Oracle® Cloud

Using Oracle Globally Distributed Exadata Database on Exascale Infrastructure

G17320-05
February 2026

Oracle Cloud Using Oracle Globally Distributed Exadata Database on Exascale Infrastructure,

G17320-05

Contents

3 Create and Manage Oracle Globally Distributed Exadata Database on Exascale Infrastructure

4 Create and Manage Private Endpoints

1

Overview of Oracle Globally Distributed Exadata Database on Exascale Infrastructure

Learn about the Oracle Cloud Infrastructure Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) service.

The topics that follow explain key capabilities of Distributed ExaDB-XS and describe the concepts you need to know about the service.

About Oracle Globally Distributed Exadata Database on Exascale Infrastructure

Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) brings the power of sharded databases to Oracle Exadata Database on Exascale Infrastructure.

Distributed ExaDB-XS is a cloud-based, fully-managed database service that enables the sharding of data across globally distributed converged databases. It is designed to support large-scale, mission-critical applications. It is a highly available, fault-tolerant, and scalable database service that enables organizations to store and process massive amounts of data with high performance and reliability.

The Distributed ExaDB-XS is built on top of Oracle's Exascale software services technology, which further empowers Exadata to meet the most demanding corporate and cloud computing requirements by decoupling Oracle Database and GI clusters from the underlying Exadata storage servers. Exascale software services can manage a large fleet of Exadata storage servers connected by the Exadata RDMA Network Fabric, providing storage services to multiple GI clusters and databases while enabling:

- Secure sharing of storage resources with strict data isolation, allowing different users and databases to share a large pool of storage while ensuring that data is inaccessible to users without the appropriate privileges
- Flexible and dynamic storage provisioning for many users and databases
- Increased storage utilization and efficiency while reducing storage costs
- Sharing of otherwise idle storage processing resources to improve performance

For more information about Oracle Exadata Exascale on Oracle Cloud, see Oracle Cloud Exadata Database Service on Exascale Infrastructure .

For a detailed discussion of sharded database features, see Oracle Globally Distributed Database Overview for Oracle AI Database 26ai.

Globally Distributed Database Concepts

To gain a greater understanding of Globally Distributed Database concepts, familiarize yourself with the following terminology.

- **Catalog** - an Oracle Database that supports automated shard deployment, centralized management of the distributed database, and multi-shard queries.

  A Catalog serves following purposes:

  - Serves as an administrative server for the entire distributed database
  - Stores a gold copy of the database schema
  - Manages multi-shard queries with a multi-shard query coordinator
  - Stores a gold copy of duplicated table data
- **Shard** - A distributed database is a collection of **shards**.

  Each shard in a distributed database is an independent Oracle Database instance that hosts subset of the distributed database data. Shared storage is not required across the shards.

  Shards can all be placed in one region or can be placed in different regions. For best performance, it is recommended that you place all of the shards in one region.

  Shards are replicated for high availability and disaster recovery with Raft replication. For high availability, Raft replicated shards can be placed in different availability domains within a region.
- **Shardspace** - A shardspace is a shard that stores data corresponding to a range or list of key values in a user-managed data distribution configuration. A shardspace consists of a shard and its replica.
- **Shard director** - A network listener that enable high performance connection routing based on a sharding key. In addition, a shard director is a set of processes known collectively as a Global Service Manager (GSM) that acts as a regional listener for clients that connect to a Globally Distributed Database.

  The shard director maintains a current topology map of the distributed database. Based on the sharding key passed during a connection request, the director routes the connections to the appropriate shard.
- **Global service** - A database service that is used to access data in the distributed database.

  A global service is an extension to the notion of the traditional database service. All of the properties of traditional database services are supported for global services.

For more in depth information about distributed database components and schema objects see Architecture and Concepts in *Oracle Globally Distributed Database*.

Exadata Database Service on Exascale Infrastructure Concepts

Exadata Database Service on Exascale Infrastructure (ExaDB-XS) provides a cloud service experience similar to Exadata Database Service on Dedicated Infrastructure. You can start with a small virtual machine (VM) cluster, and easily scale as needs grow. Oracle manages all of the physical infrastructure in a shared multitenancy infrastructure service model. Exascale is the underlying technology that serves as the foundation for this service.

Storage for database files resides in an Oracle Exadata Exascale Storage Vault. The Storage Vault provides high performance and scalable Exadata smart storage. Storage can be scaled online as needed, with a single command, and that storage becomes available for immediate use. Unlike Dedicated Infrastructure Exadata Database Service on Exascale Infrastructure does not require you to manage adding storage servers to the system, or manage storage allocations.

