.0.0.0",
        "subnetId": "ocid1.subnet.oc1.....unique_id",
        "networkType": "CLIENT"
      }
    }
  },
```

```
  "timeCreated": "2022-12-15T16:31:31.979Z"
}
```

This is a reference event for Application Virtual IP (VIP) - Create End:

```
{
  "id":
"ocid1.eventschema.oc1.phx.c1ok1948lwge4il6m85ta4jdlbnh1yjrjltrabujyv52calb0el
p263oyqrm",
  "serviceName": "Database",
  "displayName": "Application Virtual IP (VIP) - Create End",
  "eventType": "com.oraclecloud.databaseservice.createapplicationvip.end",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2022-12-15T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "id",
      "type": "string"
    },
    {
      "name": "definedTags",
      "type": [
        "null",
        "Map<String, Map<String, Object>>"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "Map<String, String>"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",
      "type": [
        "null",
        "string"
      ]
    },
    {
```

```
      "name": "hostnameLabel",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "cloudVmClusterId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "compartmentId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "vcnIpId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "ipAddress",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "subnetId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "networkType",
      "type": [
        "null",
        "string"
      ]
    }
  ],
  "exampleEvent": {
    "eventType": "com.oraclecloud.databaseservice.createapplicationvip.end",
    "cloudEventsVersion": "0.1",
    "eventTypeVersion": "2.0",
    "source": "databaseservice",
    "contentType": "application/json",
    "eventID": "bc122d87-ac42-8731-ccd1-09ab320eef11",
```

```
    "eventTime": "2022-12-16T21:16:04.000Z",
    "data": {
      "resourceId": "ocid1.applicationvip.oc1.....unique_id",
      "resourceName": "my_application_vip",
      "tagSlug": null,
      "compartmentId": "ocid1.compartment.oc1.....unique_id",
      "request": {
        "id": "195eb9b5-b5a0-474d-a1c3-86189d8eeb2c",
        "path": null,
        "action": null,
        "parameters": null,
        "headers": null
      },
      "response": {
        "status": null,
        "responseTime": null,
        "headers": null,
        "payload": null,
        "message": ""
      },
      "stateChange": {
        "previous": null,
        "current": {
          "lifecycleState": "AVAILABLE",
          "hostnameLabel": "my_application_vip",
          "freeTags": {},
          "definedTags": {}
        }
      },
      "eventGroupingId":
"6CEB05B6C81E4B19855AD716E90F5BC3/070ECF4976BDD89B16849A92B95564A6/1418EDD7590
B8D5DDFF947FC3161F358",
      "eventName": "CreateApplicationVip",
      "availabilityDomain": "",
      "resourceVersion": null,
      "additionalDetails": {
        "id": "ocid1.applicationvip.oc1.....unique_id",
        "freeformTags": {},
        "definedTags": {},
        "timeCreated": "2022-12-15T21:17:59.000Z",
        "timeUpdated": "2022-12-15T21:18:04.389Z",
        "lifecycleState": "AVAILABLE",
        "lifecycleDetails": "",
        "hostnameLabel": "my_application_vip",
        "cloudVmClusterId": "ocid1.cloudvmcluster.oc1.....unique_id",
        "compartmentId": "ocid1.compartment.oc1.....unique_id",
        "vcnIpId": "ocid1.privateip.oc1.....unique_id",
        "ipAddress": "10.0.0.0",
        "subnetId": "ocid1.subnet.oc1.....unique_id",
        "networkType": "CLIENT"
      }
    }
  },
  "timeCreated": "2022-12-15T16:31:31.979Z"
}
```

This is a reference event for Application Virtual IP (VIP) - Delete Begin:


```
{
  "id":
"ocid1.eventschema.oc1.phx.m2gheil6f1nfzb9ggpkkv17wdomdks8zin9nntqlghui6bckh17
yu0m5jcqt",
  "serviceName": "Database",
  "displayName": "Application Virtual IP (VIP) - Delete Begin",
  "eventType": "com.oraclecloud.databaseservice.deleteapplicationvip.begin",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2022-12-15T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "id",
      "type": "string"
    },
    {
      "name": "definedTags",
      "type": [
        "null",
        "Map<String, Map<String, Object>>"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "Map<String, String>"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "hostnameLabel",
      "type": [
        "null",
        "string"
```

```
      ]
    },
    {
      "name": "cloudVmClusterId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "compartmentId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "vcnIpId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "ipAddress",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "subnetId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "networkType",
      "type": [
        "null",
        "string"
      ]
    }
  ],
  "exampleEvent": {
    "eventType": "com.oraclecloud.databaseservice.deleteapplicationvip.begin",
    "cloudEventsVersion": "0.1",
    "eventTypeVersion": "2.0",
    "source": "databaseservice",
    "contentType": "application/json",
    "eventID": "e32cb1fe-123d-8341-de13-2be5f18ab31e",
    "eventTime": "2022-12-16T21:16:04.000Z",
    "data": {
      "resourceId": "ocid1.applicationvip.oc1.....unique_id",
      "resourceName": "my_application_vip",
```

```
      "tagSlug": null,
      "compartmentId": "ocid1.compartment.oc1.....unique_id",
      "request": {
        "id": "23a08e08-6b1e-40f0-a027-f2601dfd44ea",
        "path": null,
        "action": null,
        "parameters": null,
        "headers": null
      },
      "response": {
        "status": null,
        "responseTime": null,
        "headers": null,
        "payload": null,
        "message": ""
      },
      "stateChange": {
        "previous": null,
        "current": {
          "lifecycleState": "TERMINATING",
          "hostnameLabel": "my_application_vip",
          "freeTags": {},
          "definedTags": {}
        }
      },
      "eventGroupingId": "csidb3f42d234534bc8bc8849b892e84/
fbd51970d2a2486f94671614b5ea0571/9DFE1BEB5433FF69BABCCB7E34F2EAF4",
      "eventName": "DeleteApplicationVip",
      "availabilityDomain": "",
      "resourceVersion": null,
      "additionalDetails": {
        "id": "ocid1.applicationvip.oc1.....unique_id",
        "freeformTags": {},
        "definedTags": {},
        "timeCreated": "2022-12-15T21:17:59.000Z",
        "timeUpdated": "2022-12-15T21:18:04.389Z",
        "lifecycleState": "TERMINATING",
        "lifecycleDetails": "",
        "hostnameLabel": "my_application_vip",
        "cloudVmClusterId": "ocid1.cloudvmcluster.oc1.....unique_id",
        "compartmentId": "ocid1.compartment.oc1.....unique_id",
        "vcnIpId": "ocid1.privateip.oc1.....unique_id",
        "ipAddress": "10.0.0.0",
        "subnetId": "ocid1.subnet.oc1.....unique_id",
        "networkType": "CLIENT"
      }
    }
  },
  "timeCreated": "2022-12-15T16:31:31.979Z"
}
```


This is a reference event for Application Virtual IP (VIP) - Delete End:


```
{
  "id":
```

```
"ocid1.eventschema.oc1.phx.9d1tjgkavhn0rq4qdlmofrjro9npvugu73dp07uht0igxs9732x
6yar1m5l5",
  "serviceName": "Database",
  "displayName": "Application Virtual IP (VIP) - Delete End",
  "eventType": "com.oraclecloud.databaseservice.deleteapplicationvip.end",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2022-12-15T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "id",
      "type": "string"
    },
    {
      "name": "definedTags",
      "type": [
        "null",
        "Map<String, Map<String, Object>>"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "Map<String, String>"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "hostnameLabel",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "cloudVmClusterId",
```

```
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "compartmentId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "vcnIpId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "ipAddress",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "subnetId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "networkType",
      "type": [
        "null",
        "string"
      ]
    }
  ],
  "exampleEvent": {
    "eventType": "com.oraclecloud.databaseservice.deleteapplicationvip.end",
    "cloudEventsVersion": "0.1",
    "eventTypeVersion": "2.0",
    "source": "databaseservice",
    "contentType": "application/json",
    "eventID": "17619ca1-07ae-4e2d-a818-5b5f1fcd4f70",
    "eventTime": "2022-12-16T21:16:04.000Z",
    "data": {
      "resourceId": "ocid1.applicationvip.oc1.....unique_id",
      "resourceName": "my_application_vip",
      "tagSlug": null,
      "compartmentId": "ocid1.compartment.oc1.....unique_id",
      "request": {
        "id": "1b0d242b-b3cd-4d61-9779-2de23e0e6742",
```

```
        "path": null,
        "action": null,
        "parameters": null,
        "headers": null
      },
      "response": {
        "status": null,
        "responseTime": null,
        "headers": null,
        "payload": null,
        "message": ""
      },
      "stateChange": {
        "previous": null,
        "current": {
          "lifecycleState": "TERMINATED",
          "hostnameLabel": "my_application_vip",
          "freeTags": {},
          "definedTags": {}
        }
      },
      "eventGroupingId": "csid80b16d4d459eaaa60ad25a9829d8/
b3e19f76a81549e6b7bf1d8619f7c191/C683214FCB0BF3CEC1C8B23C2FEE983E",
      "eventName": "DeleteApplicationVip",
      "availabilityDomain": "",
      "resourceVersion": null,
      "additionalDetails": {
        "id": "ocid1.applicationvip.oc1.....unique_id",
        "freeformTags": {},
        "definedTags": {},
        "timeCreated": "2022-12-15T21:17:59.000Z",
        "timeUpdated": "2022-12-15T21:18:04.389Z",
        "lifecycleState": "TERMINATED",
        "lifecycleDetails": "",
        "hostnameLabel": "my_application_vip",
        "cloudVmClusterId": "ocid1.cloudvmcluster.oc1.....unique_id",
        "compartmentId": "ocid1.compartment.oc1.....unique_id",
        "vcnIpId": "ocid1.privateip.oc1.....unique_id",
        "ipAddress": "10.0.0.0",
        "subnetId": "ocid1.subnet.oc1.....unique_id",
        "networkType": "CLIENT"
      }
    }
  },
  "timeCreated": "2022-12-15T16:31:31.979Z"
}
```


Interim Software Updates Event Types

These are the event types that Interim Software Updates in Oracle Cloud Infrastructure emit.

Friendly Name	Event Type
Oneoff Patch - Create Begin	`com.oraclecloud.databaseservice.createoneoffpatch.begin`

Friendly Name	Event Type
Oneoff Patch - Create End	com.oraclecloud.databaseservice.createoneoffpatch.end
Oneoff Patch - Delete Begin	com.oraclecloud.databaseservice.deleteoneoffpatch.begin
Oneoff Patch - Delete End	com.oraclecloud.databaseservice.deleteoneoffpatch.end
Oneoff Patch - Download Begin	com.oraclecloud.databaseservice.downloadoneoffpatch.begin
Oneoff Patch - Download End	com.oraclecloud.databaseservice.downloadoneoffpatch.end

Interim Software Updates Event Types Examples:

This is a reference event for This is a reference event for Oneoff Patch - Create Begin:


```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljrsllp7rfneajgq2knxbqopwux24za7qzoe3mfj2bzfxt
nwqcxpbcq",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.createoneoffpatch.begin",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID"
    },
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID",
      "compartmentName": "example_name",
      "resourceName": "my_oneoffpatch",
      "resourceId": "OneOffPatch-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2020-08-26T12:00:00.000Z",
        "displayName": "testDisplayName",
        "databaseVersion": "19.6.0.0",
        "patchSet": "test_patch_set"
      }
    }
  },
  "serviceName": "Database",
  "displayName": "Oneoff Patch - Create Begin",
  "eventType": "com.oraclecloud.databaseservice.createoneoffpatch.begin",
  "additionalDetails": [
    { "name": "id", "type": "string" },
```

```
    { "name": "lifecycleState", "type": "string" },
    { "name": "timeCreated", "type": "string" },
    { "name": "displayName", "type": "string" },
    { "name": "dbVersion", "type": "string" },
    { "name": "patchType", "type": "string" },
    { "name": "patchShapeFamily", "type": "string" },
    { "name": "releaseUpdate", "type": "string" }
  ],
  "timeCreated": "2020-06-26T13:31:31.979Z"
}
```

This is a reference event for Oneoff Patch - Create End:

```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljrj4vvuph4qvj5eateeel6axblhkq3caqndgmjvwl3sld
pgb255j2q",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.createoneoffpatch.end",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID"
    },
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID",
      "compartmentName": "example_name",
      "resourceName": "my_oneoffpatch",
      "resourceId": "OneOffPatch-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2020-08-26T12:00:00.000Z",
        "displayName": "testDisplayName",
        "databaseVersion": "19.6.0.0",
        "patchSet": "test_patch_set"
      }
    }
  },
  "serviceName": "Database",
  "displayName": "Oneoff Patch - Create End",
  "eventType": "com.oraclecloud.databaseservice.createoneoffpatch.end",
  "additionalDetails": [
    { "name": "id", "type": "string" },
    { "name": "lifecycleState", "type": "string" },
    { "name": "timeCreated", "type": "string" },
    { "name": "displayName", "type": "string" },
    { "name": "dbVersion", "type": "string" },
```

```
    { "name": "patchType", "type": "string" },
    { "name": "patchShapeFamily", "type": "string" },
    { "name": "releaseUpdate", "type": "string" }
  ],
  "timeCreated": "2020-06-26T13:31:31.979Z"
}
```

This is a reference event for Oneoff Patch - Delete Begin:

```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljrdripga5rryplwmv4ws6hqzr3pjyl7wfvoaqutvg2ey2
vtycn5onq",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.deleteoneoffpatch.begin",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID"
    },
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID",
      "compartmentName": "example_name",
      "resourceName": "my_oneoffpatch",
      "resourceId": "OneOffPatch-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2020-08-26T12:00:00.000Z",
        "displayName": "testDisplayName",
        "databaseVersion": "19.6.0.0",
        "patchSet": "test_patch_set"
      }
    }
  },
  "serviceName": "Database",
  "displayName": "Oneoff Patch - Delete Begin",
  "eventType": "com.oraclecloud.databaseservice.deleteoneoffpatch.begin",
  "additionalDetails": [
    { "name": "id", "type": "string" },
    { "name": "lifecycleState", "type": "string" },
    { "name": "timeCreated", "type": "string" },
    { "name": "displayName", "type": "string" },
    { "name": "dbVersion", "type": "string" },
    { "name": "patchType", "type": "string" },
    { "name": "patchShapeFamily", "type": "string" },
    { "name": "releaseUpdate", "type": "string" }
  ],
```

```
  "timeCreated": "2020-06-26T13:31:31.979Z"
}
```


This is a reference event for Oneoff Patch - Delete End:


```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljrgwk2gvx5lmx6fiwotgdy32mdmrnkyzznz37dpb4mmeh
gzt37vl7a",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.deleteoneoffpatch.end",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID"
    },
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID",
      "compartmentName": "example_name",
      "resourceName": "my_oneoffpatch",
      "resourceId": "OneOffPatch-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2020-08-26T12:00:00.000Z",
        "displayName": "testDisplayName",
        "databaseVersion": "19.6.0.0",
        "patchSet": "test_patch_set"
      }
    }
  },
  "serviceName": "Database",
  "displayName": "Oneoff Patch - Delete End",
  "eventType": "com.oraclecloud.databaseservice.deleteoneoffpatch.end",
  "additionalDetails": [
    { "name": "id", "type": "string" },
    { "name": "lifecycleState", "type": "string" },
    { "name": "timeCreated", "type": "string" },
    { "name": "displayName", "type": "string" },
    { "name": "dbVersion", "type": "string" },
    { "name": "patchType", "type": "string" },
    { "name": "patchShapeFamily", "type": "string" },
    { "name": "releaseUpdate", "type": "string" }
  ],
  "timeCreated": "2020-06-26T13:31:31.979Z"
}
```

This is a reference event for Oneoff Patch - Download Begin:


```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljr3vkb7klt5hkbsnqzjaxmszsqomanlbqmr2tsrcq7xaf
cv2c74l2q",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.downloadoneoffpatch.begin",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID"
    },
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID",
      "compartmentName": "example_name",
      "resourceName": "my_oneoffpatch",
      "resourceId": "OneOffPatch-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2020-08-26T12:00:00.000Z",
        "displayName": "testDisplayName",
        "databaseVersion": "19.6.0.0",
        "patchSet": "test_patch_set"
      }
    }
  },
  "serviceName": "Database",
  "displayName": "Oneoff Patch - Download Begin",
  "eventType": "com.oraclecloud.databaseservice.downloadoneoffpatch.begin",
  "additionalDetails": [
    { "name": "id", "type": "string" },
    { "name": "lifecycleState", "type": "string" },
    { "name": "timeCreated", "type": "string" },
    { "name": "displayName", "type": "string" },
    { "name": "dbVersion", "type": "string" },
    { "name": "patchType", "type": "string" },
    { "name": "patchShapeFamily", "type": "string" },
    { "name": "releaseUpdate", "type": "string" }
  ],
  "timeCreated": "2020-06-26T13:31:31.979Z"
}
```


This is a reference event for Oneoff Patch - Download End:


```
{
  "id":
```

```
"ocid1.eventschema.oc1.phx.abyhqljrn2lruez55ah56kqksi5qfg6m7igvven7o2qkahlk5tk
wrj5ll3oa",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.downloadoneoffpatch.end",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID"
    },
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID",
      "compartmentName": "example_name",
      "resourceName": "my_oneoffpatch",
      "resourceId": "OneOffPatch-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2020-08-26T12:00:00.000Z",
        "displayName": "testDisplayName",
        "databaseVersion": "19.6.0.0",
        "patchSet": "test_patch_set"
      }
    }
  },
  "serviceName": "Database",
  "displayName": "Oneoff Patch - Download End",
  "eventType": "com.oraclecloud.databaseservice.downloadoneoffpatch.end",
  "additionalDetails": [
    { "name": "id", "type": "string" },
    { "name": "lifecycleState", "type": "string" },
    { "name": "timeCreated", "type": "string" },
    { "name": "displayName", "type": "string" },
    { "name": "dbVersion", "type": "string" },
    { "name": "patchType", "type": "string" },
    { "name": "patchShapeFamily", "type": "string" },
    { "name": "releaseUpdate", "type": "string" }
  ],
  "timeCreated": "2020-06-26T13:31:31.979Z"
}
```


Serial Console Connection Event Types

Review the list of event types that serial console connection emits.

Table 6-4 Serial Console Connection Events

Friendly Name	Event Type
DB Node Console Connection - Create Begin	com.oraclecloud.databaseservice.created bnodeconsoleconnection.begin
DB Node Console Connection - Create End	com.oraclecloud.databaseservice.created bnodeconsoleconnection.end
DB Node Console Connection - Delete Begin	com.oraclecloud.databaseservice.deleted bnodeconsoleconnection.begin
DB Node Console Connection - Delete End	com.oraclecloud.databaseservice.deleted bnodeconsoleconnection.end
DB Node Console Connection - Update	com.oraclecloud.databaseservice.updated bnodeconsoleconnection
DB Node - Update	com.oraclecloud.databaseservice.updated bnode

Example 6-64 Serial Console Connection Event Types Examples

This is a reference event for DB Node Console Connection - Create Begin:

```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.createdbnodeconsoleconnection.begin",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID",
    "resourceId": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "CREATING",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocid1.dbnode.oc1..unique_ID",
      "tenantId": "ocid1.tenant.oc1..unique_ID",
      "compartmentId": "ocid1.compartment.oc1..unique_ID"
    }
  }
}
```

This is a reference event for DB Node Console Connection - Create End:


```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.createdbnodeconsoleconnection.end",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID",
    "resourceId": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "ACTIVE",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocid1.dbnode.oc1..unique_ID",
      "tenantId": "ocid1.tenant.oc1..unique_ID",
      "compartmentId": "ocid1.compartment.oc1..unique_ID"
    }
  }
}
```


This is a reference event for DB Node Console Connection - Delete Begin:


```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.deletedbnodeconsoleconnection.begin",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID",
    "resourceId": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "DELETING",
      "timeCreated": "2019-08-29T12:00:00.000Z",
```

```
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocid1.dbnode.oc1..unique_ID",
      "tenantId": "ocid1.tenant.oc1..unique_ID",
      "compartmentId": "ocid1.compartment.oc1..unique_ID"
    }
  }
}
```


This is a reference event for DB Node Console Connection - Delete End:


```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.deletedbnodeconsoleconnection.end",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID",
    "resourceId": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "DELETED",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocid1.dbnode.oc1..unique_ID",
      "tenantId": "ocid1.tenant.oc1..unique_ID",
      "compartmentId": "ocid1.compartment.oc1..unique_ID"
    }
  }
}
```


This is a reference event for DB Node Console Connection - Update:


```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.updatedbnodeconsoleconnection",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID"
  },
```

```
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID",
    "resourceId": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "ACTIVE",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocid1.dbnode.oc1..unique_ID",
      "tenantId": "ocid1.tenant.oc1..unique_ID",
      "compartmentId": "ocid1.compartment.oc1..unique_ID"
    }
  }
}
```


This is a reference event for DB Node - Update:


```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType": "com.oraclecloud.databaseservice.updatedbnode",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_dbnode",
    "resourceId": "DbNode-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.id..oc1...unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2019-08-26T12:00:00.000Z",
      "timeUpdated": "2019-08-26T12:30:00.000Z",
      "dbSystemId": "ocid1.dbsystem.oc1.phx.unique_ID",
      "lifecycleDetails": "detail message",
      "vmClusterId": "VmCluster-unique_ID",
      "dbHostId": "dbHost-unique_ID",
      "nodeNumber": 2,
      "powerAction": "HardReset",
      "hostName": "testHostName"
    }
  }
}
```

Viewing Audit Log Events

Oracle Cloud Infrastructure Audit service provides records of API operations performed against supported services as a list of log events.

An audit event is generated when you connect to the serial console using a Secure Shell (SSH) connection. Navigate to Audit in the Console and search for `VmConsoleConnected`. When you navigate to Audit in the Console, a list of results is generated for the current compartment. Audit logs are organized by compartment, so if you are looking for a particular event, you must know which compartment the event occurred in. You can filter the list in the following ways:

- Date and time
- Request Action Types (operations)
- Keywords

For more information, see *Viewing Audit Log Events*.

Example 6-65 Serial Console Connection Audit Event Example

This is a reference event for Serial Console Connection:


```
{

  "eventType": "VmConsoleConnected",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "VmConsoleConnectionAPI",
  "eventId": "2367d627-cff8-11ed-bfd3-02001714f979",
  "eventTime": "2023-03-31T19:13:37.120Z",
  "contentType": "application/json",

  "data": {
    "eventGroupingId": "2367d62d-cff8-11ed-bfd3-02001714f979",
    "eventName": "VmConsoleConnected",
    "compartmentId": "ocid1.compartment.oc1..<TRUNCATED>aaaaxxxxx",
    "compartmentName": "exacc-dev",
    "resourceName": "",
    "resourceId":
"ocid1.dbnodeconsoleconnection.oc1.iad.<TRUNCATED>aaaaaaxxxxx",
    "availabilityDomain": null,
    "freeformTags": null,
    "definedTags": null,

    "identity": {
      "principalName": "dsaes",
      "principalId": "ocid1.user.oc1..<TRUNCATED>aaaaaaaaaxxxxxxxxxx",
      "authType": "Native",
      "callerName": null,
      "callerId": null,
      "tenantId": "ocid1.tenancy.oc1..<TRUNCATED>aaaaaaxxxxx",
      "ipAddress": null,
```

```
      "credentials": null,
      "userAgent": null,
      "consoleSessionId": null
    },

    "request": {
      "id": "",
      "path": "",
      "action": "",
      "parameters": null,
      "headers": null
   },

    "response": {
      "status": "",
      "responseTime": "0001-01-01T00:00:00.000Z",
      "headers": null,
      "payload": null,
      "message": ""
    },

    "stateChange": null,

    "additionalDetails": {
      "DBNodeId": "ocid1.dbnode.oc1.iad.<TRUNCATED>aaaaaxxxxxxx"
    }
  }
}
```


Related Topics

- Overview of Audit
- Viewing Audit Log Events
- Setting Audit Log Retention Period

Policy Details for Exadata Cloud Infrastructure

This topic covers details for writing policies to control access to Exadata Cloud Infrastructure resources.

> **Note:**
>
> For more information on Policies, see "How Policies Work".
>
> For a sample policy, see "Let database admins manage Exadata Cloud Infrastructure instances".

- About Resource-Types
  Learn about resource-types you can use in your policies.
- Resource-Types for Exadata Cloud Service Instances

- Supported Variables
  Use variables when adding conditions to a policy.
- Details for Verb + Resource-Type Combinations
  Review the list of permissions and API operations covered by each verb.

Related Topics

- How Policies Work
- Let database admins manage Exadata Cloud Infrastructure instances

About Resource-Types

Learn about resource-types you can use in your policies.

An aggregate resource-type covers the list of individual resource-types that directly follow. For example, writing one policy to allow a group to have access to the `database-family` is equivalent to writing separate policies for the group that would grant access to the `cloud-exadata-infrastructures`, `cloud-vmclusters`, `db-nodes`, `db-homes`, `databases`, `database-software-image`, and `backups` resource-types. For more information, see Resource-Types.

Resource-Types for Exadata Cloud Service Instances

Aggregate Resource-Type
`database-family`

Individual Resource-Types:
`cloud-exadata-infrastructures`

`cloud-vmclusters`

`db-nodes`

`db-homes`

`databases`

`pluggable-databases`

`db-backups`

`application-vips`

`dbnode-console-connection`

Supported Variables

Use variables when adding conditions to a policy.

Exadata Cloud Infrastructure supports only the general variables. For more information, see "General Variables for All Requests".

Related Topics

- General Variables for All Requests

Details for Verb + Resource-Type Combinations

Review the list of permissions and API operations covered by each verb.

For more information, see "Permissions", "Verbs", and "Resource-Types".

- Database-Family Resource Types
  Understand the level of access of each verb.
- cloud-exadata-infrastructures
  Review the list of permissions and API operations for `cloud-exadata-infrastructures` resource-type.
- cloud-vmclusters
  Review the list of permissions and API operations for `cloud-vmclusters` resource-type.
- db-nodes
  Review the list of permissions and API operations for `db-nodes` resource-type.
- dbnode-console-connection
  Review the list of permissions and API operations for `dbnode-console-connection` resource-type.
- db-homes
  Review the list of permissions and API operations for `db-homes` resource-type.
- dbServers
  Review the list of permissions and API operations for `dbServers` resource-type.
- database-software-images
  Review the list of permissions and API operations for `database-software-images` resource-type.
- pluggable-databases (PDBs)
  Review the list of permissions and API operations for `pluggable-databases` resource-type.
- databases (CDBs)
  Review the list of permissions and API operations for `databases` resource-type.
- db-backups
  Review the list of permissions and API operations for `db-backups` resource-type.
- data-guard-association
  Review the list of permissions and API operations for `data-guard-association` resource-type.
- key-stores
  Review the list of permissions and API operations for `key-store` resource-type.
- application-vips
  Review the list of permissions and API operations for `application-vips` resource-type.
- oneoffPatch
  Review the list of permissions and API operations for `oneoffPatch` resource-type.
- Permissions Required for Each API Operation
  The following tables list the API operations for Exadata Cloud Infrastructure instances in a logical order, grouped by resource type.

Related Topics

- Permissions
- Verbs
- Resource-Types

Database-Family Resource Types

Understand the level of access of each verb.

The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

For example, the `read` verb for the `vmclusters` resource-type covers no extra permissions or API operations compared to the `inspect` verb. However, the `use` verb includes one more permission, fully covers one more operation, and partially covers another additional operation.

cloud-exadata-infrastructures

Review the list of permissions and API operations for `cloud-exadata-infrastructures` resource-type.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	`CLOUD_EXADATA_INFRASTRUCTURE_INSPECT`	`ListCloudExadataInfrastructures` `GetCloudExadataInfrastructures`	none
read	*no extra*	*no extra*	none
use	`CLOUD_EXADATA_INFRASTRUCTURE_UPDATE`	*no extra*	`ChangeCloudExadataInfrastructureCompartment` (also needs `use cloud-vmclusters`, `use db-homes`, `use databases`, and `inspect db-backups`)

manage	USE +
	`CLOUD_EXADATA_INFRASTRUCTURE_CREATE`
	`CLOUD_EXADATA_INFRASTRUCTURE_DELETE`
		`UpdateCloudExadataInfrastructure`	`CreateCloudExadataInfrastructure`, `DeleteCloudExadataInfrastructure`, `AddStorageCapacityCloudExadataInfrastructure` (also needs `use cloud-vmclusters`)

cloud-vmclusters

Review the list of permissions and API operations for `cloud-vmclusters` resource-type.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	CLOUD_VM_CLUSTER_INSPECT	ListCloudVmClusters GetCloudVmCluster ListCloudVmClusterUpdates ListCloudVmClusterUpdateHistoryEntries GetCloudVmClusterUpdate GetCloudVmClusterUpdateHistoryEntry	*none*
read	*no extra*	*no extra*	*none*
use	CLOUD_VM_CLUSTER_UPDATE	*no extra*	ChangeCloudVmClusterCompartment (also needs use db-homes, use databases, and inspect db-backups)
manage	USE + CLOUD_VM_CLUSTER_CREATE CLOUD_VM_CLUSTER_DELETE	UpdateCloudVmCluster	CreateCloudVmCluster, DeleteCloudVmCluster (both also need manage db-homes, manage databases, use vnics, and use subnets); RemoveVmFromCloudVmCluster, AddVmToCloudVmCluster (both also need use cloud_exadata_infrastructure_update

db-nodes

Review the list of permissions and API operations for db-nodes resource-type.

> **Note:**
>
> For Exadata Cloud Infrastructure VM clusters, the database node is sometimes referred to as a virtual machine.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DB_NODE_INSPECT DB_NODE_QUERY	GetDbNode	*none*
read	*no extra*	*no extra*	*none*
use	DB_NODE_UPDATE	UpdateDbNode	*none*

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
manage	*USE* + `DB_NODE_POWER_ACTIONS`	`DbNodeAction`	*none*

dbnode-console-connection

Review the list of permissions and API operations for `dbnode-console-connection` resource-type.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	`DBNODE_CONSOLE_CONNECTION_INSPECT`	`GetDbNodeConsoleConnection` `ListDbNodeConsoleConnections`	*none*
read	*no extra*	*no extra*	*none*

use	*READ* +
	`DBNODE_CONSOLE_CONNECTION_UPDATE`
	`PLUGGABLE_DATABASE_UPDATE`
		`UpdateDbNodeConsoleConnection`	*none*

manage	*USE* +
	`DBNODE_CONSOLE_CONNECTION_CREATE`
	`DBNODE_CONSOLE_CONNECTION_DELETE`
		`CreateDbNodeConsoleConnection`
		`DeleteDbNodeConsoleConnection`
			none

db-homes

Review the list of permissions and API operations for `db-homes` resource-type.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered

inspect	`DB_HOME_INSPECT`	`ListDBHome`
		`GetDBHome`
		`ListDbHomePatches`
		`ListDbHomePatchHistoryEntries`
		`GetDbHomePatch`
		`GetDbHomePatchHistoryEntry`
			none

read	*no extra*	*no extra*	*none*
use	`DB_HOME_UPDATE`	`UpdateDBHome`	`ChangeCloudVmClusterCompartment` (also needs `use cloud-vmclusters`, `use databases`, and `inspect backups`)

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
manage	USE + DB_HOME_CREATE DB_HOME_DELETE	no extra	CreateCloudVmCluster, DeleteCloudVmCluster (both also need manage cloud-vmclusters, manage databases, use vnics, and use subnets). If automatic backups are enabled on the default database, also needs manage backups CreateDbHome, (also needs use cloud-vmclusters and manage databases). If creating the Database Home by restoring from a backup, also needs read backups DeleteDbHome, (also needs use cloud-vmclusters and manage databases). If automatic backups are enabled on the default database, also needs manage backups. If the performFinalBackup option is selected, also needs manage backups and read databases.

dbServers

Review the list of permissions and API operations for dbServers resource-type.

Table 6-5 INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
EXADATA_INFRASTRUCTURE_INSPECT	*none*	GetDbServer ListDbServers

Table 6-6 READ

Permissions	APIs Fully Covered	APIs Partially Covered
No extra	*none*	*none*

Table 6-7 USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ + `VM_CLUSTER_UPDATE` `EXADATA_INFRASTRUCTURE_UPDATE`	*none*	`AddVirtualMachineToVmCluster`, `RemoveVirtualMachineFromVmCluster`

Table 6-8 MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
No extra	*none*	*none*

database-software-images

Review the list of permissions and API operations for `database-software-images` resource-type.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	`DB_SOFTWARE_IMG_INSPECT`	`ListDatabaseSoftwareImages` `GetDatabaseSoftwareImage`	*none*
read	no extra	none	*none*
use	*READ* + `DB_SOFTWARE_IMG_UPDATE`	`UpdateDatabaseSoftwareImage` `ChangeDatabaseSoftwareImageCompartment`	*none*
manage	*USE* + `DB_SOFTWARE_IMG_CREATE` `DB_SOFTWARE_IMG_DELETE`	`CreateDatabaseSoftwareImage` `DeleteDatabaseSoftwareImage`	*none*

pluggable-databases (PDBs)

Review the list of permissions and API operations for `pluggable-databases` resource-type.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	PLUGGABLE_DATABASE_INSPECT	ListPluggableDatabases GetPluggableDatabase	UpdatePluggableDatabase StartPluggableDatabase StopPluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase RefreshPluggableDatabase ConvertRefreshablePluggableDatabase
	DATABASE_INSPECT	*no extra*	CreatePluggableDatabase DeletePluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase
read	*INSPECT* + PLUGGABLE_DATABASE_CONTENT_READ	*no extra*	CreatePluggableDatabase (Additional permissions are required if auto-backups are enabled on the CDB and includes this PDB.) UpdatePluggableDatabase (Additional permissions are required if auto-backups are enabled on the CDB and includes this PDB.) LocalClonePluggableDatabase RemoteClonePluggableDatabase
use	*READ* + PLUGGABLE_DATABASE_CONTENT_WRITE	*no extra*	LocalClonePluggableDatabase RemoteClonePluggableDatabase

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
	PLUGGABLE_DATABASE_UPDATE	*no extra*	UpdatePluggableDatabase StartPluggableDatabase StopPluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase RefreshPluggableDatabase ConvertRefreshablePluggableDatabase
	DATABASE_UPDATE	*no extra*	CreatePluggableDatabase DeletePluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase
manage	*USE* + PLUGGABLE_DATABASE_CREATE	*no extra*	CreatePluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase
	PLUGGABLE_DATABASE_DELETE	*no extra*	DeletePluggableDatabase

databases (CDBs)

Review the list of permissions and API operations for `databases` resource-type.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DATABASE_INSPECT	ListDatabases GetDatabase ListDataGuardAssociations GetDataGuardAssociation	enableDatabaseManagement disableDatabaseManagement updateDatabaseManagement
read	*INSPECT*+ DATABASE_CONTENT_READ	*no extra*	*no extra*

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
use	*READ* + DATABASE_CONTENT_WRITE DATABASE_UPDATE	UpdateDatabase SwitchoverDataGuardAssociation FailoverDataGuardAssociation ReinstateDataGuardAssociation	CreateDataGuardAssociation ChangeCloudVmClusterCompartment (also needs use cloud-vmclusters, use db-homes, and inspect db-backups) enableDatabaseManagement disableDatabaseManagement updateDatabaseManagement
manage	*USE* + DATABASE_CREATE DATABASE_DELETE	*no extra*	CreateDatabase (also needs use cloud-vmclusters, use db-homes, and if automatic backups to be enabled, also needs manage backups) DeleteDatabase (also needs use cloud-vmclusters, use db-homes, and if automatic backups to be enabled, also needs manage backups) CreateCloudVmCluster, DeleteCloudVmCluster (both also need manage cloud-vmclusters, manage db-homes, use vnics, and use subnets)

db-backups

Review the list of permissions and API operations for db-backups resource-type.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DB_BACKUP_INSPECT	GetBackup ListBackups	ChangeCloudVmClusterCompartment (also needs use cloud-vmclusters, use db-homes, and use databases)

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
read	*INSPECT* + `DB_BACKUP_CONTENT_READ`	*none*	`RestoreDatabase` (also needs `use databases`)
use	*no extra*	*no extra*	*none*
manage	*USE* + `DB_BACKUP_CREATE` `DB_BACKUP_DELETE`	`DeleteBackup`	`CreateBackup` (also needs `read databases`)

data-guard-association

Review the list of permissions and API operations for `data-guard-association` resource-type.

Table 6-9 INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
`DATABASE_INSPECT`	`ListDataGuardAssociations`, `GetDataGuardAssociation`	`CreateDataGuardAssociation`

Table 6-10 READ

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	*no extra*	*no extra*

Table 6-11 USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ + `VM_CLUSTER_UPDATE` + `DB_HOME_UPDATE` `DATABASE_UPDATE`	`DeleteDatabase` `SwitchoverDataGuardAssociation`,`FailoverDataGuardAssociation`, `ReinstateDataGuardAssociation`	`CreateDataGuardAssociation`

Table 6-12 MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE + `DATABASE_DELETE`	`DeleteDatabase`	*none*

key-stores

Review the list of permissions and API operations for `key-store` resource-type.

Table 6-13 INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
KEY_STORE_INPSECT	GetKeyStore	ChangeKeyStoreCompartment
AUTONOMOUS_CONTAINER_DATABASE_INSPECT	GetAutonomousContainerDatabase	RotateAutonomousContainerDatabaseKey
AUTONOMOUS_DATABASE_INSPECT	GetAutonomousDatabase	
AUTONOMOUS_DB_BACKUP_INSPECT	GetAutonomousDatabaseBackup	

Table 6-14 READ

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	*no extra*	*no extra*

Table 6-15 USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ + KEY_STORE_UPDATE +	UpdateKeyStore	ChangeKeyStoreCompartment
AUTONOMOUS_VM_CLUSTER_UPDATE +	*none*	CreateAutonomousContainerDatabase
AUTONOMOUS_CONTAINER_DATABASE_UPDATE	*none*	RotateAutonomousContainerDatabaseKey
AUTONOMOUS_DATABASE_UPDATE	RotateAutonomousDatabaseKey	*none*

Table 6-16 MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE + KEY_STORE_CREATE +	CreateKeyStore	*none*
KEY_STORE_DELETE +	DeleteKeyStore	*none*
AUTONOMOUS_CONTAINER_DATABASE_CREATE	CreateAutonomousContainerDatabase	*none*

application-vips

Review the list of permissions and API operations for `application-vips` resource-type.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	APPLICATION_VIP_INSPECT	ListApplicationVips GetApplicationVips	*none*
read	*INSPECT* +	*no extra*	*none*
use	*READ* +	*no extra*	*none*

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
manage	*USE* + APPLICATION_VIP_CREATE APPLICATION_VIP_DELETE	CreateApplicationVip DeleteApplicationVip	*none*

oneoffPatch

Review the list of permissions and API operations for `oneoffPatch` resource-type.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	ONEOFF_PATCH_INSPECT	DownloadOneoffPatch GetOneoffPatch ListOneoffPatches	CreateOneoffPatch DeleteOneoffPatch UpdateOneoffPatch ChangeOneoffPatchCompartment
read	*INSPECT* + *no extra*	DownloadOneoffPatch	*none*
use	*READ* + ONEOFF_PATCH_UPDATE	*no extra*	UpdateOneoffPatch ChangeOneoffPatchCompartment
manage	*USE* + ONEOFF_PATCH_CREATE ONEOFF_PATCH_DELETE	*no extra*	CreateOneoffPatch DeleteOneoffPatch

Related Topics

- OneoffPatch Reference

Permissions Required for Each API Operation

The following tables list the API operations for Exadata Cloud Infrastructure instances in a logical order, grouped by resource type.

Database API Operations

For information about permissions, see:

Permissions.

The following tables list of API operations and permissions by API peration.

Table 6-17 Cloud Exadata Infrastructure Resource

API Operation	Permissions Required to Use the Operation
ListCloudExadataInfrastructures	CLOUD_EXADATA_INFRASTRUCTURE_INSPECT
GetCloudExadataInfrastructure	CLOUD_EXADATA_INFRASTRUCTURE_INSPECT
CreateCloudExadataInfrastructure	CLOUD_EXADATA_INFRASTRUCTURE_CREATE

Table 6-17 (Cont.) Cloud Exadata Infrastructure Resource

API Operation	Permissions Required to Use the Operation
UpdateCloudExadataInfrastructure	CLOUD_EXADATA_INFRASTRUCTURE_UPDATE
ChangeCloudExadataInfrastructureCompartment	CLOUD_EXADATA_INFRASTRUCTURE_UPDATE
DeleteCloudExadataInfrastructure	CLOUD_EXADATA_INFRASTRUCTURE_DELETE
AddStorageCapacityCloudExadataInfrastructure	CLOUD_EXADATA_INFRASTRUCTURE_UPDATE

Table 6-18 Cloud VM Cluster

API Operation	Permissions Required to Use the Operation
ListCloudVmClusters	CLOUD_VM_CLUSTER_INSPECT
GetCloudVmCluster	CLOUD_VM_CLUSTER_INSPECT
CreateCloudVmCluster	CLOUD_VM_CLUSTER_CREATE and CLOUD_EXADATA_INFRASTRUCTURE_UPDATE and VNIC_CREATE and VNIC_ATTACH and SUBNET_ATTACH and (needed if Private DNS is used: DNS_ZONE_READ, DNS_RECORD_UPDATE, DNS_ZONE_CREATE DNS_VIEW_INSPECT)
ChangeCloudVmClusterCompartment	CLOUD_VM_CLUSTER_UPDATE
UpdateCloudVmCluster	CLOUD_VM_CLUSTER_UPDATE and CLOUD_EXADATA_INFRASTRUCTURE_UPDATE
GetCloudVmClusterIormConfig	CLOUD_VM_CLUSTER_INSPECT
UpdateCloudVmClusterIormConfig	CLOUD_VM_CLUSTER_UPDATE
DeleteCloudVmCluster	CLOUD_VM_CLUSTER_DELETE and CLOUD_EXADATA_INFRASTRUCTURE_UPDATE and DB_HOME_DELETE and VNIC_DELETE and SUBNET_DETACH and VNIC_DETACH and (needed if Private DNS is used: DNS_ZONE_READ, DNS_RECORD_UPDATE, DNS_ZONE_DELETE)
AddVmToCloudVmCluster	CLOUD_VM_CLUSTER_UPDATE and CLOUD_EXADATA_INFRASTRUCTURE_UPDATE and (needed if Private DNS is used: DNS_ZONE_READ, DNS_RECORD_UPDATE, DNS_ZONE_CREATE, DNS_VIEW_INSPECT)
RemoveVmFromCloudVmCluster	CLOUD_VM_CLUSTER_UPDATE and CLOUD_EXADATA_INFRASTRUCTURE_UPDATE and (needed if Private DNS is used: DNS_ZONE_READ, DNS_RECORD_UPDATE, DNS_ZONE_DELETE)

Table 6-19 Cloud VM Cluster Maintenance Updates and Update History

API Operation	Permissions Required to Use the Operation
ListCloudVmClusterUpdates	CLOUD_VM_CLUSTER_INSPECT
GetCloudVmClusterUpdate	CLOUD_VM_CLUSTER_INSPECT
ListCloudVmClusterUpdateHistoryEntries	CLOUD_VM_CLUSTER_INSPECT

Table 6-19 (Cont.) Cloud VM Cluster Maintenance Updates and Update History

API Operation	Permissions Required to Use the Operation
GetCloudVmClusterUpdateHistoryEntry	CLOUD_VM_CLUSTER_INSPECT

Table 6-20 Virtual Machines / Nodes

API Operation	Permissions Required to Use the Operation
ListDbNodes	DB_NODE_INSPECT
GetDbNode	DB_NODE_INSPECT
DbNodeAction	DB_NODE_POWER_ACTIONS

Table 6-21 Database Homes

API Operation	Permissions Required to Use the Operation
ListDbHomes	DB_HOME_INSPECT
GetDbHome	DB_HOME_INSPECT
ListDbHomePatches	DB_HOME_INSPECT
ListDbHomePatchHistoryEntries	DB_HOME_INSPECT
GetDbHomePatch	DB_HOME_INSPECT
GetDbHomePatchHistoryEntry	DB_HOME_INSPECT
CreateDbHome	DB_SYSTEM_INSPECT and DB_SYSTEM_UPDATE and DB_HOME_CREATE and DATABASE_CREATE To enable automatic backups for the database, also need DB_BACKUP_CREATE and DATABASE_CONTENT_READ
UpdateDbHome	DB_HOME_UPDATE
DeleteDbHome	DB_SYSTEM_UPDATE and DB_HOME_DELETE and DATABASE_DELETE If automatic backups are enabled, also need DELETE_BACKUP If performing a final backup on termination, also need DB_BACKUP_CREATE and DATABASE_CONTENT_READ

Table 6-22 Databases (CDB)

API Operation	Permissions Required to Use the Operation
ListDatabases	DATABASE_INSPECT
GetDatabase	DATABASE_INSPECT
CreateDatabase	DATABASE_UPDATE To enable automatic backups, also need DB_BACKUP_CREATE and DATABASE_CONTENT_READ

Table 6-22 (Cont.) Databases (CDB)

API Operation	Permissions Required to Use the Operation
UpdateDatabase	DATABASE_UPDATE To enable automatic backups, also need DB_BACKUP_CREATE and DATABASE_CONTENT_READ
DeleteDatabase	For new resource model using VM cluster resource: CLOUD_VM_CLUSTER_INSPECT and DB_HOME_UPDATE and DATABASE_DELETE
enableDatabaseManagement	DATABASE_INSPECT and DATABASE_UPDATE
disableDatabaseManagement	DATABASE_INSPECT and DATABASE_UPDATE
disableDatabaseManagement	DATABASE_INSPECT and DATABASE_UPDATE

Table 6-23 Pluggable Databases (PBDs)

API Operation	Permissions Required to Use the Operation
ListPluggableDatabase	PLUGGABLE_DATABASE_INSPECT
GetPluggableDatabase	PLUGGABLE_DATABASE_INSPECT
CreatePluggableDatabase	PLUGGABLE_DATABASE_CREATE and DATABASE_INSPECT and DATABASE_UPDATE
UpdatePluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE
StartPluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE
StopPluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE
DeletePluggableDatabase	PLUGGABLE_DATABASE_DELETE and DATABASE_INSPECT and DATABASE_UPDATE
LocalClonePluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE and PLUGGABLE_DATABASE_CONTENT_READ and PLUGGABLE_DATABASE_CONTENT_WRITE and PLUGGABLE_DATABASE_CREATE and DATABASE_INSPECT and DATABASE_UPDATE
RemoteClonePluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE and PLUGGABLE_DATABASE_CONTENT_READ and PLUGGABLE_DATABASE_CONTENT_WRITE and PLUGGABLE_DATABASE_CREATE and DATABASE_INSPECT and DATABASE_UPDATE
enableDatabaseManagement	DATABASE_INSPECT and DATABASE_UPDATE
disableDatabaseManagement	DATABASE_INSPECT and DATABASE_UPDATE
disableDatabaseManagement	DATABASE_INSPECT and DATABASE_UPDATE

Table 6-24 System Shapes and Database Versions

API Operation	Permissions Required to Use the Operation
ListDbSystemShapes	(no permissions required; available to anyone)
ListDbVersions	(no permissions required; available to anyone)

Table 6-25 Oracle Data Guard Associations

API Operation	Permissions Required to Use the Operation
GetDataGuardAssociation	DATABASE_INSPECT
ListDataGuardAssociations	DATABASE_INSPECT
CreateDataGuardAssociation	DB_SYSTEM_UPDATE and DB_HOME_CREATE and DB_HOME_UPDATE and DATABASE_CREATE and DATABASE_UPDATE
SwitchoverDataGuardAssociation	DATABASE_UPDATE
FailoverDataGuardAssociation	DATABASE_UPDATE
ReinstateDataGuardAssociation	DATABASE_UPDATE

Table 6-26 Backups and Database Restore

API Operation	Permissions Required to Use the Operation
GetBackup	DB_BACKUP_INSPECT
ListBackups	DB_BACKUP_INSPECT
CreateBackup	DB_BACKUP_CREATE and DATABASE_CONTENT_READ
DeleteBackup	DB_BACKUP_DELETE and DB_BACKUP_INSPECT
RestoreDatabase	DB_BACKUP_INSPECT and DB_BACKUP_CONTENT_READ and DATABASE_CONTENT_WRITE

Table 6-27 Application VIP

API Operation	Permissions Required to Use the Operation
CreateApplicationVip	APPLICATION_VIP_CREATE and CLOUD_VM_CLUSTER_UPDATE and PRIVATE_IP_CREATE and PRIVATE_IP_ASSIGN and VNIC_ASSIGN and SUBNET_ATTACH
DeleteApplicationVip	APPLICATION_VIP_DELETE and CLOUD_VM_CLUSTER_UPDATE and PRIVATE_IP_DELETE and PRIVATE_IP_UNASSIGN and VNIC_UNASSIGN and SUBNET_DETACH
ListApplicationVips	APPLICATION_VIP_INSPECT
ListApplicationVips	APPLICATION_VIP_INSPECT

Table 6-28 Serial Console Access to VM

API Operation	Permissions Required to Use the Operation
AddVirtualMachineToVmCluster	VM_CLUSTER_UPDATE and EXADATA_INFRASTRUCTURE_UPDATE
RemoveVirtualMachineFromVmCluster	VM_CLUSTER_UPDATE and EXADATA_INFRASTRUCTURE_UPDATE
CreateDbNodeConsoleConnection	DBNODE_CONSOLE_CONNECTION_CREATE and DBNODE_CONSOLE_CONNECTION_INSPECT
GetDbNodeConsoleConnection	DBNODE_CONSOLE_CONNECTION_INSPECT
ListDbNodeConsoleConnections	DBNODE_CONSOLE_CONNECTION_INSPECT
DeleteDbNodeConsoleConnection	DBNODE_CONSOLE_CONNECTION_DELETE
UpdateDbNodeConsoleConnection	DBNODE_CONSOLE_CONNECTION_UPDATE
UpdateDbNode	DB_NODE_UPDATE

Managing Exadata Resources with Oracle Enterprise Manager Cloud Control

To manage and monitor Exadata Cloud Infrastructure and Exadata Database Service on Cloud@Customer resources, use Oracle Enterprise Manager Cloud Control.

For complete documentation and Oracle By Example tutorials, see the following documentation resources: *Oracle Enterprise Manager Cloud Control for Oracle Exadata Cloud* and *Setting Up Oracle Enterprise Manager 13.4 on Oracle Cloud Infrastructure*.

- Overview of Oracle Enterprise Manager Cloud Control
  Oracle Enterprise Manager Cloud Control provides a complete lifecycle management solution for Oracle Cloud Infrastructure's Exadata Cloud Infrastructure (ExaDB-D) and Exadata Database Service on Cloud@Customer (ExaDB-C@C) services.
- Features of Enterprise Manager Cloud Control
  Familiarize yourself with the features of Enterprise Manager Cloud Control to manage and monitor Exadata Cloud and Exadata Cloud@Customer resources.
- Analyzing Exadata Database Service Database Performance
  This topic describes how to use Database Metrics and Performance Hub to monitor, analyze, and tune the performance of OCI user-managed databases, including Oracle Exadata Database Service on Dedicated Infrastructure databases and databases running on virtual machine and bare metal systems.

Related Topics

- Oracle Enterprise Manager Cloud Control for Oracle Exadata Cloud
- Setting Up Oracle Enterprise Manager 13.4 on Oracle Cloud Infrastructure

Overview of Oracle Enterprise Manager Cloud Control

Oracle Enterprise Manager Cloud Control provides a complete lifecycle management solution for Oracle Cloud Infrastructure's Exadata Cloud Infrastructure (ExaDB-D) and Exadata Database Service on Cloud@Customer (ExaDB-C@C) services.

Enterprise Manager Cloud Control discovers ExaDB-D and ExaDB-C@C services as a single target and automatically identifies and organizes all dependent components. Using Enterprise Manager Cloud Control you can then:

- Monitor and manage all Exadata, ExaDB-D and ExaDB-C@C systems, along with any other targets, from a single interface
- Visualize storage and compute data
- View performance metrics of your Exadata components

Features of Enterprise Manager Cloud Control

Familiarize yourself with the features of Enterprise Manager Cloud Control to manage and monitor Exadata Cloud and Exadata Cloud@Customer resources.

Enterprise Manager Target for Exadata Cloud

The target for Oracle Cloud Infrastructure Exadata resources, which covers both Exadata Cloud and Exadata Cloud@Customer does the following:

- Automatically identifies and organizes related targets.
- Provides a high-level integration point for Enterprise Manager framework features such as incident rules, groups, notifications, and monitoring templates.

Improved Performance Monitoring

Enterprise Manager Cloud Control enhances performance monitoring in the following ways:

- Adds Exadata Storage Server and Exadata Storage Grid targets.
- Offers visualization of storage and compute performance for your Exadata Cloud and Exadata Cloud@Customer resources.
- Enables use of the same Maximum Availability Architecture (MAA) key performance indicators (KPI) developed for Oracle Exadata Database Machine.

Scripted CLI-based Discovery

Enterprise Manager Cloud Control uses scripts to discover Oracle Cloud Infrastructure Exadata resources. Scripts search the existing hosts, clusters, ASM, databases and related targets, and add the storage server targets.

"Single Pane of Glass" View of On-Premises and Oracle Cloud Infrastructure Exadata Resources

Enterprise Manager Cloud Control 's use of a single Exadata target type provides a consistent Enterprise Manager experience across on-premises, Exadata Cloud, and Exadata Cloud@Customer resources. The common Exadata target menu allows you to easily navigate to, monitor and manage all of your Exadata systems.

Visualization

Enterprise Manager Cloud Control allows you to visualize the database and related targets associated with each Exadata Cloud and Exadata Cloud@Customer system.

Analyzing Exadata Database Service Database Performance

This topic describes how to use Database Metrics and Performance Hub to monitor, analyze, and tune the performance of OCI user-managed databases, including Oracle Exadata Database Service on Dedicated Infrastructure databases and databases running on virtual machine and bare metal systems.

With this tool, you can view real-time and historical performance data. For information about using Performance Hub, see Using Performance Hub to Analyze Database Performance.

To use Database Metrics and Performance Hub for Exadata Cloud Infrastructure, Virtual Machine, and Bare Metal databases, Database Management must be enabled for the database. When enabling a database, the database administrator can choose from two database management options: Basic Management and Full Management. For information about using Database Metrics and Performance Hub with Virtual Machine, Bare Metal, Exadata Cloud Infrastructure and external databases, see Enable Database Management.

> **Note:**
>
> Using Identity and Access Management (IAM), you can create a policy that grants users access to Performance Hub while limiting actions they can take on Autonomous Databases, databases running on virtual machine and bare metal systems, Oracle Database Cloud Service, Exadata Cloud Infrastructure, and external databases. For information about IAM policies and ExaDB-D databases, see *Required IAM Policy* . For information about policies and how to use them, see How Policies Work.

Related Topics

- Required IAM Policy for Exadata Cloud Infrastructure
  Review the identity access management (IAM) policy for provisioning Oracle Exadata Database Service on Dedicated Infrastructure systems.

Observability and Management for Exadata Database Service on Dedicated Infrastructure

- Metrics for Exadata Cloud Infrastructure in the Database Management Service
  Database Management provides comprehensive database performance diagnostics and management capabilities for Oracle Databases.
- Oracle Cloud Infrastructure Operations Insights
  Oracle Cloud Infrastructure Operations Insights allows you to use the Capacity Planning and SQL Warehouse functionality to gain insight into Oracle Databases deployed in Oracle Cloud (Bare Metal, Virtual Machine VM, and Exadata Cloud Infrastructure).
- Monitor Metrics to Diagnose and Troubleshoot Problems with Pluggable Databases
  Enable Database Management service to view metrics to diagnose and troubleshoot problems with pluggable databases.

Metrics for Exadata Cloud Infrastructure in the Database Management Service

Database Management provides comprehensive database performance diagnostics and management capabilities for Oracle Databases.

This article describes the metrics emitted by the Exadata Cloud Infrastructure Database service in the `oracle_oci_database` namespace for Oracle Databases.

To use database metrics for these Oracle Databases in Exadata Cloud Infrastructure Database Service, you must enable Database Management for the database you want to monitor. You can enable either Basic Management or Full Management for your database. See Enable Database Management for instructions.

Dimensions

All the metrics discussed in this topic include the following dimensions.

- RESOURCEID - The OCID of the database.
- RESOURCENAME - The name of the database.
- DEPLOYMENTTYPE - The deployment type of the database.

Note: Valid alarm intervals are 5 minutes or greater due to the frequency at which these metrics are emitted. See To create an alarm for details on creating alarms.

The database metrics can be provided for the basic and full Database Management options.

NOT_SUPPORTED

The metrics listed in the following table are automatically available for Oracle Databases when the **Basic Database Management** option is enabled.

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
`BlockChanges`	**DB Block Changes**	changes per second	The average number of blocks changed per second. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
`CpuUtilization`	**CPU Utilization**	percent	The CPU utilization expressed as a percentage, aggregated across all consumer groups. The utilization percentage is reported with respect to the number of CPUs the database is allowed to use, which is two times the number of OCPUs. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
`CurrentLogons`	**Current Logons**	count	The number of successful logons during the selected interval. Statistics: Sum Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
`ExecuteCount`	**Execute Count**	count	The number of user and recursive calls that executed SQL statements during the selected interval. Statistic: Sum Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
`OcpusAllocated`	**OCPU Allocated**	integer	The actual number of OCPUs allocated by the service during the selected interval of time. Statistic: Count Interval: 1 minute	5 minutes	N/A

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
ParseCount	**Parse Count (Total)**	count	The number of hard and soft parses during the selected interval. Statistic: Sum Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
StorageAllocated	**Allocated Storage Space**	GB	The maximum amount of space allocated by tablespace during the interval. For container databases, this metric provides data for root container tablespaces. Statistic: Max Interval: 30 minutes	30 minutes	N/A
StorageAllocatedByTablespace	**Allocated Storage Space By Tablespace**	GB	The maximum amount of space allocated by tablespace during the interval. For container databases, this metric provides data for root container tablespaces. Statistic: Max Interval: 30 minutes	30 minutes	tablespaceName tablespaceType
StorageUsed	**Storage Space Used**	GB	The maximum amount of space used during the interval. Statistic: Max Interval: 30 minutes	30 minutes	N/A

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
`StorageUsedByTablespace`	**Storage Space Used By Tablespace**	GB	The maximum amount of space used by tablespace during the interval. For container databases, this metric provides data for root container tablespaces. Statistic: Max Interval: 30 minutes	30 minutes	tablespaceName tablespaceType
`StorageUtilization`	**Storage Utilization**	percent	The percentage of provisioned storage capacity currently in use. Represents the total allocated space for all tablespaces. Statistic: Mean Interval: 30 minutes	30 minutes	N/A
`StorageUtilizationByTablespace`	**Storage Space Utilization By Tablespace**	percent	The percentage of the space utilized, by tablespace. For container databases, this metric provides data for root container tablespaces. Statistic: mean Interval: 30 minutes	30 minutes	tablespaceName tablespaceType
`TransactionCount`	**Transaction Count**	count	The combined number of user commits and user rollbacks during the selected interval. Statistic: Sum Interval: 1 minute	5 minutes	instanceNumber instanceName hostName

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
UserCalls	**User Calls**	count	The combined number of logons, parses, and execute calls during the selected interval. Statistic: Sum Interval: 1 minute	5 minutes	instanceNumber instanceName hostName

NOT_SUPPORTED

The metrics listed in the following table are automatically available for Oracle Databases when the **Full Database Management** option is enabled.

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
AllocatedStorageUtilizationByTablespace	**Allocated Space Utilization By Tablespace**	percent	The percentage of space used by tablespace, out of all allocated. For container databases, this metric provides data for root container tablespaces. Statistic: Mean Interval: 30 minutes	30 minutes	tablespaceName tablespaceType
AvgGCCRBlockReceiveTime	**Average GC CR Block Receive Time**	milliseconds	The average global cache CR (consistent-read) block receive time. Statistic: Mean Interval: 5 minutes *For RAC / cluster databases only.*	5 minutes	instanceNumber instanceName hostName

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
`BlockingSessions`	**Blocking Sessions**	count	Current blocking sessions. Statistic: Max Interval: 15 minutes *Not applicable for container databases.*	15 minutes	N/A
`CPUTime`	**CPU Time**	seconds per second	The average rate of accumulation of CPU time by foreground sessions in the database instance over the time interval. The CPU time component of Average Active Sessions. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
`DbmgmtJobExecutionsCount`	??	??	The number of SQL job executions on a single managed database or a database group, and their status. Status dimensions can be the following values: "Succeeded," "Failed," "InProgress." Statistic: Sum Interval: 1 minute	??	managedDatabaseId managedDatabaseGroupId jobId status

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
`DBTime`	**DB Time**	seconds per second	The average rate of accumulation of database time (CPU + Wait) by foreground sessions in the database instance over the time interval. Also known as Average Active Sessions. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
`FRASpaceLimit`	**Flash Recovery Area Limit**	GB	The flash recovery area space limit. Statistic: Max Interval: 15 minutes *Not applicable for pluggable databases.*	15 minutes	N/A
`FRAUtilization`	**Flash Recovery Area Utilization**	percent	The flash recovery area utilization. Statistic: Mean Interval: 15 minutes *Not applicable for pluggable databases.*	15 minutes	N/A
`GCCRBlocksReceived`	**GC CR Blocks Received**	blocks per second	The global cache CR (consistent-read) blocks received per second. Statistic: Mean Interval: 5 minutes *For RAC / cluster databases only.*	5 minutes	instanceNumber instanceName hostName

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
`GCCurrentBlocksReceived`	**GC Current Blocks Received**	blocks per second	Represents global cache current blocks received per second. Statistic reports the mean value. Statistic: Mean Interval: 5 minutes *For Real Application Cluster (RAC) databases only.*	5 minutes	instanceNumber instanceName hostName
`InterconnectTraffic`	**Average Interconnect Traffic**	MB per second	The average internode data transfer rate. Statistic: Mean Interval: 5 minutes *For RAC / cluster databases only.*	5 minutes	instanceNumber instanceName hostName
`InvalidObjects`	**Invalid Objects**	count	Invalid database objects count. Statistic: Max Interval: 24 hours *Not applicable for container databases.*	24 hours	N/A
`IOPS`	**IOPS**	operations per second	The average number of input-output operations per second. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName ioType (Read, Write)
`IOThroughput`	**IO Throughput**	MB per second	The average throughput in MB per second. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName ioType (Read, Write)

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
`LogicalBlocksRead`	**Logical Reads**	reads per second	The average number of blocks read from SGA/ Memory (buffer cache) per second. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
`MaxTablespaceSize`	**Max Tablespace Size**	GB	The maximum possible tablespace size. For container databases, this metric provides data for root container tablespaces. Statistic: Max Interval: 30 minutes	30 minutes	tablespaceName tablespaceType
`MemoryUsage`	**Memory Usage**	MB	Memory pool total size in MB. Statistic: Mean Interval: 15 minutes	15 minutes	instanceNumber instanceName hostName memoryType (SGA, PGA) memoryPool (AllocatedPGA, Buffercachel, FixedSGA, JavaPool, LargePool, LogBuffer, OtherPools, SharedPool, StreamsPool)
`MonitoringStatus`	**Monitoring Status**	not applicable	The monitoring status of the resource. If a metric collection fails, error information is captured in this metric. Statistic: Mean Interval: 5 minutes	5 minutes	collectionName errorSeverity errorCode

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
NonReclaimableFRA	**Non Reclaimable Fast Recovery Area**	percent	The Non-reclaimable fast recovery area. Statistic: Mean Interval: 15 minutes *Not applicable for pluggable databases.*	15 minutes	N/A
ParsesByType	**Parses By Type**	parses per second	The number of hard or soft parses per second. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName parseType (HardParse, SoftParse)
ProblematicScheduledDBMSJobs	**Problematic Scheduled DBMS Jobs**	count	The problematic scheduled database jobs count. Statistic: Max Interval: 15 minutes *Not applicable for container databases.*	15 minutes	type (Broken, Failed)
Processes	**Process Count**	count	The database processes count. Statistic: Max Interval: 1 minute *Not applicable for pluggable databases.*	5 minutes	instanceNumber instanceName hostName
ProcessLimitUtilization	**Process Limit Utilization**	percent	The process limit utilization. Statistic: Mean Interval: 1 minute *Not applicable for pluggable databases.*	5 minutes	instanceNumber instanceName hostName

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
ReclaimableFRA	**Reclaimable Fast Recovery Area**	percent	The reclaimable fast recovery area. Statistic: Mean Interval: 15 minutes *Not applicable for pluggable databases.*	15 minutes	N/A
ReclaimableFRASpace	**Flash Recovery Area Reclaimable Space**	GB	The flash recovery area reclaimable space. Statistic: Mean Interval: 15 minutes *Not applicable for pluggable databases.*	15 minutes	N/A
RedoSize	**Redo Generated**	MB per second	The average amount of redo generated, in MB per second. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
SessionLimitUtilization	**Session Limit Utilization**	percent	The session limit utilization. Statistic: Mean Interval: 1 minute *Not applicable for pluggable databases.*	5 minutes	instanceNumber instanceName hostName
Sessions	**Sessions**	count	The number of sessions in the database. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
TransactionsByStatus	**Transactions By Status**	transactions per second	The number of committed or rolled back transactions per second. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName transactionStatus (Committed, RolledBack)

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
`UnusableIndexes`	**Unusable Indexes**	count	Unusable indexes count in database schema. Statistic: Max Interval: 24 hours *Not applicable for container databases.*	24 hours	schemaName
`UsableFRA`	**Usable Fast Recovery Area**	percent	The useable fast recovery area. Statistic: Mean Interval: 15 minutes *Not applicable for pluggable databases.*	15 minutes	N/A
`UsedFRASpace`	**Flash Recovery Area Usage**	GB	The flash recovery area space usage. Statistic: Max Interval: 15 minutes *Not applicable for pluggable databases.*	15 minutes	N/A
`WaitTime`	**Wait Time**	seconds per second	The average rate of accumulation of non-idle wait time by foreground sessions in the database instance over the time interval. The wait time component of Average Active Sessions. Statistic: Mean Interval: 5 minutes	5 minutes	instanceNumber instanceName hostName waitClass

Oracle Cloud Infrastructure Operations Insights

Oracle Cloud Infrastructure Operations Insights allows you to use the Capacity Planning and SQL Warehouse functionality to gain insight into Oracle Databases deployed in Oracle Cloud (Bare Metal, Virtual Machine VM, and Exadata Cloud Infrastructure).

Using Operations Insights on Oracle Cloud Databases allows you to:

- Analyze resource usage of databases across cloud databases
- Forecast future demand for database resources such as CPU, memory, and storage based on historical trends
- Improve resource utilization by identifying under and over utilized resources
- Identify Exadata systems projected to reach high utilization
- Identify total lead time to expand capacity using machine learning based forecast based on long term historic data to project future resource growth
- Compare SQL performance across databases and identify common patterns

Related Topics

- Enabling Database Cloud Service Databases

Monitor Metrics to Diagnose and Troubleshoot Problems with Pluggable Databases

Enable Database Management service to view metrics to diagnose and troubleshoot problems with pluggable databases.

- About Database Management
- Using the Console to Enable Database Management for a Container Database (CDB)
  To enable Database Management for a container database (CDB), use this procedure.
- Using the Console to Enable Database Management for a Pluggable Database (PDB)
  To enable Database Management for a pluggable database (PDB), use this procedure.
- Using the Console to Edit Database Management for a Pluggable Database (PDB)
  To edit the Database Management configuration for a pluggable database (PDB), use this procedure.
- Using the Console to Disable Database Management for a Pluggable Database (PDB)
  To disable Database Management for a pluggable database (PDB), use this procedure.
- Using the Console to View Performance Hub for a Container Database (CDB)
  To view Performance Hub for a container database (CDB), use this procedure. You must first enable Database Management to view the performance report.
- Using the Console to View Performance Hub for a Pluggable Database (PDB)
  To view Performance Hub for a pluggable database (PDB), use this procedure. You must first enable Database Management to view the performance report.
- Using the API to Enable, Disable, or Update Database Management Service
- Oracle Cloud Database Metrics
  Use the metrics to diagnose and troubleshoot issues.

About Database Management

As a Database Administrator, you can use the Oracle Cloud Infrastructure Database Management service to monitor and manage Oracle Databases. For more information, see *About Database Management*.