The architecture consists of the following elements:

1. A single Exascale Vault, which provides storage for the databases
2. A set of VMs run on Oracle-managed multitenant physical database servers
3. VM filesystems, which are centrally hosted by Oracle
4. A virtual client network (VCN), which provides client and backup network connectivity

Data Replication Solutions

Oracle's Globally Distributed Database services offer data replication solutions to ensure high availability, disaster recovery, and additional scalability for reads.

Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) offers catalog replication with Oracle Data Guard. Raft replication is available for the shards.

Distributed ExaDB-XS automatically deploys the specified replication topology to the procured systems, and enables data replication.

Catalog Replication with Oracle Data Guard

The catalog is a database. Oracle Data Guard replication to a physical standby databases can be used to provide high availability. Replication is automatically configured and deployed when the sharded database is created.

Oracle Data Guard is tightly integrated with Distributed ExaDB-XS to provide high availability and disaster recovery with strict data consistency and zero data loss. Oracle Data Guard replication maintains a synchronized copy (standby database) of the catalog (the primary database) for high availability and data protection. A standby can be deployed locally or remotely.

Chunk Set-level Replication with Raft Replication

Instead of replication at the whole shard level using additional databases for standbys, the Raft replication feature in a Distributed ExaDB-XS creates sets of chunks of data from each shard and distributes them automatically among the shards to handle chunk assignment, chunk movement, workload distribution, and balancing upon scaling (addition or removal of shards), including planned or unplanned shard availability changes.

Raft replication is built into the Distributed ExaDB-XS to provide a consensus-based, high-performance, low-overhead availability solution, with distributed replicas and fast failover with zero data loss, while automatically maintaining the replication factor if shards fail. With Raft replication management overhead does not increase with the number of shards. If you are used to NoSQL databases and do not expect to know anything about how replication works, native replication just works.

Unlike Data Guard replication, Raft replication does not need to be reconfigured when shards are added or removed, and replicas do not need to be actively managed.

For more details about how Raft replication works see Using Raft Replication in Oracle Globally Distributed Database.

Resource Identifiers

Oracle's Globally Distributed Database services resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID).

Distributed database resources are listed in the following table.

Resource	Identifier
Distributed Database	`osddistributeddb`
Distributed Database Private Endpoint	`osddistributeddbprivateendpoint`
OSD Work Request	`osdworkrequest`

For example, the OCID format for `osddistributeddb` is `ocid1.osddistributeddb.oc1.iad.<UNIQUE ID>`.

For information about the OCID format and other ways to identify your resources, see Resource Identifiers.

Metering and Billing

Metering and billing for Globally Distributed Exadata Database on Exascale Infrastructure is based on the number of ECPU per hour.

Because Globally Distributed Exadata Database on Exascale Infrastructure provisions Oracle Exadata Database Service on Exascale Infrastructure virtual machine (VM) clusters, billing and metering is handled by Oracle Exadata Database Service on Exascale Infrastructure usage. However, the cluster will be billed under the Globally Distributed Database SKU. See Metering Frequency and Per-Second Billingfor details.

> **Note**
>
> Once you tag a cluster for use in a Globally Distributed Database, it will continue to bill for the Globally Distributed Database SKU until the cluster is deleted.

Service Limits

Service Limits specific to Globally Distributed Exadata Database on Exascale Infrastructure can be set for Distributed Database Count and Distributed Database Private Endpoint Count.

Exadata Database virtual machine instances, ECPU count, and storage need to have limits set for the Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS).

To understand the resource capacity of the ExaDB-XS service, see Capacity Limits for Exadata Database Service on Exascale Infrastructure.

Integrated Services

Oracle's Globally Distributed Database services are integrated with various Oracle Cloud Infrastructure services and features.

- IAM
- Work Requests
- Monitoring

IAM

Oracle Globally Distributed Database services are integrated with the Identity and Access Management (IAM) service for authentication and authorization for the Console, SDK, CLI, and REST API.

To learn more about IAM, see IAM Overview.

Work Requests

Oracle's Globally Distributed Database services use their own APIs for Work Requests.

The permissions required for using the APIs are documented in Permissions for Globally Distributed Database APIs.

Monitoring

Oracle Cloud Infrastructure Monitoring lets you actively and passively monitor your Globally Distributed Database resources and alarms.

Globally Distributed Database metrics capture CPU utilization, OCPU consumption, memory utilization, deployment health, and inbound and outbound lag. You can view these metrics using the Monitoring service.