Performance Hub provides a visual representation of diagnostic data that you can leverage to fix performance issues or tune the database to improve performance. For more information about Performance Hub, see *Performance Hub*.

Related Topics

- About Database Management
- Performance Hub

Using the Console to Enable Database Management for a Container Database (CDB)

To enable Database Management for a container database (CDB), use this procedure.

> **Note:**
>
> You can also enable Database Management for a database from the Database Management Administration page. For more information, see *Enable Database Management for Oracle Cloud Databases*.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata on Oracle Public Cloud**.
2. Choose your **Compartment**.

   A list of Exadata VM Clusters is displayed.
3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the database for which you want to enable Database Management.

   Exadata VM Cluster Details page is displayed.

   Under **Resources**, **Databases** is selected by default.
4. In the list of databases, click the database for which you want to enable Database Management.

   Database Details page is displayed.
5. In the **Database Information** section, under the **Associated Services**, check the status of Database Management.

   If the Database Management is displayed as **Not Enabled**, perform the following steps:

Enable Database Management

1. Click **Enable**.

   Enable Database Management window is displayed.
2. In the **Database information** section, provide the following details:
   - **Database type**: Read-only. Type of the database.

- **Exadata VM Cluster**: Read-only. Compartment in which the database is located.
- **Database home**: Read-only. Database home of the database.
- **Database name**: Read-only. Name of the database.
- **Service name**: The unique service name of the database. A default unique name is displayed, which can be changed if required.
- **Protocol**: Select either TCP or TCPS to connect to the Oracle Cloud Database. TCP is selected by default.

  > **Note:**
  >
  > – If Oracle Data Guard is enabled after Database Management was enabled for an Exadata VM Cluster using the TCPS protocol, then TCPS will have to be reconfigured. Enabling Oracle Data Guard is causing TCPS configuration to be overwritten, and it's recommended that TCPS is configured on an Exadata VM Cluster after enabling Oracle Data Guard.
  >
  > – Database Management currently does not support Oracle Data Guard configuration and Database Management features are not available for standby databases.

- **Port**: Specify the port number.
  If TCP is selected in the **Protocol** field, then the port number `1521` is displayed by default. You can change it if required. You can select the port number from a range of 1 to 65535.
- **Database wallet secret**: This field is only displayed if TCPS is selected in the **Protocol** field.

  a. Select the secret that contains the database wallet from the drop-down list. If an existing database wallet secret is not available, then select **Create new secret...** from the drop-down list.
  The Create database wallet secret panel is displayed and you can create a new secret.

  For information on database wallets and creating a secret in the Vault service, see *Oracle Cloud Database-related Prerequisite Tasks*.

  b. If the Database Management (`dpd`) service policy that grants Database Management permission to read the secret that contains the database wallet is not created, then the `System policies are required...` message is displayed. You can click **Add policy** to view and automatically create the service policy.
  For information on Vault service permissions required to use existing secrets or create new secrets, see *Permissions Required to Enable Database Management for Oracle Cloud Databases*.

3. In the **Specify credentials for the connection** section, provide the following details:
   - **Database user name**: Enter the database user name.
   - **Database user password secret**:

     a. Select the secret that contains the database user password from the drop-down list. If the compartment in which the secret resides is different from the compartment displayed, then click **Change compartment** and select another

compartment. If an existing secret with the database user password is not available, then select **Create new secret...** from the drop-down list.
The Create password secret panel is displayed and you can create a new secret.

For information on database monitoring user credentials and saving the database user password as a secret in the Vault service, see *Oracle Cloud Database-related Prerequisite Tasks*.

b. If the Database Management (`dpd`) service policy that grants Database Management permission to read the secret that contains the database wallet is not created, then the `System policies are required...` message is displayed. You can click **Add policy** to view and automatically create the service policy.
For information on Vault service permissions required to use existing secrets or create new secrets, see *Permissions Required to Enable Database Management for Oracle Cloud Databases*.

4. In the **Private endpoint information** section, select the private endpoint that will act as a representation of Database Management in the VCN in which the Oracle Cloud Database can be accessed.

   You can choose the private endpoint from a different compartment as well. You must ensure that the appropriate Database Management private endpoint is available.

   Here are the two types of Database Management private endpoints:

   - Private endpoint for single instance Databases in the bare metal and virtual machine DB systems.
   - Private endpoint for Oracle RAC Databases in the virtual machine DB system.

   If a Database Management private endpoint is not available, then you must create one.

   For information on how to create a private endpoint, see *Create a Database Management Private Endpoint*.

5. In the **Management options** section, choose between the following options:

   - **Full management**: This includes fleet management, advanced Performance Hub, and other SKU features along with basic management capabilities.
   - **Basic management**: This includes basic monitoring metrics and the ASH Analytics and SQL Monitoring features in Performance Hub for container databases.
     For more information on management options, see *About Management Options*.

6. Click **Enable Database Management**.

7. A confirmation message with a link to the Oracle Cloud Database's **Work requests** section on the **Database information** page is displayed. Click the link to monitor the progress of the work request.

8. In the **Database Information** section, under the **Associated Services**, verify if the status of **Database Management** is **Enabled**.

   The **Disable** option is also displayed, which you can click to disable Database Management.

If you encounter issues when enabling Database Management, see *Issues Encountered When Enabling Database Management for Oracle Cloud Databases* for likely causes and solutions.

Related Topics

- Permissions Required to Enable Database Management for Oracle Cloud Databases
- Oracle Cloud Database-related Prerequisite Tasks
- Enable Database Management for Oracle Cloud Databases

- Issues Encountered When Enabling Database Management for Oracle Cloud Databases

Using the Console to Enable Database Management for a Pluggable Database (PDB)

To enable Database Management for a pluggable database (PDB), use this procedure.

> **Note:**
>
> You can also enable Database Management for a database from the Database Management Administration page. For more information, see *Enable Database Management for Oracle Cloud Databases*.

Prerequisite

To enable the Database Management for a pluggable database, enable Database Management for the associated database with the **Full Management** option.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata on Oracle Public Cloud**.
2. Choose your **Compartment**.

   A list of Exadata VM Clusters is displayed.
3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the pluggable database for which you want to enable Database Management.

   Exadata VM Cluster Details page is displayed.

   Under **Resources**, **Databases** is selected by default.
4. In the list of databases, click the database that contains the pluggable database for which you want to enable Database Management.

   Database Details page is displayed.
5. Under **Resources**, click **Pluggable Databases**.
6. In the list of pluggable databases, click the pluggable database for which you want to enable Database Management.

   Pluggable Database Details page is displayed.
7. In the **Database Information** section, under the **Associated Services**, check the status of Database Management.

   If the Database Management is displayed as **Not Enabled**, perform the following steps:

Enable Database Management

1. Click **Enable**.

   Enable Database Management window is displayed.
2. In the **Database information** section, provide the following details:
   - **Database type**: Read-only. Type of the database.
   - **Exadata VM Cluster**: Read-only. Compartment in which the database is located.
   - **Database home**: Read-only. Database home of the database.
   - **Pluggable Database name**: Read-only. Name of the database.

- **Service name**: The unique service name of the database. A default unique name is displayed, which can be changed if required.
- **Protocol**: Select either TCP or TCPS to connect to the Oracle Cloud Database. TCP is selected by default.

> **Note:**
>
> – If Oracle Data Guard is enabled after Database Management was enabled for an Exadata VM Cluster using the TCPS protocol, then TCPS will have to be reconfigured. Enabling Oracle Data Guard is causing TCPS configuration to be overwritten, and it's recommended that TCPS is configured on an Exadata VM Cluster after enabling Oracle Data Guard.
>
> – Database Management currently does not support Oracle Data Guard configuration and Database Management features are not available for standby databases.

- **Port**: Specify the port number.
  If TCP is selected in the **Protocol** field, then the port number `1521` is displayed by default. You can change it if required. You can select the port number from a range of 1 to 65535.
- **Database wallet secret**: This field is only displayed if TCPS is selected in the **Protocol** field.
  a. Select the secret that contains the database wallet from the drop-down list. If an existing database wallet secret is not available, then select **Create new secret...** from the drop-down list.
  The Create database wallet secret panel is displayed and you can create a new secret.

  For information on database wallets and creating a secret in the Vault service, see *Oracle Cloud Database-related Prerequisite Tasks*.

  b. If the Database Management (`dpd`) service policy that grants Database Management permission to read the secret that contains the database wallet is not created, then the `System policies are required...` message is displayed. You can click **Add policy** to view and automatically create the service policy.
  For information on Vault service permissions required to use existing secrets or create new secrets, see *Permissions Required to Enable Database Management for Oracle Cloud Databases*.

3. In the **Specify credentials for the connection** section, provide the following details:
   - **Database user name**: Enter the database user name.
   - **Database user password secret**:
     a. Select the secret that contains the database user password from the drop-down list. If the compartment in which the secret resides is different from the compartment displayed, then click **Change compartment** and select another compartment. If an existing secret with the database user password is not available, then select **Create new secret...** from the drop-down list.
     The Create password secret panel is displayed and you can create a new secret.

For information on database monitoring user credentials and saving the database user password as a secret in the Vault service, see *Oracle Cloud Database-related Prerequisite Tasks*.

b. If the Database Management (`dpd`) service policy that grants Database Management permission to read the secret that contains the database wallet is not created, then the `System policies are required...` message is displayed. You can click **Add policy** to view and automatically create the service policy. For information on Vault service permissions required to use existing secrets or create new secrets, see *Permissions Required to Enable Database Management for Oracle Cloud Databases*.

4. In the **Private endpoint information** section, select the private endpoint that will act as a representation of Database Management in the VCN in which the Oracle Cloud Database can be accessed.

   You can choose the private endpoint from a different compartment as well. You must ensure that the appropriate Database Management private endpoint is available.

   Here are the two types of Database Management private endpoints:

   - Private endpoint for single instance Databases in the bare metal and virtual machine DB systems.
   - Private endpoint for Oracle RAC Databases in the virtual machine DB system.

   If a Database Management private endpoint is not available, then you must create one.

   For information on how to create a private endpoint, see *Create a Database Management Private Endpoint*.

5. In the **Management options** section, choose between the following options:

   - **Full management**: This includes fleet management, advanced Performance Hub, and other SKU features along with basic management capabilities.
   - **Basic management**: This includes basic monitoring metrics and the ASH Analytics and SQL Monitoring features in Performance Hub for container databases. For more information on management options, see *About Management Options*.

6. Click **Enable Database Management**.

7. A confirmation message with a link to the Oracle Cloud Database's **Work requests** section on the **Database information** page is displayed. Click the link to monitor the progress of the work request.

8. In the **Database Information** section, under the **Associated Services**, verify if the status of **Database Management** is **Enabled**.

   The **Disable** option is also displayed, which you can click to disable Database Management.

If you encounter issues when enabling Database Management, see *Issues Encountered When Enabling Database Management for Oracle Cloud Databases* for likely causes and solutions.

Related Topics

- Permissions Required to Enable Database Management for Oracle Cloud Databases
- Oracle Cloud Database-related Prerequisite Tasks
- Enable Database Management for Oracle Cloud Databases
- Issues Encountered When Enabling Database Management for Oracle Cloud Databases

Using the Console to Edit Database Management for a Pluggable Database (PDB)

To edit the Database Management configuration for a pluggable database (PDB), use this procedure.

> **Note:**
>
> You can also enable Database Management for a database from the Database Management Administration page. For more information, see *Enable Database Management for Oracle Cloud Databases*.

Prerequisite

To enable the Database Management for a pluggable database, enable Database Management for the associated database with the **Full Management** option.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata on Oracle Public Cloud**.
2. Choose your **Compartment**.

   A list of Exadata VM Clusters is displayed.
3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the pluggable database for which you want to edit Database Management.

   Exadata VM Cluster Details page is displayed.

   Under **Resources**, **Databases** is selected by default.
4. In the list of databases, click the database that contains the pluggable database for which you want to edit Database Management.

   Database Details page is displayed.
5. Under **Resources**, click **Pluggable Databases**.
6. In the list of pluggable databases, click the pluggable database for which you want to edit Database Management.

   Pluggable Database Details page is displayed.
7. In the **Database Information** section, under the **Associated Services**, check the status of Database Management.

   If the Database Management is displayed as **Enabled**, perform the following steps to edit Database Management:

Edit Database Management

1. Click **Enable**.

   Edit Database Management window is displayed.
2. In the **Database information** section, provide the following details:
   - **Database type**: Read-only. Type of the database.
   - **Exadata VM Cluster**: Read-only. Compartment in which the database is located.
   - **Database home**: Read-only. Database home of the database.

- **Pluggable Database name**: Read-only. Name of the database.
- **Service name**: The unique service name of the database. A default unique name is displayed, which can be changed if required.
- **Protocol**: Select either TCP or TCPS to connect to the Oracle Cloud Database. TCP is selected by default.

> **Note:**
>
> – If Oracle Data Guard is enabled after Database Management was enabled for an Exadata VM Cluster using the TCPS protocol, then TCPS will have to be reconfigured. Enabling Oracle Data Guard is causing TCPS configuration to be overwritten, and it's recommended that TCPS is configured on an Exadata VM Cluster after enabling Oracle Data Guard.
>
> – Database Management currently does not support Oracle Data Guard configuration and Database Management features are not available for standby databases.

- **Port**: Specify the port number.
  If TCP is selected in the **Protocol** field, then the port number `1521` is displayed by default. You can change it if required. You can select the port number from a range of 1 to 65535.
- **Database wallet secret**: This field is only displayed if TCPS is selected in the **Protocol** field.
  a. Select the secret that contains the database wallet from the drop-down list. If an existing database wallet secret is not available, then select **Create new secret...** from the drop-down list.
  The Create database wallet secret panel is displayed and you can create a new secret.

  For information on database wallets and creating a secret in the Vault service, see *Oracle Cloud Database-related Prerequisite Tasks*.

  b. If the Database Management (`dpd`) service policy that grants Database Management permission to read the secret that contains the database wallet is not created, then the `System policies are required...` message is displayed. You can click **Add policy** to view and automatically create the service policy.
  For information on Vault service permissions required to use existing secrets or create new secrets, see *Permissions Required to Enable Database Management for Oracle Cloud Databases*.

3. In the **Specify credentials for the connection** section, provide the following details:
   - **Database user name**: Enter the database user name.
   - **Database user password secret**:
     a. Select the secret that contains the database user password from the drop-down list. If the compartment in which the secret resides is different from the compartment displayed, then click **Change compartment** and select another compartment. If an existing secret with the database user password is not available, then select **Create new secret...** from the drop-down list.
     The Create password secret panel is displayed and you can create a new secret.

For information on database monitoring user credentials and saving the database user password as a secret in the Vault service, see *Oracle Cloud Database-related Prerequisite Tasks*.

b. If the Database Management (`dpd`) service policy that grants Database Management permission to read the secret that contains the database wallet is not created, then the `System policies are required...` message is displayed. You can click **Add policy** to view and automatically create the service policy. For information on Vault service permissions required to use existing secrets or create new secrets, see *Permissions Required to Enable Database Management for Oracle Cloud Databases*.

4. In the **Private endpoint information** section, select the private endpoint that will act as a representation of Database Management in the VCN in which the Oracle Cloud Database can be accessed.

   You can choose the private endpoint from a different compartment as well. You must ensure that the appropriate Database Management private endpoint is available.

   Here are the two types of Database Management private endpoints:

   - Private endpoint for single instance Databases in the bare metal and virtual machine DB systems.
   - Private endpoint for Oracle RAC Databases in the virtual machine DB system.

   If a Database Management private endpoint is not available, then you must create one.

   For information on how to create a private endpoint, see *Create a Database Management Private Endpoint*.

5. In the **Management options** section, choose between the following options:

   - **Full management**: This includes fleet management, advanced Performance Hub, and other SKU features along with basic management capabilities.
   - **Basic management**: This includes basic monitoring metrics and the ASH Analytics and SQL Monitoring features in Performance Hub for container databases. For more information on management options, see *About Management Options*.

6. Click **Enable Database Management**.

7. A confirmation message with a link to the Oracle Cloud Database's **Work requests** section on the **Database information** page is displayed. Click the link to monitor the progress of the work request.

8. In the **Database Information** section, under the **Associated Services**, verify if the status of **Database Management** is **Enabled**.

   The **Disable** option is also displayed, which you can click to disable Database Management.

If you encounter issues when enabling Database Management, see *Issues Encountered When Enabling Database Management for Oracle Cloud Databases* for likely causes and solutions.

Related Topics

- Permissions Required to Enable Database Management for Oracle Cloud Databases
- Oracle Cloud Database-related Prerequisite Tasks
- Enable Database Management for Oracle Cloud Databases
- Issues Encountered When Enabling Database Management for Oracle Cloud Databases

Using the Console to Disable Database Management for a Pluggable Database (PDB)

To disable Database Management for a pluggable database (PDB), use this procedure.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata on Oracle Public Cloud**.
2. Choose your **Compartment**.

   A list of Exadata VM Clusters is displayed.
3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the pluggable database for which you want to disable Database Management.

   Exadata VM Cluster Details page is displayed.

   Under **Resources**, **Databases** is selected by default.
4. In the list of databases, click the database that contains the pluggable database for which you want to disable Database Management.

   Database Details page is displayed.
5. Under **Resources**, click **Pluggable Databases**.
6. In the list of pluggable databases, click the pluggable database for which you want to disable Database Management.

   Pluggable Database Details page is displayed.
7. In the **Database Information** section, under the **Associated Services**, check the status of Database Management.
8. If the Database Management is displayed as **Enabled**, perform the following steps to disable Database Management:
   a. Click **Disable**.
   b. A confirmation message with a link to the **Work requests** section on the **Database information** page is displayed. Click the link to monitor the progress of the work request.
   c. In the **Database Information** section, under the **Associated Services**, verify if the status of Database Management is **Disabled**.

Using the Console to View Performance Hub for a Container Database (CDB)

To view Performance Hub for a container database (CDB), use this procedure. You must first enable Database Management to view the performance report.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata on Oracle Public Cloud**.
2. Choose your **Compartment**.

   A list of Exadata VM Clusters is displayed.
3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the database for which you want to view Performance Hub.

   Exadata VM Cluster Details page is displayed.

   Under **Resources**, **Databases** is selected by default.

4. In the list of databases, click the database for which you want to view Performance Hub.

   Database Details page is displayed.

5. Click **Performance Hub**.

With Basic Management, Performance Hub provides **ASH Analytics** and **SQL Monitoring**. Advanced Management will additionally provide **ADDM**, **Workload**, and **Blocking Sessions**.

Performance Hub allows you to download reports for your managed databases. For more information about downloading reports, see *Automatic Workload Repository (AWR) Report*, *Active Sessions History (ASH) Report*, and *Performance Hub Report*.

Related Topics

- Automatic Workload Repository (AWR) Report
- Active Sessions History (ASH) Report
- Performance Hub Report

Using the Console to View Performance Hub for a Pluggable Database (PDB)

To view Performance Hub for a pluggable database (PDB), use this procedure. You must first enable Database Management to view the performance report.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata on Oracle Public Cloud**.
2. Choose your **Compartment**.

   A list of Exadata VM Clusters is displayed.

3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the pluggable database for which you want to view Performance Hub.

   Exadata VM Cluster Details page is displayed.

   Under **Resources**, **Databases** is selected by default.

4. In the list of databases, click the database that contains the pluggable database.

   Database Details page is displayed.

5. Under **Resources**, click **Pluggable Databases**.
6. In the list of pluggable databases, click the pluggable database that you're interested in.

   Pluggable Database Details page is displayed.

7. Click **Performance Hub**.

With Basic Management, Performance Hub provides **ASH Analytics** and **SQL Monitoring**. Advanced Management will additionally provide **ADDM**, **Workload**, and **Blocking Sessions**.

Performance Hub allows you to download reports for your managed databases. For more information about downloading reports, see *Automatic Workload Repository (AWR) Report*, *Active Sessions History (ASH) Report*, and *Performance Hub Report*.

Related Topics

- Automatic Workload Repository (AWR) Report
- Active Sessions History (ASH) Report
- Performance Hub Report

Using the API to Enable, Disable, or Update Database Management Service

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.

Use these API operations to configure the Database Management service.

- Enable Database Management service for an Oracle Database located in Oracle Cloud Infrastructure to access tools including Metrics and Performance hub: `enableDatabaseManagement`
- Disable Database Management service: `disableDatabaseManagement`
- Update Database Management configuration: `updateDatabaseManagement`

Oracle Cloud Database Metrics

Use the metrics to diagnose and troubleshoot issues.

The metrics for Oracle Cloud Databases help measure useful quantitative data, such as CPU and storage utilization, the number of successful and failed database logon and connection attempts, database operations, SQL queries, transactions, and so on.

For more information, see *Oracle Cloud Database Metrics*.

- Using the Console View Metrics for a Container Database (CDB)
  To view metrics for a container database (CDB), you must first enable Database Management with the **Full Management** option.
- Using the Console to View Metrics for a Pluggable Database (PDB)
  To view metrics for a Pluggable Database (PDB), the following prerequisites must be met:

Related Topics

- Oracle Cloud Database Metrics

Using the Console View Metrics for a Container Database (CDB)

To view metrics for a container database (CDB), you must first enable Database Management with the **Full Management** option.

To enable Database Management for databases, see *Using the Console to Enable Database Management for a Database*.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata on Oracle Public Cloud**.
2. Choose your **Compartment**.

   A list of Exadata VM Clusters is displayed.
3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the database for which you want to view the metrics.

   Exadata VM Cluster Details page is displayed.

   Under **Resources**, **Databases** is selected by default.
4. In the list of databases, click the database for which you want to view the metrics.

   Database Details page is displayed.

5. Under **Resources**, click **Metrics**.

Related Topics

- Using the Console to Enable Database Management for a Container Database (CDB)
  To enable Database Management for a container database (CDB), use this procedure.

Using the Console to View Metrics for a Pluggable Database (PDB)

To view metrics for a Pluggable Database (PDB), the following prerequisites must be met:

- Enable Database Management for databases with the **Full Management** option.
- Enable Database Management for pluggable databases.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata on Oracle Public Cloud**.
2. Choose your **Compartment**.

   A list of Exadata VM Clusters is displayed.
3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the pluggable database for which you want to view the metrics.

   Exadata VM Cluster Details page is displayed.

   Under **Resources**, **Databases** is selected by default.
4. In the list of databases, click the database that contains the pluggable database.

   Database Details page is displayed.
5. Under **Resources**, click **Pluggable Databases**.
6. In the list of pluggable databases, click the pluggable database for which you want to view the metrics.

   Pluggable Database Details page is displayed.
7. Under **Resources**, click **Metrics**.
8. Select a namespace from the **Metric namespace** from where you wish to view metrics.

> **Note:**
>
> - When Database Management is enabled, then you can view metrics only from the `oracle_oci_database` namespace.
> - When Database Management is disabled, then a banner, "`Database management must be enabled to provide data for metrics.`" is displayed.

With Basic Management, Performance Hub provides **ASH Analytics** and **SQL Monitoring**. Advanced Management will additionally provide **ADDM**, **Workload**, and **Blocking Sessions**.

Related Topics

- Using the Console to Enable Database Management for a Container Database (CDB)
  To enable Database Management for a container database (CDB), use this procedure.
- Using the Console to Enable Database Management for a Pluggable Database (PDB)
  To enable Database Management for a pluggable database (PDB), use this procedure.

Security Guide for Oracle Exadata Database Service on Dedicated Infrastructure

This guide describes security for an Exadata Cloud Infrastructure. It includes information about the best practices for securing the Exadata Cloud Infrastructure.

Part 1: Security Configurations and Default Enabled Features

Responsibilities

Exadata Cloud Infrastructure is jointly managed by the customer and Oracle.

The Exadata Cloud Infrastructure deployment is divided into two areas of responsibility:

Customer accessible services: components that the customer can access as part of their subscription to Exadata Cloud Infrastructure

- Customer accessible virtual machines (VM)
- Customer accessible database services

Oracle Managed Infrastructure: components that are owned and operated by Oracle to run customer accessible services

- Power Distribution Units (PDUs)
- Out of band (OOB) management switches » InfiniBand switches
- Exadata Storage Servers
- Physical Exadata database servers
- Datacenter security which hosts Exadata Servers with customer information

Customers control and monitor access to customer services, including network access to their VMs (through OCI Virtual Cloud Networks and OCI Security Lists), authentication to access the VM, and authentication to access databases running in the VMs. Oracle controls and monitors access to Oracle Managed Infrastructure components and physical server security. Oracle staff are not authorized to access customer services, including customer VMs and databases except where customers are unable to access the customer VM. See the Exadata Cloud Service Security Controls document, https://www.oracle.com/a/ocom/docs/engineered-systems/exadata/exadata-cloud-service-security.pdf, Exception Workflows .

Customers access Oracle databases (DB) running on Exadata Cloud Infrastructure via client and backup VCNs to the databases running in the customer VM using standard Oracle database connection methods, such as Oracle Net on port 1521. Customer's access the VM running the Oracle databases via standard Oracle Linux methods, such as token based ssh on port 22.

Infrastructure Security

Secutrity features offered by Exadata Cloud Infrastructure.

- **Oracle Cloud Physical Security**
  Oracle Cloud data centers align with Uptime Institute and Telecommunications Industry Association (TIA) ANSI/TIA-942-A Tier 3 (99.982% Availability) or Tier 4 (99.995% Availability) standards and follow a N2 ('N' stands for Need) redundancy methodology for critical equipment operation. Data centers housing Oracle Cloud Infrastructure services use redundant power sources and maintain generator backups in case of widespread electrical outage. Server rooms are closely monitored for air temperature and humidity, and fire-suppression systems are in place. Data center staff are trained in incident response and escalation procedures to address security and availability events that may arise. For more information see *Oracle Cloud Infrastructure Security Guide*. For further details on Oracle Cloud Infrastructure Data Center compliance, and see *Oracle Cloud Compliance*.
- **Operator access to customer systems**
  Oracle access protocols include:
  - Physical access to facilities is limited to certain Oracle employees, contractors, and authorized visitors.
  - Oracle employees, subcontractors, and authorized visitors are issued identification cards that must be worn while on Oracle premises.
  - Visitors are required to sign a visitor's register, be escorted and/or observed when they are on Oracle premises, and/or be bound by the terms of a confidentiality agreement with Oracle.
  - Security monitors the possession of keys/access cards and the ability to access facilities. Staff leaving Oracle's employment must return keys/cards and key/cards are deactivated upon termination.

  - Security authorizes all repairs and modifications to the physical security barriers or entry controls at service locations.
  - Oracle use a mixture of 24/7 onsite security officers or patrol officers, depending on the risk/protection level of the facility. In all cases officers are responsible for patrols, alarm response, and recording of security incidents.
  - Oracle has implemented centrally managed electronic access control systems with integrated intruder alarm capability. The access logs are kept for a minimum of six months. Furthermore, the retention period for CCTV monitoring and recording ranges from 30-90 days minimum, depending on the facility's functions and risk level.

- **Hypervisor Customer Isolation**
  The hypervisor is the software that manages virtual devices in a cloud environment, handling server and network virtualization. In traditional virtualization environments, the hypervisor manages network traffic, enabling it to flow between VM instances and between VM instances and physical hosts. This adds considerable complexity and computational overhead in the hypervisor. Proof-of concept computer security attacks, such as virtual machine escape attacks, have highlighted the substantial risk that can come with this design. These attacks exploit hypervisor complexity by enabling an attacker to "breakout" of a VMinstance, access the underlying operating system, and gain control of the hypervisor. The attacker can then access other hosts, sometimes undetected. Oracle Cloud Infrastructure reduces this risk by decoupling network virtualization from the hypervisor. We've implemented network virtualization as a highly customized hardware and software layer that moves cloud control away from the hypervisor and host, and puts it on its own network. This hardened and monitored layer of control is what enables isolated network virtualization. Isolated network virtualization reduces risk by limiting the attack surface. Even if a malicious actor succeeds with a VM escape attack on a single host, it's designed so they can't reach other hosts in the cloud infrastructure. The attack is effectively contained to the one host. Isolated network virtualization is implemented in every data center in every region, which means that all Oracle Cloud Infrastructure tenants benefit from this reduced risk.

Figure 6-1 Isolated Network Virualization Reduces Risk in Oracle Generation 2 Cloud

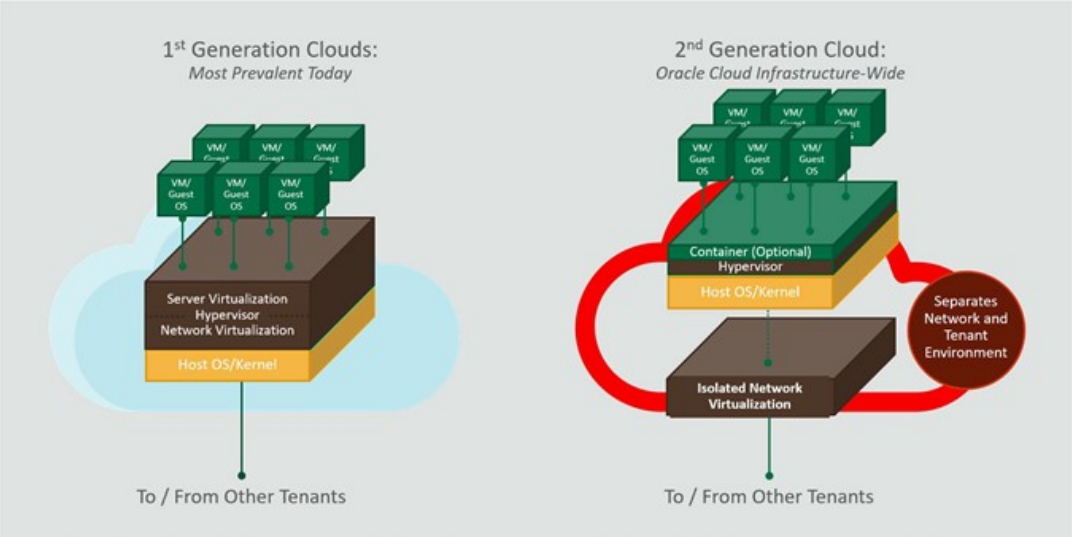

- **Multitenant Security**
  Consistent with our security philosophy of Defense in Depth, Multitenant has a comprehensive isolation architecture.

There are four major categories to this, with several important features in each category.

1. Access Control Mechanism
2. Prevent Unauthorized Admin Access
3. Protect from direct access to Data Files
4. Resource Isolation

Figure 6-2 Multitenant's Comprehensive Isolation Architecture

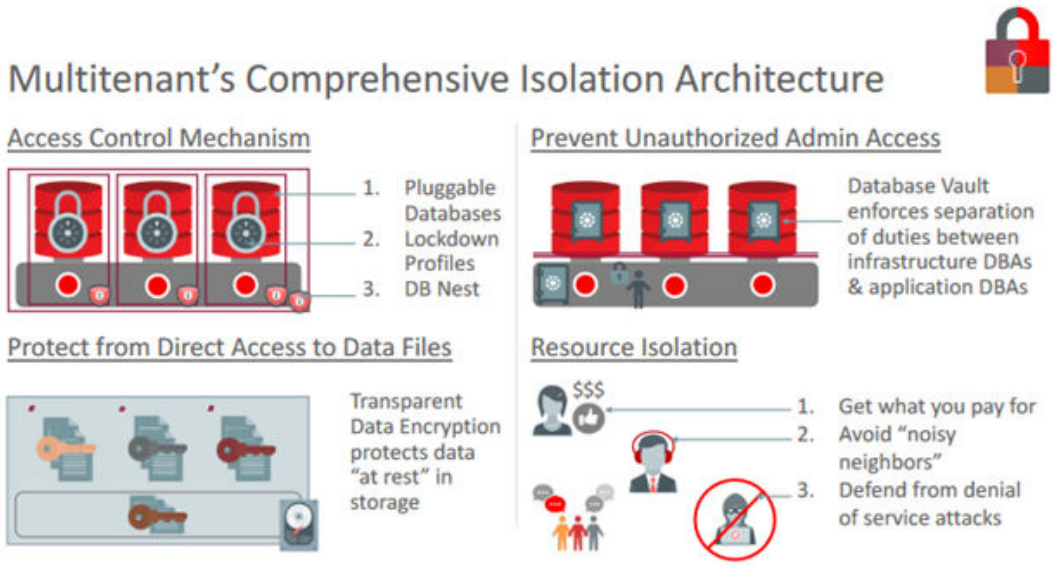

Related Topics

- Oracle Cloud Infrastructure Security Architecture
- Oracle Cloud Infrastructure Security Guide
- Data Security: Physical and Environmental Controls
- Oracle Multitenant with Oracle Database 19c
- Oracle Cloud Compliance

Guiding Principles Followed for Security Configuration Defaults

- **Defense in Depth** Exadata Cloud Infrastructure offers a number of controls to ensure confidentiality, integrity, and availability throughout the service.
  First, Exadata Cloud Infrastructure is built from the hardened operating system image provided by Exadata Database Machine (https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmsq/exadata-security-overview.html). This secures the core operating environment by restricting the installation image to only the required software packages, disabling unnecessary services, and implementing secure configuration parameters throughout the system.

  In addition to inheriting all the strength of Exadata Database Machine's mature platform, because Exadata Cloud Infrastructure provisions systems for customers, additional secure default configuration choices are implemented in the service instances. For example, all database tablespaces require transparent data encryption (TDE), strong password enforcement for initial database users and superusers, and enhanced audit and event rules.

Exadata Cloud Infrastructure also constitutes a complete deployment and service, so it is subjected to industry-standard external audits such as PCI, HIPPA and ISO27001. These external audit requirements impose additional value-added service features such as antivirus scanning, automated alerting for unexpected changes to the system, and daily vulnerability scans for all Oracle-managed infrastructure systems in the fleet.

- **Least privilege**
  Oracle Secure Coding Standards require software processes run at the minimum privilege level to implement their functionality.

  Each process and daemon, must run as a normal, unprivileged user unless it can prove a requirement for a higher level of privilege. This helps contain any unforeseen issues or vulnerabilities to unprivileged user space and not compromise an entire system.

  This principle also applies to Oracle operations team members who use individual named accounts to access the Exadata Cloud Infrastructure for maintenance or troubleshooting. Only when necessary will they use the audited access to higher levels of privilege to solve or resolve an issue. Most issues are resolved through automation, so we also employ least privilege by not permitting human operators to access a system unless the automation is unable to resolve the issue.
- **Auditing and accountability**
  When required, access to the system is allowed, but all access and actions are logged and tracked for accountability.

  Exadata Cloud Infrastructure audit logs are controlled by Oracle and used for security monitoring and compliance purposes. Oracle can share relevant audit logs with customers per Oracle Incident Response Practices and the Oracle Data Processing Agreement.

  Auditing capabilities are provided at all infrastructure components to ensure all actions are captured. Customers also have ability to configure auditing for their database and guest VM configuration and may choose to integrate those with other enterprise auditing systems.
- **Automating cloud operations**
  By eliminating manual operations required to provision, patch, maintain, troubleshoot, and configure systems, the opportunity for error is reduced.

Security Features

- **Hardened OS image**
  - Minimal package installation:

    Only the necessary packages required to run an efficient system are installed. By installing a smaller set of packages, the attack surface of the operating system is reduced and the system remains more secure.
  - Secure configuration:

    Many non-default configuration parameters are set during installation to enhance the security posture of the system and its content. For example, SSH is configured to only listen on certain network interfaces, sendmail is configured to only accept localhost connections, and many other similar restrictions are implemented during installation.
  - Run only necessary services:

    Any services that may be installed on the system, but not required for normal operation, are disabled by default. For example, while NFS is a service often configured by customers for various application purposes, it is disabled by default as it is not required for normal database operations. Customers may choose to optionally configure services per their requirements.

- **Minimized attack surface**

  As part of the hardened image, attack surface is reduced installing and running only the software required to deliver the service.

- **Additional security features enabled (grub passwords, secure boot)**
  - Leveraging Exadata image capabilities, ExaDB-D enjoys the features integrated into the base image such as grub passwords and secure boot in addition to many others.
  - Through customization, customers may wish to further enhance their security posture with additional configurations.
- **Secure access methods**
  - Accessing database servers via SSH using strong cryptographic ciphers. Weak ciphers are disabled by default.
  - Accessing databases via encrypted Oracle Net connections. By default, our services are available using encrypted channels and a default configured Oracle Net client will use encrypted sessions.
  - Accessing diagnostics via Exadata MS web interface (https).
- **Auditing and logging**
  - By default, auditing is enabled for administrative operations and those audit records are communicated to OCI internal systems for automated review and alerting when required.

Guest VM Default Fixed Users

Several user accounts regularly manage the components of Exadata Cloud Infrastructure. These users are required and may not be modified.

In all ExaDB-D machines, Oracle uses and recommends token-based SSH login.

There are three classes of users:

- Default Users: No Logon Privileges
- Default Users WITH RESTRICTED SHELL Access
  These users are used for accomplishing a defined task with a restricted shell login. These users should not be removed as the defined task will fail in case these users are deleted.
- Default Users with Login Permissions
  These privileged users are used for accomplishing most of the tasks in the system. These users should never be altered or deleted as it would have significant impact on the running system.

Default Users: No Logon Privileges

This user list consists of default operating system users along with some specialized users like exawatch and dbmsvc. These users should not be altered. These users cannot login to the system as all are set to /sbin/nologin.

In the list of users below, most are either standard Linux OS users or related to standard Linux packages except for the exawatch and dbmsvc users.

- exawatch: The exawatch user is responsible for collecting and archiving system statistics on both the database servers and the storage servers

- dbmsvc: User is used for Management Server on the database node service in Oracle Exadata System

NOLOGIN Users

```
bin:x:1:1:bin:/bin:/sbin/nologin
Daemon:x:2:2:daemon:/sbin:/sbin/nologin
adm:x:3:4:adm:/dev/null:/sbin/nologin
mail:x:8:12:mail:/var/spool/mail:/sbin/nologin
nobody:x:99:99:Nobody:/:/sbin/nologin
systemd-network:x:192:192:systemd Network Management:/:/sbin/nologin
dbus:x:81:81:System message bus:/:/sbin/nologin
rpm:x:37:37::/var/lib/rpm:/sbin/nologin
sshd:x:74:74:Privilege-separated SSH:/var/empty/sshd:/sbin/nologin
rpc:x:32:32:Rpcbind Daemon:/var/lib/rpcbind:/sbin/nologin
unbound:x:999:997:Unbound DNS resolver:/etc/unbound:/sbin/nologin
nscd:x:28:28:NSCD Daemon:/:/sbin/nologin
tss:x:59:59:Account used by the trousers package to sandbox the tcsd
daemon:/dev/null:/sbin/nologin
saslauth:x:998:76:Saslauthd user:/run/saslauthd:/sbin/nologin
mailnull:x:47:47::/var/spool/mqueue:/sbin/nologin
smmsp:x:51:51::/var/spool/mqueue:/sbin/nologin
chrony:x:997:996::/var/lib/chrony:/sbin/nologin
rpcuser:x:29:29:RPC Service User:/var/lib/nfs:/sbin/nologin
nslcd:x:65:55:LDAP Client User:/:/sbin/nologin
uucp:x:10:14:Uucp user:/var/spool/uucp:/sbin/nologin
tcpdump:x:72:72::/:/sbin/nologin
exawatch:x:1010:510::/opt/oracle.ExaWatcher:/sbin/nologin
sssd:x:996:508:User forsssd:/:/sbin/nologin
dbmsvc:x:2001:2001::/:/sbin/nologin
clamupdate:x:995:504:Clamav database update user:/var/lib/clamav:/sbin/nologin
```

Default Users WITH RESTRICTED SHELL Access

These users are used for accomplishing a defined task with a restricted shell login. These users should not be removed as the defined task will fail in case these users are deleted.

`dbmmonitor` password is set to a random string during deployment, which must change on first use.

- `dbmmonitor`: The `dbmmonitor` user is used for DBMCLI Utility. For more details refer to Using the DBMCLI Utility

Restricted Shell Users