See Monitoring a Globally Distributed Database for more details about monitoring the health and performance of a distributed database.

2

Getting Started With Globally Distributed Exadata Database on Exascale Infrastructure

The following topics give you the information and prerequisites you need to get started with Globally Distributed Exadata Database on Exascale Infrastructure.

- Configuring the Tenancy
  Before you can use Oracle's Globally Distributed Database services to create and manage a distributed database, you must perform these preparatory tasks to organize your tenancy, create policies for the various resources, and then procure and configure the network, security, and infrastructure resources.
- Interfaces for Oracle Globally Distributed Exadata Database on Exascale Infrastructure
  You can use the Oracle Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) service through the Oracle Cloud Interface Console (a browser based interface), REST APIs, or Oracle Cloud Infrastructure Software Development Kits and Command Line Interface.

Configuring the Tenancy

Before you can use Oracle's Globally Distributed Database services to create and manage a distributed database, you must perform these preparatory tasks to organize your tenancy, create policies for the various resources, and then procure and configure the network, security, and infrastructure resources.

- Task 1. Subscribe to Ashburn Region
- Task 2. Create Compartments
- Task 3. Create User Access Constraints
- Task 4. Configure Network Resources
- Task 5. Configure Security Resources
- Task 6. Create Exadata Resources
- (Optional) Create API Key and User Constraints

Task 1. Subscribe to Ashburn Region

As the tenant administrator, subscribe to Ashburn (IAD) region and all of the regions required to run your Globally Distributed Exadata Database on Exascale Infrastructure implementation.

1. Subscribe to the Ashburn (IAD) region.
   - To use the service, you must subscribe to the Ashburn region.
   - Your tenancy Home Region does not have to be the Ashburn region, but you must subscribe to the Ashburn region to use Oracle's Globally Distributed Database services.
2. Subscribe to any other region where you will be placing a database.

- Subscribe to any regions where you plan to place databases for your implementation; this includes databases for the catalog, shards, and optional Oracle Data Guard standby database for the catalog.

> **Note**
>
> Globally Distributed Exadata Database on Exascale Infrastructure only supports creating shards in two regions. Also, for the best performance, distributed databases using Raft replication should have shards in the same region.

For more information, see Managing Regions.

Task 2. Create Compartments

As the tenant administrator, create compartments in your tenancy for all of the resources required by Oracle Globally Distributed Exadata Database on Exascale Infrastructure service.

Oracle recommends the following structure, and these compartments are referenced throughout the setup tasks:

- A "parent" compartment for the entire deployment. This is **gdd** in the examples.
- "Child" compartments for each of the various kinds of resources:
  - **gdd_certs_vaults_keys** for certificate authorities, certificates, certificate bundles, vaults, and keys
  - **gdd_databases** for databases, VM clusters, VCNs, subnets, private endpoints, and Globally Distributed Exadata Database on Exascale Infrastructure resources.
  - **gdd_instances** for compute instances for application servers (edge node/jump host to act as bastion to connect to the database)

The resulting compartment structure will resemble the following:

```
tenant /
     gdd /
          gdd_certs_vaults_keys
          gdd_databases
          gdd_instances
```

For more information, see Working with Compartments.

Task 3. Create User Access Constraints

Formulate an access control plan, and then institute it by creating appropriate IAM (Identity and Access Management) resources. Accordingly, access control within a distributed database is implemented at various levels, which are defined by the groups and policies here.

The user groups, dynamic groups, and policies described in the following tables should guide the creation of your own user access control plan for your distributed database implementation.

As the tenant administrator, create the following recommended groups, dynamic groups, and policies to grant permissions to the previously defined roles. The examples and documentation links assume that your tenancy uses identity domains.

- Understanding Role Separation

- Dynamic Groups
- User Groups
- Policies

Understanding Role Separation

You need to ensure that your cloud users have access to use and create only the appropriate kinds of cloud resources to perform their job duties. A best practice is to define roles for the purposes of role separation.

The roles and responsibilities described in the following table should guide your understanding of how to define user groups, dynamic groups, and policies for your Distributed ExaDB-XS implementation. The example roles presented here are used throughout the environment setup, resource creation, and management instructions.