```
dbmmonitor:x:2003:2003::/home/dbmmonitor:/bin/rbash
```

Default Users with Login Permissions

These privileged users are used for accomplishing most of the tasks in the system. These users should never be altered or deleted as it would have significant impact on the running system.

SSH keys are used for login by customer staff and cloud automation software.

Customer-added SSH keys may be added by the `UpdateVmCluster` method, or by customers directly accessing the customer VM and managing SSH keys inside of the customer VM.

Customers are responsible for adding comments to keys to make them identifiable. When a customer adds the SSH key by the `UpdateVmCluster` method, the key is only added to the `authorized_keys` file of the `opc` user.

Cloud automation keys are temporary, specific to a given cloud automation task, for example, VM Cluster Memory resize, and unique. Cloud automation access keys can be identified by the following comments: `OEDA_PUB`, `EXACLOUD_KEY`, `ControlPlane`. Cloud automation keys are removed after the cloud automation task completes so the `authorized_keys` files of the `root`, `opc`, `oracle`, and `grid` accounts should only contain cloud automation keys while the cloud automation actions are running.

Privileged Users

```
root:x:0:0:root:/root:/bin/bash
oracle:x:1001:1001::/home/oracle:/bin/bash
grid:x:1000:1001::/home/grid:/bin/bash
opc:x:2000:2000::/home/opc:/bin/bash
dbmadmin:x:2002:2002::/home/dbmadmin:/bin/bash
```

- `root`: Linux requirement, used sparingly to run local privileged commands. `root` is also used for some processes like Oracle Trace File Analyzer Agent and `ExaWatcher`.
- `grid`: Owns Oracle Grid Infrastructure software installation and runs Grid Infastructure processes.
- `oracle`: Owns Oracle database software installation and runs Oracle Database processes.
- `opc`:
  - Used by Oracle Cloud automation for automation tasks.
  - Has the ability to run certain privileged commands without further authentication (to support automation functions).
  - Runs the local agent, also known as “DCS Agent” that performs lifecycle operations for Oracle Database and Oracle Grid Infastructure software (patching, create database, and so on).
- `dbmadmin`:
  - The `dbmadmin` user is used for Oracle Exadata Database Machine Command-Line Interface (DBMCLI) utility.
  - The `dbmadmin` user should be used to run all services on the database server. For more information, see Using the DBMCLI Utility.

Related Topics

- Using the DBMCLI Utility

Default Security Settings: Customer VM

In addition to all of the Exadata features explained in Security Features of Oracle Exadata Database Machine, the following security settings are also applicable, to Exadata Cloud Infrastructureinstances.

- Custom database deployment with non-default parameters.
  The command `host_access_control` is to configure Exadata security settings:
  - Implementing password aging and complexity policies.
  - Defining account lockout and session timeout policies.

  - Restricting remote root access.
  - Restricting network access to certain accounts.
  - Implementing login warning banner.
- `account-disable`: Disables a user account when certain configured conditions are met.
- `pam-auth`: Various PAM settings for password changes and password authentication.
- `rootssh`: Adjusts the `PermitRootLogin` value in `/etc/ssh/sshd_config`, which permits or denies the `root` user to login through SSH.
  - By default, `PermitRootLogin` is set to `no`.
  - PermitRootLogin=without-password is required for the cloud automation to perform some lifecycle management operations, disabling root login will cause that service functionality to fail.
- `session-limit`: Sets the `* hard maxlogins` parameter in `/etc/security/limits.conf`, which is the maximum number of logins for all users. This limit does not apply to a user with `uid=0`.
  Defaults to `* hard maxlogins 10` and it is the recommended secure value.
- `ssh-macs`: Specifies the available Message Authentication Code (MAC) algorithms.
  - The MAC algorithm is used in protocol version 2 for data integrity protection.
  - Defaults to `hmac-sha1`, `hmac-sha2-256`, `hmac-sha2-512` for both server and client.
  - Secure recommended values: `hmac-sha2-256`, `hmac-sha2-512` for both server and client.
- `password-aging`: Sets or displays the current password aging for interactive user accounts.
  - `-M`: Maximum number of days a password may be used.
  - `-m`: Minimum number of days allowed between password changes.
  - `-W`: Number of days warning given before a password expires.
  - Defaults to `-M 99999`, `-m 0`, `-W 7`
  - `--strict_compliance_only-M 60`, `-m 1`, `-W 7`
  - Secure recommended values: `-M 60`, `-m 1`, `-W 7`

Related Topics

- Security Features of Oracle Exadata Database Machine

Default Processes on Customer VM

A list of the processes that run by default on the customer VM, also called DOMU, or Guest VM and Guest OS

- **Exadata Cloud Infrastructure VM agent:**
  Cloud agent for handling database lifecycle operations.
  - Runs as `opc` user
  - Process table shows it running as a Java process with `jar` names - `dbcs-agent-VersionNumber-SNAPSHOT.jar` and `dbcs-admin-VersionNumver-SNAPSHOT.jar`.
- **Oracle Trace File Analyzer agent:**

Oracle Trace File Analyzer (TFA) provides a number of diagnostic tools in a single bundle, making it easy to gather diagnostic information about the Oracle database and clusterware, which in turn helps with problem resolution when dealing with Oracle Support

  - Runs as `root` user
  - Runs as initd demon (`/etc/init.d/init.tfa`)
  - Process tables show a Java application (`oracle.rat.tfa.TFAMain`)
  - Runs as `root` and `exawatch` users.
  - Runs as backgroud script, `ExaWatcher.sh` and all its child process run as a Perl process.
  - Process table shows as multiple Perl applications.`ExaWatcher`:

- **Database and GI (clusterware):**
  - Runs as `dbmsvc` and `grid` users
  - Process table shows following applications:
    * `oraagent.bin`, `apx_*` and `ams_*` as `grid` user
    * `dbrsMain`, and Java applications - `derbyclient.jar`, `weblogic.Server` as `oracle` user.
- **Management Server (MS):**
  Part of Exadata image software for managing and monitoring the image functions.
  - Runs as `dbmadmin`.
  - Process table shows it running as a Java process.
- Guest VM Network Security
- Compliance Requirements

Guest VM Network Security

Table 6-29 Default Port Matrix for Guest VM Services

Type of interface	Name of interface	Port	Process running
Bridge on client VLAN	`bondeth0`	22	sshd
		1521	Oracle TNS listener
		Optionally, customers can assign a SCAN listener port (TCP/IP) in the range between 1024 and 8999. Default is 1521.	
		5000	Oracle Trace File Analyzer Collector
		7879	Jetty Management Server

Table 6-29 (Cont.) Default Port Matrix for Guest VM Services

Type of interface	Name of interface	Port	Process running
	bondeth0:1	1521 Optionally, customers can assign a SCAN listener port (TCP/IP) in the range between 1024 and 8999. Default is 1521.	Oracle TNS listener
	bondeth0:2	1521 Optionally, customers can assign a SCAN listener port (TCP/IP) in the range between 1024 and 8999. Default is 1521.	Oracle TNS listener
Bridge on backup VLAN	bondeth1	7879	Jetty Management Server
Oracle Clusterware running on each cluster node communicates through these interfaces.	clib0/clre0	1525	Oracle TNS listener
		3260	Synology DSM iSCSI
		5054	Oracle Grid Interprocess Communication
		7879	Jetty Management Server
		Dynamic Port: 9000-65500 Ports are controlled by the configured ephemeral range in the operating system and are dynamic.	System Monitor service (osysmond)
		Dynamic Port: 9000-65500 Ports are controlled by the configured ephemeral range in the operating system and are dynamic.	Cluster Logger service (ologgerd)
	clib1/clre1	5054	Oracle Grid Interprocess communication
		7879	Jetty Management Server
Cluster nodes use these interfaces to access storage cells (ASM disks). However, the IP/ports 7060/7070 attached to the storage interfaces are used to access DBCS agent from the Control Plane server.	stib0/stre0	7060	dbcs-admin
		7070	dbcs-agent
	stib1/stre1	7060	dbcs-admin
		7070	dbcs-agent

Table 6-29 (Cont.) Default Port Matrix for Guest VM Services

Type of interface	Name of interface	Port	Process running
Control Plane server to domU	`eth0`	22	sshd
Loopback	`lo`	22	sshd
		2016	Oracle Grid Infrastructure
		6100	Oracle Notification Service (ONS), part of Oracle Grid Infrastructure
		7879	Jetty Management Server
		Dynamic Port 9000-65500	Oracle Trace File Analyzer

> **Note:**
>
> TNS listener opens dynamic ports after initial contact to well known ports (1521, 1525).

Default iptables rules for Guest VM:

The default iptables are setup to ACCEPT connections on input, forward, and output chains.

```
#iptables -L -n -v
Chain INPUT (policy ACCEPT 0 packets, 0 bytes)
 pkts bytes target     prot opt in     out     source
destination

Chain FORWARD (policy ACCEPT 0 packets, 0 bytes)
 pkts bytes target     prot opt in     out     source
destination

Chain OUTPUT (policy ACCEPT 0 packets, 0 bytes)
 pkts bytes target     prot opt in     out     source
destination
```

Compliance Requirements

PII (Personally Identifiable Information) This information is considered confidential and sensitive, and must be protected to prevent unauthorized use of personal information for the purposes of legal regulation, financial liability, and personal reputation.

You must configure a set of explicit rules to prevent Personally Identifiable Information (PII) from being displayed in your data.

The default Application Performance Monitoring rules hide PII in URLs by recognizing monetary values, bank-account numbers, and dates. However, the default rules only catch

obvious PII and are not exhaustive. You must evaluate the default rules and further configure rules to ensure correct reporting in your environment and ensure that PII is not displayed in your data.

For more information, see Hide Personally Identifiable Information and Security and Personally Identifiable Information

Backup Retention

When you enable the Automatic Backup feature, the service creates daily incremental backups of the database to Object Storage. The first backup created is a level 0 backup. Then, level 1 backups are created every day until the next weekend. Every weekend, the cycle repeats, starting with a new level 0 backup.

If you choose to enable automatic backups, you can choose one of the following preset retention periods: 7 days, 15 days, 30 days, 45 days, or 60 days. The system automatically deletes your incremental backups at the end of your chosen retention period.

For more information, see Manage Database Backup and Recovery on Oracle Exadata Database Service on Dedicated Infrastructure

Audit Log Retention Period

The OCI Audit service provides records of API operations performed against supported services as a list of log events. By default, Audit service records are retained for 365 days.

For more information, see Audit Log Retention Period

Service Log Retention

Oracle Cloud Infrastructure services, such as API Gateway, Events, Functions, Load Balancing, Object Storage, and VCN Flow Logs emit service logs. Each of these supported services has a Logs resource that allows you to enable or disable logging for that service. By default, Log retention is 1 month, but it can be set until 6 months.

Logs groups can be used to limit access to sensitive logs generated by services using IAM policy. You don't have to rely on complex compartment hierarchies to secure your logs. For example, say the default log group in a single compartment is where you store logs for the entire tenancy. You grant access to the compartment for log administrators with IAM policy as you normally would. However, let's say some projects contain personally identifiable information (PII) and those logs can only be viewed by a select group of log administrators. Log groups allow you to put logs that contain PII into a separate log group, and then use IAM policy to restrict access to all but a few log administrators.

For more information, see Service Logs and Managing Logs and Log Groups

Default Database Security Configuration

Default database security features enabled and used:

- Transparent Database Encryption (TDE) is used for database tablespaces created by Oracle Database Cloud tools.
  - CDB$ROOT: users tablespace is encrypted
  - PDBs: all tablespaces encrypted
  - Wallet password is provided during initial DB creation. Wallet passwords may be changed using `dbaascli`. Customers should change this password periodically.
- Users in the database

  - No additional users are created in the database.
  - After DB creation, all DB users are locked except for SYS, SYSTEM and DBSNMP.
  - Auditing is enabled for the following operations:
    * `DATABASE LINK`
    * `PUBLIC DATABASE LINK`
    * `PUBLIC SYNONYM`
    * `DROP ANY PROCEDURE`
    * `PROCEDURE`
    * `ALTER SYSTEM`
    * `TRIGGER`
    * `CREATE DATABASE LINK`
    * `ALTER DATABASE LINK`
    * `CREATE PROCEDURE`
    * `ALTER SYSTEM`
    * `CREATE TRIGGER`
    * `CREATE ANY TRIGGER`
    * `SELECT ANY DICTIONARY`
    * **DB VERSION_11_2:** `EXEMPT REDACTION POLICY`
    * **DB VERSION_12_1 or DB VERSION_12_2:** `BECOME USER`
    * **DB VERSION_12_1:** `SESSION`
    * `DBAASSECURE` profile is created and it is set as default profile for database user account.

- Native SQL*Net encryption for all network connections - Relevant `sqlnet.ora` parameters set in Exadata Cloud Infrastructure by default are:
  - `SQLNET.ENCRYPTION_TYPES_SERVER = (AES256, AES192, AES128)`
  - `SQLNET.ENCRYPTION_SERVER = requested`
  - `SQLNET.CRYPTO_CHECKSUM_SERVER = accepted`
  - `SQLNET.CRYPTO_CHECKSUM_TYPES_SERVER = (SHA256, SHA384, SHA512)`
- TCPS protocol offered for network connection to the database on port 2484 (wallet configured at `/var/opt/oracle/dbaas_acfs/grid/tcps_wallets`). Relevant `sqlnet.ora` parameters set in Exadata Cloud Infrastructure by default are:
  - ```
    SSL_CIPHER_SUITES = (SSL_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,
    SSL_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,
    SSL_ECDHE_RSA_WITH_AES_128_GCM_SHA256,
    SSL_ECDHE_RSA_WITH_AES_256_GCM_SHA384)
    ```
    ```

  - `WALLET_LOCATION = (SOURCE=(METHOD=FILE)(METHOD_DATA=(DIRECTORY=/var/opt/oracle/dbaas_acfs/grid/tcps_wallets)))`
  - `SQLNET.IGNORE_ANO_ENCRYPTION_FOR_TCPS = TRUE`
  - `SSL_VERSION = 1.2`

- Remote listener registration - listeners run from GI home. Exadata Cloud Infrastructure deployments the Grid Infrastructure vresion speified in *Oracle Support Document 2333222.1 (Exadata Cloud Service Software Versions)*. Exadata Cloud Infrastructure default configuration includes `listener.ora` parameter `VALID_NODE_CHECKING_REGISTRATION_LISTENER=SUBNET` combined with `REMOTE_REGISTRATION_ADDRESS_<SCANLISTENER>=<value>` to restrict remote listener registrations for security purposes.
- OCI Vault integration - TDE encryption key may be stored in OCI Vault (a Key Management System). For more information and instructions to configure principals, vaults, etc. see *Customer-Managed Keys in Exadata Cloud Infrastructure*. Both private vs shared vault types are supported for Exadata Cloud Infrastructure OCI Vault integration. DB user authentication is not integrated with OCI Vault.

Related Topics

- Securing Database
- Customer-Managed Keys in Exadata Cloud Infrastructure
  Customer-managed keys for Exadata Cloud Infrastructure is a feature of Oracle Cloud Infrastructure (OCI) Vault service that enables you to encrypt your data using encryption keys that you control.
- Oracle Support Document 2333222.1 (Exadata Cloud Service Software Versions)

Default Backup Security Configuration

OS/VM backups:

Oracle does a full backup of guest VM weekly and maintains one or more backup copies. These backups are full disk snapshots of the guest VM (local OS filesystems, not ASM disk groups which reside on Exadata storage). This backup is triggered at a preset time every week. The backups are stored locally in the dom0. Customers can request Oracle to restore the guest VM image from the most recent backup by filing a My Oracle Support (MOS) Service Request (SR). Oracle cannot restore specific files from the image backup. Customers should perform file level backups in the guest VM if they require the ability to perform single-file restore.

Managed DB backups:

- Weekly full backup (level 0)
- Daily rolling incremental backup (level 1) on seven day cycle
- Automatic backups daily at a specific time set during the database deployment creation process

Retention period for backups vary from 30 days (on Object Storage) to 7 days (on local storage)

Encryption:

- Both Object Storage and local storage: All backups to cloud storage are encrypted.
- Object Storage only: All backups to cloud storage are encrypted.

All backups can be configured via CP UI or CP API.

All backups are encrypted with the same master key used for Transparent Data Encryption (TDE) wallet encryption.

Operator Access to Customer System and Customer Data

Only automated tooling is permitted to access guest VM for purposes of lifecycle automation.

One specific use case is when guest VM is unable to boot. In this case, customers must provide permission to access the guest VM for recovery purposes. Details to handle this scenario are described in section "Exception Workflows" of Exadata Cloud Service Security Controls.

Customers control and monitor access to customer services, including network access to their guest VMs (through layer 2 VLANs and firewalls implemented in the guest VM), authentication to access the guest VM, and authentication to access databases running in the guest VMs. Oracle controls and monitors access to Oracle-managed infrastructure components. Oracle staff are not authorized to access customer services, including guest VMs and databases.

Compliance Requirements

PII (Personally Identifiable Information) This information is considered confidential and sensitive, and must be protected to prevent unauthorized use of personal information for the purposes of legal regulation, financial liability, and personal reputation.

You must configure a set of explicit rules to prevent Personally Identifiable Information (PII) from being displayed in your data.

The default Application Performance Monitoring rules hide PII in URLs by recognizing monetary values, bank-account numbers, and dates. However, the default rules only catch obvious PII and are not exhaustive. You must evaluate the default rules and further configure rules to ensure correct reporting in your environment and ensure that PII is not displayed in your data.

For more information, see Hide Personally Identifiable Information and Security and Personally Identifiable Information

Backup Retention

When you enable the Automatic Backup feature, the service creates daily incremental backups of the database to Object Storage. The first backup created is a level 0 backup. Then, level 1 backups are created every day until the next weekend. Every weekend, the cycle repeats, starting with a new level 0 backup.

If you choose to enable automatic backups, you can choose one of the following preset retention periods: 7 days, 15 days, 30 days, 45 days, or 60 days. The system automatically deletes your incremental backups at the end of your chosen retention period.

For more information, see Manage Database Backup and Recovery on Oracle Exadata Database Service on Dedicated Infrastructure

Audit Log Retention Period

The OCI Audit service provides records of API operations performed against supported services as a list of log events. By default, Audit service records are retained for 365 days.

For more information, see Audit Log Retention Period

Service Log Retention

Oracle Cloud Infrastructure services, such as API Gateway, Events, Functions, Load Balancing, Object Storage, and VCN Flow Logs emit service logs. Each of these supported services has a Logs resource that allows you to enable or disable logging for that service. By default, Log retention is 1 month, but it can be set until 6 months.

Logs groups can be used to limit access to sensitive logs generated by services using IAM policy. You don't have to rely on complex compartment hierarchies to secure your logs. For example, say the default log group in a single compartment is where you store logs for the entire tenancy. You grant access to the compartment for log administrators with IAM policy as you normally would. However, let's say some projects contain personally identifiable information (PII) and those logs can only be viewed by a select group of log administrators. Log groups allow you to put logs that contain PII into a separate log group, and then use IAM policy to restrict access to all but a few log administrators.

For more information, see Service Logs and Managing Logs and Log Groups

Break Glass Procedure for Accessing Customer's Guest VM

There are situations where some problems can only be resolved by Oracle logging into the customer guest VM.

Below are situations where customer's guest VM access is require and recommended procedures for accessing guest VM:

1. Situations where the **starter database is not yet created and customer do not have ssh access to their guest VM yet.** An example would be SR opened by customer to troubleshoot why customer is unable to create a starter database. In this situation, customer never had access to guest VM and no database have yet been created and hence no customer data exists in guest VM.

   As per the security policy associated with ExaDB-D service, Oracle personnel are prohibited to access customer guest VM without customer’s explicit permission. To comply with this policy, Oracle requires to get Customer permission to access guest VM by asking the following question.

   “In order for Oracle to resolve the issue described *in this SR, we need customer’s explicit permission allowing us to login to customer guest VM. By giving us explicit permission to access guest VM, you are confirming that there is no confidential data that is stored in customer guest VM or associated databases and customer security team is authorizing Oracle to have access to customer guest VM in order for Oracle to help fix this issue. Do I have your explicit permission to access guest VM?"*

   After affirmative response by customer, Oracle support staff can login to customer guest VM to resolve the issue.

2. Situations where a **number of databases exist in customer system and customer have access to guest VM but now support needs to login to guest VM to resolve one of the many situations**
   We have encountered (Nodes doesn’t start because of changes on guest VM, eg. Non-existing mounts in fstab, need to run fsck, Hugepage / sysctl conf modification or lvm backup not completed successfully, fstab has wrong entries for non-existing mounts, customer changed the sshd configurations or permissions in /etc/ssh/sshd_config file, etc.) or simply because customer wants Oracle to help resolve the issue they are facing.

   This case is more serious than the first one as there could be some sensitive data in customer guest VM file system or database. In this case, our support staff will be required

to ask the customer to open a new explicit SR specifically to get this permission with the following SR title and content.

As per the security policy associated with ExaDB-D service, Oracle personnel are prohibited to access customer guest VM without customer's explicit permission. For Oracle to comply with this policy, We are required to ask you to open a new SR with exact language as shown below granting Oracle an explicit permission to access guest VM.Please note any modification to the language below may delay resolution of your SR.

New SR Title: SR granting Oracle explicit permission to access DomU of ExaDB-C@C with AK serial number AK99999999

New SR Content: We are opening this SR to grant explicit permission to Oracle to access our DomU in order for support to help resolve issue described in SR# 1-xxxxxxxx.

We acknowledge that by providing this permission, we understand that Oracle will have access to ALL FILES in DomU and agree that there are no confidential

files stored in any of the file systems in DomU. In addition, we also agree that customer security team has authorized Oracle to have access to customer DomU

in order to resolve the issue described in the above SR.

After affirmative response by customer in the above SR, Oracle support staff can login to customer guest VM to resolve the issue.

Part 2: Additional Procedures for Updating Security Posture

Customer Responsibilities

A list of Oracle Cloud Operations responsibilities and customer responsibilities for various operations by components

Table 6-30 Oracle Cloud Ops and Customer Responsibilities for various operations

Operations	Oracle Cloud Ops responsibilities for ORACLE CLOUD PLAFTORM	Customer responsibilities for ORACLE CLOUD PLAFTORM	Oracle Cloud Ops responsibilities for CUSTOMER / TENANT INSTANCES	Customer responsibilities for CUSTOMER / TENANT INSTANCES
DATABASE DEPLOYMENT	Software instrastructure and guidance for ExaCS deployment	**Network Admin:** Configure cloud network infraestructure (VCN, Backup/ Client subnet, Gateway, etc)**Database Admin:** Setup database requirements (Memory, Storage, Computation, Database version, Database type, etc)	Install Operating System, Database and Grid Infraestructure	**Database Admin:** Mantain customer hardware requirements based on workloads
MONITORING	Physical Security, Infraestructure, Control Plane, Hardware Faults, Availability, Capacity	Nothing required	Infrastructure availability to support customer monitoring of customer services	**Database Admin:** Monitoring of Customer Operating System, Databases, Apps and Grid Infraestructure
INCIDENT MANAGEMENT & RESOLUTION	Incident Managment and RemediationSpare parts and field dispatch	Nothing required	Support for any incidents related to the underlying platform	**Database Admin:** Incident Management and resolution for Customer's apps
PATCH MANAGEMENT	Proactive patching of hardware, IaaS/ PaaS control stack	Nothing required	Staging of available patches, for example, Oracle Database patch set	**Database Admin:** Patching of tenant instancesTesting
BACKUP & RESTORATION	Infrastructure and Control Plane backup and recovery, recreate customer VMs	Nothing required	Provide running and customer accessible VM	**Database Admin:** Snapshots / backup and recovery of customer's IaaS and PaaS data using Oracle native or third-party capabiltiy

Enabling additional security capabilities

- KMS Integration (HSM keys)
- Using non-default encryption algorithms for TDE tablespace encryption

KMS Integration (HSM keys)

Oracle Exadata Cloud Service (ExaCS) has integration with the OCI Vault service to protect data at rest for its databases. Users now have the control to create and manage TDE master keys within the OCI Vault that protect your Exadata databases.

With this feature, users have the option to start using the OCI vault service to store and manage the master encryption keys. The OCI Vault keys used for protecting databases are stored in a highly available,

durable, and managed service. OCI vault integration for ExaCS is only available after Oracle Database 11g release 2 (11.2.0.4).

With OCI Vault integration with ExaDB-D, customers can now:

- Centrally control and manage your TDE master keys
- Have their TDE master keys stored in a highly available, durable and managed service wherein the keys are protected by hardware security modules (HSM) that meet Federal Information Processing Standards (FIPS) 140-2 Security Level 3 security certification.
- Rotate their encryption keys periodically to maintain security compliance and regulatory needs.
- Migrate from Oracle-managed keys to customer-managed keys for their existing databases.
- The Key version will only be assigned to the container database (CDB), and not to its pluggable database (PDB). PDB will be assigned an automatically generated new key version.

Related Topics

- Announcing Customer-Managed Encryption Keys for Oracle Exadata Cloud Service
- Manage Databases on Exadata Cloud Infrastructure

Using non-default encryption algorithms for TDE tablespace encryption

In the published Oracle Advanced Security Guide (section Encrypting Columns in Tables the methodology to create a table to encrypt columns using a non-default encryption algorithm ids described.

Troubleshooting Exadata Cloud Infrastructure Systems

These topics cover some common issues you might run into and how to address them.

- Known Issues for Exadata Cloud Infrastructure
  General known issues.
- Troubleshoot Network Connectivity
  To determine if a VM Cluster is properly configured to access the Oracle Cloud Infrastructure (OCI) Services Network, you need to perform the following steps on each virtual machine in the VM Cluster.
- Backup Failures in Exadata Database Service on Dedicated Infrastructure
  If your Exadata managed backup does not successfully complete, you can use the procedures in this topic to troubleshoot and fix the issue.
- Troubleshooting Oracle Data Guard
  Learn to identify and resolve Oracle Data Guard issues.
- Patching Failures on Exadata Cloud Infrastructure Systems

Known Issues for Exadata Cloud Infrastructure

General known issues.

CPU Offline Scaling Fails

Description: CPU offline scaling fails with the following error:

```
** CPU Scale Update **An error occurred during module execution. Please refer to the log file for more information
```

Cause: After provisioning a VM cluster, the `/var/opt/oracle/cprops/cprops.ini` file, which is automatically generated by the database as a service (DBaaS) is not updated with the `common_dcs_agent_bindHost` and `common_dcs_agent_port` parameters and this causes CPU offline scaling to fail.

Action: As the `root` user, manually add the following entries in the `/var/opt/oracle/cprops/cprops.ini` file.

```
common_dcs_agent_bindHost=<IP_Address>
common_dcs_agent_port=7070
```

> **Note:**
>
> The `common_dcs_agent_port` value is **7070** always.

Run the following command to get the IP address:

```
netstat -tunlp | grep 7070
```

For example:

```
netstat -tunlp | grep 7070
tcp 0 0 <IP address 1>:7070 0.0.0.0:* LISTEN 42092/java
tcp 0 0 <IP address 2>:7070 0.0.0.0:* LISTEN 42092/java
```

You can specify either of the two IP addresses, *<IP address 1>* or *<IP address 2>* for the `common_dcs_agent_bindHost` parameter.

Adding a VM to a VM Cluster Fails

Description: When adding a VM to a VM cluster, you might encounter the following issue:

```
[FATAL] [INS-32156] Installer has detected that there are non-readable files in oracle
home.
CAUSE: Following files are non-readable, due to insufficient permission oracle.ahf/data/
scaqak03dv0104/diag/tfa/tfactl/user_root/tfa_client.trc
ACTION: Ensure the above files are readable by grid.
```

Cause: Installer has detected a non-readable trace file, `oracle.ahf/data/scaqak03dv0104/diag/tfa/tfactl/user_root/tfa_client.trc` created by Autonomous Health Framework (AHF) in Oracle home that causes adding a cluster VM to fail.

AHF ran as `root` created a `trc` file with `root` ownership, which the `grid` user is not able to read.

Action: Ensure that the AHF trace files are readable by the `grid` user before you add VMs to a VM cluster. To fix the permission issue, run the following commands as `root` on all the existing VM cluster VMs:

```
chown grid:oinstall /u01/app/19.0.0.0/grid/srvm/admin/logging.properties


chown -R grid:oinstall /u01/app/19.0.0.0/grid/oracle.ahf*


chown -R grid:oinstall /u01/app/grid/oracle.ahf*
```

Troubleshoot Network Connectivity

To determine if a VM Cluster is properly configured to access the Oracle Cloud Infrastructure (OCI) Services Network, you need to perform the following steps on each virtual machine in the VM Cluster.

Validation check for Identity and Access management connectivity:

- `ssh` to a virtual machine on your ExaDB-D VM Cluster as opc user.
- Execute the command: `curl https://identity.<region>.oci.oraclecloud.com` here <region> corresponds to the OCI region where your VM Cluster is deployed. If your VM Cluster is deployed in the Ashburn region you need to use “us-ashburn-1” for <region>. The curl command will now look like `curl https://identity.us-ashburn-1.oci.oraclecloud.com`.
- If your Virtual Cloud Network (VCN) is properly configured for accessing the OCI Services Network, you will get an immediate response that looks like

```
{
 "code" : "NotAuthorizedOrNotFound",
 "message" : "Authorization failed or requested resource not found."
}
```

- The ssh session will hang and will eventually timeout if your network is **not** configured for accessing the OCI Services
- Depending on your VCN setup, you will need to follow the steps outlined in the action section below to configure access to the OCI Services Network.

Validation check for Object Storage Service (OSS) connectivity:

- ssh to a virtual machine on your ExaDB-D VM Cluster as `opc` user.

- Execute the command: `curl https://objectstorage.<region>.oraclecloud.com`, here <region> corresponds to the OCI region where your VM Cluster is deployed. If your VM Cluster is deployed in the Ashburn region you need to use “us-ashburn-1” for <region>. The curl command will now look like `curl https://objectstorage.us-ashburn-1.oraclecloud.com`.
- If your Virtual Cloud Network (VCN) is properly configured for accessing the OCI Services Network, you will get an immediate response that looks like

```
{
 "code" : "NotAuthorizedOrNotFound",
 "message" : "Authorization failed or requested resource not found."
}
```

- The ssh session will hang and will eventually timeout if your network is **not** configured for accessing the OCI Services
- Depending on your VCN setup, you will need to follow the steps outlined in the action section below to configure access to the OCI Services Network.

Action:

- **This action is applicable to customers who have deployed their VM Cluster on a private subnet.**
  If you haven’t already configured a Service Gateway to reach the OCI Services Network, use the instructions in the documentation to configure a Service Gateway for use by the VM Cluster to reach the OCI Services https://docs.oracle.com/en/engineered-systems/exadata-cloud-service/ecscm/ecs-network-setup.html#GUID-51C3EC2C-20DA-4EE5-B882-CD500FA6F7C6
- **This action is applicable to customers who have deployed their VM Cluster on a public subnet**.
  If you haven’t already configured an Internet Gateway to reach the OCI Services Network, use the instructions in the documentation to configure the Internet Gateway for use by the VM Cluster to reach OCI Services https://docs.oracle.com/en/engineered-systems/exadata-cloud-service/ecscm/ecs-network-setup.html#GUID-D8296957-E344-4688-B626-42A99E1D164B

Once you configure your VCN to reach the OCI Services network following the above instructions, execute the steps in both the **Validation check** sections to ensure that you have established connectivity to the OCI Services network from your VM Cluster.

Additional Information:

You can find instructions to update a service gateway here (https://docs.oracle.com/en-us/iaas/Content/Network/Tasks/servicegateway.htm#switch_label)

Backup Failures in Exadata Database Service on Dedicated Infrastructure

If your Exadata managed backup does not successfully complete, you can use the procedures in this topic to troubleshoot and fix the issue.

The most common causes of backup failure are the following:

- The host cannot access Object Storage
- The database configuration on the host is not correct

The information that follows is organized by the error condition. If you already know the cause, you can skip to the section with the suggested solution. Otherwise, use the procedure in Determining the Problem to get started.

- Determining the Problem
  In the Console, a failed database backup either displays a status of **Failed** or hangs in the **Backup in Progress** or **Creating** state. If the error message does not contain enough information to point you to a solution, you can gather more information by using `dbaascli` and by viewing the log files. Then, refer to the applicable section in this topic for a solution.
- Database Service Agent Issues
  Your Oracle Cloud InfrastructureDatabase makes use of an agent framework to allow you to manage your database through the cloud platform. Use the following to check and restart the agent.
- Object Store Connectivity Issues
  Backing up your database to Oracle Cloud Infrastructure Object Storage requires that the host can connect to the applicable Swift endpoint.
- Host Issues
  One or more of the following conditions on the database host can cause backups to fail:
- Database Issues
  An improper database state or configuration can lead to failed backups.
- TDE Wallet and Backup Failures
  Learn to identify the root cause of TDE wallet and backup failures.

Determining the Problem

In the Console, a failed database backup either displays a status of **Failed** or hangs in the **Backup in Progress** or **Creating** state. If the error message does not contain enough information to point you to a solution, you can gather more information by using `dbaascli` and by viewing the log files. Then, refer to the applicable section in this topic for a solution.

NOT_SUPPORTED

Database backups can fail during the `RMAN` configuration stage or during a running `RMAN` backup job. RMAN configuration tasks include validating object store connectivity, backup module installation, and `RMAN` configuration changes. The log files you examine depend on which stage the failure occurs.

1. Log on to the host as the `root` user.
2. Check the applicable log file:
   - If the failure occurred during `RMAN` configuration, navigate to the `/var/opt/oracle/log/<database_name>/bkup/` directory and check the `bkup.log` file.
   - If the failure occurred during the backup job, navigate to the `/var/opt/oracle/log/<database_name>/obkup/` directory and check the `obkup.log` file.

> **Note:**
>
> - Each execution of `bkup` and `obkup` commands generates a separate log file but `bkup.log` and `obkup.log` are symbolic links that point to the most recently generated log file.
> - Ensure that you check the log files on all of the Exadata DB system compute nodes because all nodes send backup pieces to Object Storage.

Database Service Agent Issues

Your Oracle Cloud InfrastructureDatabase makes use of an agent framework to allow you to manage your database through the cloud platform. Use the following to check and restart the agent.

Occasionally you might need to restart the dcsagent program if it has the status of **stop/waiting** to resolve a backup failure. View the `/opt/oracle/dcs/log/dcs-agent.log` file to identify issues with the agent.

NOT_SUPPORTED

1. From a command prompt, check the status of the agent:

   ```
   systemctl status dbcsagent.service
   ```

2. If the agent is in the **stop/waiting** state, try to restart the agent:

   ```
   systemctl start dbcsagent.service
   ```

3. Check the status of the agent again to confirm that it has the **stop/running** status:

   ```
   systemctl status dbcsagent.service
   ```

Object Store Connectivity Issues

Backing up your database to Oracle Cloud Infrastructure Object Storage requires that the host can connect to the applicable Swift endpoint.

Though Oracle controls the actual Swift user credentials for the storage bucket for managed backups, verifying general connectivity to Object Storage in your region is a good indicator that object store connectivity is not the issue. You can test this connectivity by using another Swift user.

NOT_SUPPORTED

1. Create a Swift user in your tenancy. See Working with Auth Tokens.
2. With the user you created in the previous step, use the following command to verify the host can access the object store.

   ```
   curl -v -X HEAD -u <user_ID>:'<auth_token>' https://
   swiftobjectstorage.<region_name>.oraclecloud.com/v1/
   <object_storage_namespace>
   ```

See Object Storage FAQ for the correct region to use. See Understanding Object Storage Namespaces for information about your Object Storage namespace.

3. If you cannot connect to the object store, refer to Prerequisites for Backups on Exadata Cloud Service topic for information on configuring object store connectivity.

Host Issues

One or more of the following conditions on the database host can cause backups to fail:

NOT_SUPPORTED

If an interactive command such as `oraenv`, or any command that might return an error or warning message, was added to the `.bash_profile` file for the grid or oracle user, Database service operations like automatic backups can be interrupted and fail to complete. Check the `.bash_profile` file for these commands, and remove them.

NOT_SUPPORTED

Backup operations require space in the `/u01` directory on the host file system. Use the `df -h` command on the host to check the space available for backups. If the file system has insufficient space, you can remove old log or trace files to free up space.

NOT_SUPPORTED

Your system might not have the required version of the backup module (`opc_installer.jar`). See Unable to use Managed Backups in your DB System for details about this known issue. To fix the problem, you can follow the procedure in that section or simply update your DB system and database with the latest bundle patch.

NOT_SUPPORTED

Customizing the site profile file (`$ORACLE_HOME/sqlplus/admin/glogin.sql`) can cause managed backups to fail in Oracle Cloud Infrastructure. In particular, interactive commands can lead to backup failures. Oracle recommends that you not modify this file for databases hosted in Oracle Cloud Infrastructure.

Database Issues

An improper database state or configuration can lead to failed backups.

NOT_SUPPORTED

The database must be active and running (ideally on all nodes) while the backup is in progress.

NOT_SUPPORTED

Use the following command to check the state of your database, and ensure that any problems that might have put the database in an improper state are resolved:

```
srvctl status database -d <db_unique_name> -verbose
```

The system returns a message including the database's instance status. The instance status must be `Open` for the backup to succeed. If the database is not running, use the following command to start it:

```
srvctl start database -d <db_unique_name> -o open
```

If the database is mounted but does not have the `Open` status, use the following commands to access the SQL*Plus command prompt and set the status to `Open`:

```
sqlplus / as sysdba
alter database open;
```

NOT_SUPPORTED

When you provision a new database, the archiving mode is set to `ARCHIVELOG` by default. This is the required archiving mode for backup operations. Check the archiving mode setting for the database and change it to `ARCHIVELOG`, if applicable.

NOT_SUPPORTED

Open an SQL*Plus command prompt and enter the following command:

```
select log_mode from v$database;
```

If you need to set the archiving mode to `ARCHIVELOG`, start the database in `MOUNT` status (and not `OPEN` status), and use the following command at the SQL*Plus command prompt:

```
alter database archivelog;
```

Confirm that the `db_recovery_file_dest` parameter points to `+RECO`, and that the `log_archive_dest_1` parameter is set to `USE_DB_RECOVERY_FILE_DEST`.

For RAC databases, one instance must have the `MOUNT` status when enabling archivelog mode. To enable archivelog mode for a RAC database, perform the following steps:

1. Shut down all database instances:

   ```
   srvctl stop database -d
   ```

2. Start one of the database instances in mount state:

   ```
   srvctl start instance -d <db_unique_name> -i <instance_name> -o mount
   ```

3. Access the SQL*Plus command prompt:

   ```
   sqlplus / as sysdba
   ```

4. Enable archive log mode:

   ```
   alter database archivelog;
   exit;
   ```

5. Stop the database:

   ```
   srvctl stop instance -d <db_unique_name> -i <instance_name>
   ```

6. Restart all database instances:

   ```
   srvctl start database -d <db_unqiue_name>
   ```

7. At the SQL*Plus command prompt, confirm the archiving mode is set to: `ARCHIVELOG`:

   ```
   select log_mode from v$database;
   ```

NOT_SUPPORTED

Backups can fail when the database instance has a stuck archiver process. For example, this can happen when the flash recovery area (FRA) is full. You can check for this condition using the `srvctl status database -db <db_unique_name> -v` command. If the command returns the following output, you must resolve the stuck archiver process issue before backups can succeed:

```
Instance <instance_identifier> is running on node *<node_identifier>.
Instance status: Stuck Archiver
```

Refer to ORA-00257:Archiver Error (Doc ID 2014425.1) for information on resolving a stuck archiver process.

After resolving the stuck process, the command should return the following output:

```
Instance <instance_identifier> is running on node *<node_identifier>.
Instance status: Open
```

If the instance status does not change after you resolve the underlying issue with the device or resource being full or unavailable, try restarting the database using the `srvctl` command to update the status of the database in the clusterware.

NOT_SUPPORTED

Editing certain RMAN configuration parameters can lead to backup failures in Oracle Cloud Infrastructure. To check your RMAN configuration, use the `show all` command at the RMAN command line prompt.

See the following list of parameters for details about RMAN the configuration settings that should not be altered for databases in Oracle Cloud Infrastructure.

NOT_SUPPORTED

```
CONFIGURE RETENTION POLICY TO RECOVERY WINDOW OF 30 DAYS;

CONFIGURE CONTROLFILE AUTOBACKUP ON;

CONFIGURE DEVICE TYPE 'SBT_TAPE' PARALLELISM 5 BACKUP TYPE TO COMPRESSED
BACKUPSET;

CONFIGURE CHANNEL DEVICE TYPE DISK MAXPIECESIZE 2 G;
```

```
CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' PARMS  'SBT_LIBRARY=/var/opt/oracle/
dbaas_acfs/<db_name>/opc/libopc.so, ENV=(OPC_PFILE=/var/opt/oracle/dbaas_acfs/
<db_name>/opc/opc<db_name>.ora)';

CONFIGURE ARCHIVELOG DELETION POLICY TO BACKED UP 1 TIMES TO 'SBT_TAPE';

CONFIGURE CHANNEL DEVICE TYPE DISK MAXPIECESIZE 2 G;

CONFIGURE ENCRYPTION FOR DATABASE ON;
```

NOT_SUPPORTED

RMAN backups fail when an object store wallet file is lost. The wallet file is necessary to enable connectivity to the object store.

NOT_SUPPORTED

1. Get the name of the database with the backup failure using SQL*Plus:

   ```
   show parameter db_name
   ```

2. Determine the file path of the backup config parameter file that contains the RMAN wallet information at the Linux command line:

   ```
   locate opc_<database_name>.ora
   ```

   For example:

   ```
   find / -name "opctestdb30.ora" -print /var/opt/oracle/dbaas_acfs/
   testdb30/opc/opctestdb30.ora
   ```

3. Find the file path to the wallet file in the backup config parameter file by inspecting the value stored in the `OPC_WALLET` parameter. To do this, navigate to the directory containing the backup config parameter file and use the following `cat` command:

   ```
   cat opc<database_name>.ora
   ```

   For example:

   ```
   cd /var/opt/oracle/dbaas_acfs/testdb30/opc/


   ls -altr *.ora
   opctestdb30.ora


   cat opctestdb30.ora
   OPC_HOST=https://swiftobjectstorage.us-phoenix-1.oraclecloud.com/v1/
   dbbackupphx
   OPC_WALLET='LOCATION=file:/var/opt/oracle/dbaas_acfs/testdb30/opc/
   opc_wallet CREDENTIAL_ALIAS=alias_opc'
   ```

```
OPC_CONTAINER=bUG3TFsSi8QzjWfuTxqqExample
_OPC_DEFERRED_DELETE=false
```

4. Confirm that the `cwallet.sso` file exists in the directory specified in the `OPC_WALLET` parameter, and confirm that the file has the correct permissions. The file permissions should have the octal value of "600" (`-rw-------`). Use the following command:

   ```
   ls -ltr /var/opt/oracle/dbaas_acfs/<database_name>/opc/opc_wallet
   ```

   For example:

   ```
   ls -altr /var/opt/oracle/dbaas_acfs/testdb30/opc/opc_wallet
   -rw------- 1 oracle oinstall 0 Oct 29 01:59 cwallet.sso.lck
   -rw------- 1 oracle oinstall 111231 Oct 29 01:59 cwallet.sso
   ```

TDE Wallet and Backup Failures

Learn to identify the root cause of TDE wallet and backup failures.

NOT_SUPPORTED

For backup operations to work, the `$ORACLE_HOME/network/admin/sqlnet.ora` file must contain the `ENCRYPTION_WALLET_LOCATION` parameter formatted exactly as follows:

```
ENCRYPTION_WALLET_LOCATION=(SOURCE=(METHOD=FILE)
(METHOD_DATA=(DIRECTORY=/var/opt/oracle/dbaas_acfs/<database_name>/
tde_wallet)))
```

NOT_SUPPORTED

Use the `cat` command to check the TDE wallet location specification. For example:

```
$ cat $ORACLE_HOME/network/admin/sqlnet.ora
ENCRYPTION_WALLET_LOCATION=(SOURCE=(METHOD=FILE)
(METHOD_DATA=(DIRECTORY=/var/opt/oracle/dbaas_acfs/<database_name>/
tde_wallet)))
```

NOT_SUPPORTED

Database backups fail if the TDE wallet is not in the proper state. The following scenarios can cause this problem:

NOT_SUPPORTED

If the database was started using SQL*Plus, and the `ORACLE_UNQNAME` environment variable was not set, the wallet is not opened correctly.

To fix the problem, start the database using the `srvctl` utility:

```
srvctl start database -d <db_unique_name>
```

NOT_SUPPORTED

In a multitenant environment for Oracle Database versions that support PDB-level keystore, each PDB has its own master encryption key. For Oracle 18c databases, this encryption key is stored in a single keystore used by all containers. (Oracle Database 19c does not support a keystore at the PDB level.) After you create or plug in a new PDB, you must create and activate a master encryption key for it. If you do not do so, the `STATUS` column in the `v$encryption_wallet` view shows the value `OPEN_NO_MASTER_KEY`.

To check the master encryption key status and create a master key, do the following:

1. Review the the `STATUS` column in the `v$encryption_wallet` view, as shown in the following example:

```
SQL> alter session set container=pdb2;
Session altered.

SQL> select WRL_TYPE,WRL_PARAMETER,STATUS,WALLET_TYPE from
v$encryption_wallet;

WRL_TYPE   WRL_PARAMETER
STATUS             WALLET_TYPE
---------- ----------------------------------------------
------------------ -----------
FILE       /var/opt/oracle/dbaas_acfs/testdb30/tde_wallet/
OPEN_NO_MASTER_KEY AUTOLOGIN
```

2. Confirm that the PDB is in READ WRITE open mode and is not restricted, as shown in the following example:

```
SQL> show pdbs

CON_ID CON_NAME     OPEN MODE              RESTRICTED
------ ------------ ---------------------- ---------------
2      PDB$SEED     READ ONLY              NO
3      PDB1         READ WRITE             NO
4      PDB2         READ WRITE             NO
```

   The PDB cannot be open in restricted mode (the `RESTRICTED` column must show `NO`). If the PDB is currently in restricted mode, review the information in the `PDB_PLUG_IN_VIOLATIONS` view and resolve the issue before continuing. For more information on the `PDB_PLUG_IN_VIOLATIONS` view and the restricted status, review the Oracle Multitenant Administrator's Guide on pluggable database for your Oracle Database version.

3. Create and activate a master encryption key for the PDB:
   - Set the container to the PDB:

     ```
     ALTER SESSION SET CONTAINER = <pdb>;
     ```

- Create and activate a master encryption key in the PDB by executing the following command:

  ```
  ADMINISTER KEY MANAGEMENT SET KEY USING TAG '<tag>'
  FORCE KEYSTORE IDENTIFIED BY <keystore-password> WITH BACKUP USING
  '<backup_identifier>';
  ```

  Note the following:

  - The `USING TAG` clause is optional and can be used to associate a tag with the new master encryption key.
  - The `WITH BACKUP` clause is optional and can be used to create a backup of the keystore before the new master encryption key is created.

  You can also use the `dbaascli` commands `dbaascli tde status` and `dbaascli tde rotate masterkey` to investigate and manage your keys.

4. Confirm that the status of the wallet has changed from `OPEN_NO_MASTER_KEY` to OPEN by querying the `v$encryption_wallet` view as shown in step 1.

NOT_SUPPORTED

Configuration parameters related to the TDE wallet can cause backups to fail.

NOT_SUPPORTED

Confirm that the wallet status is `open` and the wallet type is `auto login` by checking the `v$encryption_wallet` view. For example:

```
SQL> select status, wrl_parameter,wallet_type from v$encryption_wallet;

STATUS  WRL_PARAMETER                                    WALLET_TYPE
------- ---------------------------------------------- --------------
OPEN    /var/opt/oracle/dbaas_acfs/testdb30/tde_wallet/ AUTOLOGIN
```

For pluggable databases (PDBs), ensure that you switch to the appropriate container before querying `v$encryption_wallet` view. For example:

```
$ sqlplus / as sysdba

SQL> alter session set container=pdb1;

Session altered.

SQL> select WRL_TYPE,WRL_PARAMETER,STATUS,WALLET_TYPE from
v$encryption_wallet;

WRL_TYPE  WRL_PARAMETER                                    STATUS   WALLET_TYPE
--------- ---------------------------------------------- -------- -----------
FILE      /var/opt/oracle/dbaas_acfs/testdb30/tde_wallet/ OPEN     AUTOLOGIN
```

NOT_SUPPORTED

The TDE wallet file (`ewallet.p12`) can cause backups to fail if it is missing, or if it has incompatible file system permissions or ownership. Check the file as shown in the following example as the `root` user:

```
# ls -altr /var/opt/oracle/dbaas_acfs/<database_name>/tde_wallet/ewallet.p12
total 76
-rw------ 1 oracle oinstall 5467 Oct 1 20:17 ewallet.p12
```

The TDE wallet file should have file permissions with the octal value "600" (`-rw-------`), and the owner of this file should be a part of the `oinstall` operating system group.

NOT_SUPPORTED

The auto login wallet file (`cwallet.sso`) can cause backups to fail if it is missing, or if it has incompatible file system permissions or ownership. Check the file as shown in the following example as the `root` user:

```
# ls -altr /var/opt/oracle/dbaas_acfs/<database_name>/tde_wallet/cwallet.sso
total 76
-rw------ 1 oracle oinstall 5512 Oct 1 20:18 cwallet.sso
```

The auto login wallet file should have file permissions with the octal value "600" (`-rw-------`), and the owner of this file should be a part of the `oinstall` operating system group.

Troubleshooting Oracle Data Guard

Learn to identify and resolve Oracle Data Guard issues.

When troubleshooting Oracle Data Guard, you must first determine whether the problem occurs during the Data Guard setup and initialization or during Data Guard operation, when lifecycle commands are entered. The steps to identify and resolve the issues are different, depending on the scenario in which they are used.

There are three lifecycle operations: switchover, failover, and reinstate. The Data Guard broker is used for all of these commands. The broker command line interface (`dgmgrl`) is the main tool used to identify and troubleshoot the issues. Although you can use logfiles to identify root causes, `dgmgrl` is faster and easier to use to check and identify an issue.

Setting up and enabling Data Guard involves multiple steps. Log files are created for each step. If any of the steps fail, review the relevant log file to identify and fix the problem.

- Validation of the primary cloud VM Cluster and database
- Validation of the standby cloud VM Cluster
- Recreating and copying files to the standby database (passwordfile and wallets)
- Creating Data Guard through Network (RMAN Duplicate command)
- Configuring Data Guard broker
- Finalizing the setup

- Troubleshooting Data Guard using logfiles
  The tools used to identify the issue and the locations of relevant logfiles are different, depending on the scenario in which they are used.
- Troubleshooting the Data Guard Setup Process
  The following errors might occur in the different steps of the Data Guard setup process. While some errors are displayed within the Console, most of the root causes can be found in the logfiles

Troubleshooting Data Guard using logfiles

The tools used to identify the issue and the locations of relevant logfiles are different, depending on the scenario in which they are used.

Use the following procedures to collect relevant log files to investigate issues. If you are unable to resolve the problem after investigating the log files, contact My Oracle Support.

> **Note:**
>
> When preparing collected files for Oracle Support, bundle them into a compressed archive, such as a ZIP file.

NOT_SUPPORTED

On each compute node associated with the Data Guard configuration, gather log files pertaining to the problem you experienced.

- Enablement stage log files (such as those documenting the Create Standby Database operation) and the logs for the corresponding primary or standby system.
- Enablement job ID logfiles. For example: 23.
- Locations of enablement log files by enablement stage and Exadata system (primary or standby).
- Database name logfiles (`db_name` or `db_unique_name`, depending on the file path).

> **Note:**
>
> Check all nodes of the corresponding primary and standby Exadata systems. Commands executed on a system may have been run on any of its nodes.

NOT_SUPPORTED

Data Guard Deployer (`DGdeployer`) is the process that performs the configuration. When configuring the primary database, it creates the `/var/opt/oracle/log/<dbname>/dgdeployer/dgdeployer.log` file.

This log should contain the root cause of a failure to configure the primary database.

NOT_SUPPORTED

- The primary log from the `dbaasapi` command-line utility is: `/var/opt/oracle/log/dbaasapi/db/dg/<job_ID>.log`. Look for entries that contain `dg_api`.

- One standby log from the `dbaasapi` command-line utility is: `/var/opt/oracle/log/dbaasapi/db/dg/<job_ID>.log`. In this log, look for entries that contain `dg_api`.
- The other standby log is: `/var/opt/oracle/log/<dbname>/dgcc/dgcc.log`. This log is the Data Guard configuration log.

NOT_SUPPORTED

- The Oracle Cloud Deployment Engine (ODCE) creates the `/var/opt/oracle/log/<dbname>/ocde/ocde.log` file. This log should contain the cause of a failure to create the standby database.
- The `dbaasapi` command line utility creates the `var/opt/oracle/log/dbaasapi/db/dg/<job_ID>.log` file. Look for entries that contain `dg_api`.
- The Data Guard configuration log file is `/var/opt/oracle/log/<dbname>/dgcc/dgcc.log`.

NOT_SUPPORTED

- `DGdeployer` is the process that performs the configuration. It creates the following `/var/opt/oracle/log/<dbname>/dgdeployer/dgdeployer.log` file. This log should contain the root cause of a failure to configure the standby database.
- The `dbaasapi` command-line utility creates the `/var/opt/oracle/log/dbaasapi/db/dg/<job_ID>.log` file. Look for entries that contain `dg_api`.
- The Data Guard configuration log is `/var/opt/oracle/log/<dbname>/dgcc/dgcc.log`.

NOT_SUPPORTED

`DGdeployer` is the process that performs the configuration. While configuring Data Guard, it creates the `/var/opt/oracle/log/<dbname>/dgdeployer/dgdeployer.log` file. This log should contain the root cause of a failure to configure the primary database.

NOT_SUPPORTED

On each node of the primary and standby sites, gather log files for the related database name (`db_name`).

> **Note:**
>
> Check all nodes on both primary and standby Exadata systems. A lifecycle management operation may impact both primary and standby systems.

NOT_SUPPORTED

- **Database alert log:** `/u02/app/oracle/diag/rdbms/<dbname>/<dbinstance>/trace/alert_<dbinstance>.log`
- **Data Guard Broker log:** `/u02/app/oracle/diag/rdbms/<dbname>/<dbinstance>/trace/drc<dbinstance>.log`
- **Cloud tooling log file for Data Guard:** `/var/opt/oracle/log/<dbname>/odg/odg.log`

Troubleshooting the Data Guard Setup Process

The following errors might occur in the different steps of the Data Guard setup process. While some errors are displayed within the Console, most of the root causes can be found in the logfiles

NOT_SUPPORTED

The password entered for enabling Data Guard didn't match the primary admin password for the SYS user. This error occurs during the Validate Primary stage of enablement.

NOT_SUPPORTED

The database may not be running. This error occurs during the Validate Primary stage of enablement. Check with `srvctl` and `sql` on the host to verify that the database is up and running on all nodes.

NOT_SUPPORTED

The primary database could not be configured. Invalid Data Guard commands or failed listener reconfiguration can cause this error.

NOT_SUPPORTED

The TDE wallet could not be created. The Oracle Transparent Database Encryption (TDE) keystore (wallet) files could not be prepared for transportation to the standby site. This error occurs during the create TDE Wallet stage of enablement. Either of the following items can cause failure at this stage:

- The TDE wallet files could not be accessed
- The enablement commands could not create an archive containing the wallet files

Troubleshooting procedure:

1. Ensure that the cluster is accessible. To check the status of a cluster, run the following command:

   ```
   crsctl check cluster -all
   ```

2. If the cluster is down, run the following command to restart it:

   ```
   crsctl start crs -wait
   ```

3. If this error occurs when the cluster is accessible, check the logs for create TDE Wallet (enablement stage) to determine cause and resolution for the error.

NOT_SUPPORTED

The archive containing the TDE wallet was likely not transmitted to the standby site. Retrying usually solves the problem.

NOT_SUPPORTED

- The primary and standby sites may not be able to communicate with each other to configure the standby database. These errors occur during the configure standby database stage of enablement. In this stage, configurations are performed on the standby database, including the rman duplicate of the primary database. To resolve this issue:

1. Verify the connectivity status for the primary and standby sites.
2. Ensure that the host can communicate from port 1521 to all ports. Check the network setup, including Network Security Groups (NSGs), Network Security Lists, and the remote VCN peering setup (if applicable). The best way to test communication between the host and other nodes is to access the databases using SQL*PLUS from the primary to standby and from the standby to the primary.

- The SCAN VIPs or listeners may not be running. Use the test above to help identify the issue.

NOT_SUPPORTED

Possible causes:

- SCAN VIPs or listeners may not be running. You can confirm this issue by using the following commands on any cluster node.

  – `[grid@exa1-****** ~]$ srvctl status scan`

  – `[grid@exa1-****** ~]$ srvctl status scan_listener`

- Databases may not be reachable. You can confirm this issue by attempting to connect using an existing Oracle Net alias.

Troubleshooting procedure:

1. As the oracle OS user, check for the existence of an Oracle Net alias for the container database (CDB). Look for an alias in $ORACLE_HOME/network/admin/<dbname>/tnsnames.ora.
   The following example shows an entry for a container database named db12c:

   ```
   cat $ORACLE_HOME/network/admin/db12c/tnsnames.ora
   DB12C = (DESCRIPTION =(ADDRESS = (PROTOCOL = TCP)(HOST = exa1-*****-
   scan.********.******.******.com)(PORT = 1521))
   (CONNECT_DATA = (SERVER = DEDICATED) (SERVICE_NAME =
   db12c.********.******.******.com)
   (FAILOVER_MODE = (TYPE = select) (METHOD = basic))))
   ```

2. Verify that you can use the alias to connect to the database. For example, as sysdba, enter the following command:

   ```
   sqlplus sys@db12c
   ```

NOT_SUPPORTED

A possible cause for this error is that the Oracle Database sys or system user passwords for the database and the TDE wallet may not be the same. To compare the passwords:

1. Connect to the database as the **sys user** and check the TDE status in

   ```
   V$ENCRYPTION_WALLET
   ```

   .

2. Connect to the database as the **system user** and check the TDE status in

   ```
   V$ENCRYPTION_WALLET
   ```

   .

3. Update the applicable passwords to match. Log on to the system host as **opc** and run the following commands:

   a. To change the SYS password:

      ```
      sudo dbaascli database changepassword --dbname <database_name>
      ```

   b. To change the TDE wallet password:

      ```
      sudo dbaascli tde changepassword --dbname <database_name>
      ```

NOT_SUPPORTED

For possible causes and resolutions to TDE wallet issues, see TDE Wallet and Backup Failures .

NOT_SUPPORTED

When the switchover, failover, and reinstate commands are run, multiple error messages may occur. Refer to the Oracle Database documentation for these error messages.

Note

Oracle recommends using the Data Guard broker command line interface (dgmgrl) to validate the configurations.

1. As the Oracle User, connect to the primary or standby database with `dgmgrl` and verify the configuration and the database:

   ```
   dgmgrl sys/<pwd>@<database>
   DGMGRL> VALIDATE CONFIGURATION VERBOSE
   DGMGRL> VALIDATE DATABASE VERBOSE <PRIMARY>
   DGMGRL> VALIDATE DATABASE VERBOSE <STANDBY>
   ```

2. Consult the Oracle Database documentation to check for the respective error message. For example:

   - **ORA-16766:** Redo apply is stopped.
   - **ORA-16853**: Apply lag has exceeded specified threshold.
   - **ORA-16664**: Unable to receive the result from a member (under the standby database).
   - **ORA-12541**: TNS: no listener (under the primary database)

For cause and resolution, review the errors in Database Error Messages.

Patching Failures on Exadata Cloud Infrastructure Systems

Patching operations can fail for various reasons. Typically, an operation fails because a database node is down, there is insufficient space on the file system, or the virtual machine cannot access the object store.

- Determining the Problem
  In the Console, you can identify a failed patching operation by viewing the patch history of an Exadata Cloud Infrastructure system or an individual database.
- Troubleshooting and Diagnosis
  Diagnose the most common issues that can occur during the patching process of any of the Exadata Cloud Infrastructure components.

Determining the Problem

In the Console, you can identify a failed patching operation by viewing the patch history of an Exadata Cloud Infrastructure system or an individual database.

A patch that was not successfully applied displays a status of `Failed` and includes a brief description of the error that caused the failure. If the error message does not contain enough information to point you to a solution, you can use the database CLI and log files to gather more data. Then, refer to the applicable section in this topic for a solution.

Troubleshooting and Diagnosis

Diagnose the most common issues that can occur during the patching process of any of the Exadata Cloud Infrastructure components.

- Database Server VM Issues
  One or more of the following conditions on the database server VM can cause patching operations to fail.
- Oracle Grid Infrastructure Issues
  One or more of the following conditions on Oracle Grid Infrastructure can cause patching operations to fail.
- Oracle Databases Issues
  An improper database state can lead to patching failures.

Database Server VM Issues

One or more of the following conditions on the database server VM can cause patching operations to fail.

Database Server VM Connectivity Problems

Cloud tooling relies on the proper networking and connectivity configuration between virtual machines of a given VM cluster. If the configuration is not set properly, this may incur in failures on all the operations that require cross-node processing. One example can be not being able to download the required files to apply a given patch.

Given the case, you can perform the following actions:

- Verify that your DNS configuration is correct so that the relevant virtual machine addresses are resolvable within the VM cluster.
- Refer to the relevant Cloud Tooling logs as instructed in the *Obtaining Further Assistance* section and contact Oracle Support for further assistance.

Oracle Grid Infrastructure Issues

One or more of the following conditions on Oracle Grid Infrastructure can cause patching operations to fail.

Oracle Grid Infrastructure is Down

Oracle Clusterware enables servers to communicate with each other so that they can function as a collective unit. The cluster software program must be up and running on the VM Cluster for patching operations to complete. Occasionally you might need to restart the Oracle Clusterware to resolve a patching failure.

In such cases, verify the status of the Oracle Grid Infrastructure as follows:

```
./crsctl check cluster
CRS-4537: Cluster Ready Services is online
CRS-4529: Cluster Synchronization Services is online
CRS-4533: Event Manager is online
```

If Oracle Grid Infrastructure is down, then restart by running the following commands:

```
crsctl start cluster -all
```

```
crsctl check cluster
```

Oracle Databases Issues

An improper database state can lead to patching failures.

Oracle Database is Down

The database must be active and running on all the active nodes so the patching operations can be completed successfully across the cluster.

Use the following command to check the state of your database, and ensure that any problems that might have put the database in an improper state are resolved:

```
srvctl status database -d db_unique_name -verbose
```

The system returns a message including the database instance status. The instance status must be **Open** for the patching operation to succeed.

If the database is not running, use the following command to start it:

```
srvctl start database -d db_unique_name -o open
```

Obtaining Further Assistance

If you were unable to resolve the problem using the information in this topic, follow the procedures below to collect relevant database and diagnostic information. After you have collected this information, contact Oracle Support.

- Collecting Cloud Tooling Logs
  Use the relevant log files that could assist Oracle Support for further investigation and resolution of a given issue.
- Collecting Oracle Diagnostics

Related Topics

- Oracle Support

Collecting Cloud Tooling Logs

Use the relevant log files that could assist Oracle Support for further investigation and resolution of a given issue.

DBAASCLI Logs

`/var/opt/oracle/log/dbaascli`

- `dbaascli.log`

Collecting Oracle Diagnostics

To collect the relevant Oracle diagnostic information and logs, run the `dbaascli diag collect` command.

For more information about the usage of this utility, see *DBAAS Tooling: Using dbaascli to Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check*.

Related Topics

- DBAAS Tooling: Using dbaascli to Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check

Standby Database Fails to Restart After Switchover in Oracle Database 11g Oracle Data Guard Setup

Description: After performing the switchover, the new standby (old primary) database remains shut down and fails to restart.

Action: After performing switchover, do the following:

1. Restart the standby database using the `srvctl start database -db <standby dbname>` command.
2. Reload the listener as `grid` user on all primary and standby cluster nodes.
   - To reload the listener using high availability, download and apply patch **25075940** to the Grid home, and then run `lsnrctl reload -with_ha`.
   - To reload the listener, run `lsrnctl reload`.

After reloading the listener, verify that the `<dbname>_DGMGRL` services are loaded into the listener using the `lsnrctl status` command.

To download patch 25075940

1. Log in to My Oracle Support.
2. Click **Patches & Updates**.

3. Select **Bug Number** from the **Number/Name or Bug Number (Simple)** drop-down list.
4. Enter the bug number **34741066**, and then click **Search**.
5. From the search results, click the name of the latest patch.
   You will be redirected to the **Patch 34741066: LSNRCTL RELOAD -WITH_HA FAILED TO READ THE STATIC ENTRY IN LISTENER.ORA** page.
6. Click **Download**.