Roles	Responsibilities
Tenant administrator	Subscribe to regions Create compartments Create dynamic groups, user groups, and policies
Infrastructure administrator	Create/Update/Delete virtual-network-family Create/Update/Delete Exadata Infrastructure Create/Update/Delete Exadata VM Clusters Tag Exadata VM Clusters Create/Update/Delete Globally Distributed Database Private Endpoints
Certificate administrator	Create/Update/Delete Vault Create/Update/Delete Keys
User	Create and manage Globally Distributed Databases using UI and APIs

Dynamic Groups

Create the following dynamic groups to control access to resources created in the Oracle Globally Distributed Exadata Database on Exascale Infrastructure compartments.

See Creating a Dynamic Group for instructions.

Dynamic Group Name	Description	Rules
gdd-cas-dg	Certificate authority resources	All resource.type='certificateauthority' resource.compartment.id = 'OCID of compartment tenant root / gdd / gdd_certs_vaults_keys'
gdd-clusters-dg	Exadata Database VM cluster resources	All resource.compartment.id = 'OCID of compartment tenant root / gdd / gdd_databases'

Dynamic Group Name	Description	Rules
gdd-instances-dg	Compute instance resources	All resource.compartment.id = 'OCID of compartment tenant root / gdd / gdd_instances'

User Groups

Create the following groups to give users permissions to use resources in the Globally Distributed Database compartments.

See Creating a Group for instructions.

User Group Name	Description
gdd-certificate-admins	Certificate administrators that create and manage keys and vaults.
gdd-infrastructure-admins	Infrastructure administrators that create and manage cloud network and infrastructure resources
gdd-users	Users that create and manage Globally Distributed Database resources using the APIs and UI

Policies

Create IAM policies to grant the groups access to resources created in the compartments for your Oracle Globally Distributed Exadata Database on Exascale Infrastructure tenancy.

Note that there is more than one Globally Distributed Database service on Oracle Cloud. These policies are specific to the Globally Distributed Exadata Database on Exascale Infrastructure service.

The following example policies, which are based on the compartment structure and groups created previously, should guide the creation of your own IAM policies for your implementation.

The identity domain (for example, Default) should be the identity domain you created the groups in.

See Creating a Policy for instructions.

gdd-certificate-admins-tenant-level

- Description: Tenant-level privileges for group gdd-certificate-admins
- Compartment: tenant
- Statements:

```
Allow group 'Default' / 'gdd-certificate-admins' to INSPECT tenancies in
tenancy
Allow group 'Default' / 'gdd-certificate-admins' to INSPECT work-requests
in tenancy
```

gdd-infrastructure-admins-tenant-level

- Description: Tenant-level privileges for group gdd-infrastructure-admins

- Compartment: tenant
- Statements:

```
Allow group 'Default' / 'gdd-infrastructure-admins' to INSPECT tenancies
in tenancy
Allow group 'Default' / 'gdd-infrastructure-admins' to INSPECT work-
requests in tenancy
Allow group 'Default' / 'gdd-infrastructure-admins' to READ limits in
tenancy
Allow group 'Default' / 'gdd-infrastructure-admins' to READ tag-namespaces
in tenancy
```

gdd-users-tenant-level

- Description: Tenant-level privileges for group gdd-users
- Compartment: tenant
- Statements:

```
Allow group 'Default' / 'gdd-users' to INSPECT tenancies in tenancy
Allow group 'Default' / 'gdd-users' to INSPECT work-requests in tenancy
Allow group 'Default' / 'gdd-users' to READ limits in tenancy
Allow group 'Default' / 'gdd-users' to READ Distributed-database in tenancy
Allow group 'Default' / 'gdd-users' to READ tag-namespaces in tenancy
```

gdd-certificate-admins

- Description: Compartment-level privileges for group gdd-certificate-admins
- Compartment: tenant/gdd
- Statements:

```
Allow group 'Default' / 'gdd-certificate-admins' to MANAGE keys in
compartment gdd
Allow group 'Default' / 'gdd-certificate-admins' to MANAGE vaults in
compartment gdd
```

gdd-infrastructure-admins

- Description: Compartment-level privileges for group gdd-infrastructure-admins
- Compartment: tenant/gdd
- Statements:

```
Allow group 'Default' / 'gdd-infrastructure-admins' to MANAGE exadb-vm-
clusters in compartment gdd
Allow group 'Default' / 'gdd-infrastructure-admins' to MANAGE instance-
family in compartment gdd
Allow group 'Default' / 'gdd-infrastructure-admins' to MANAGE Distributed-
database in compartment gdd
Allow group 'Default' / 'gdd-infrastructure-admins' to MANAGE tags in
compartment gdd
Allow group 'Default' / 'gdd-infrastructure-admins' to MANAGE virtual-
network-family in compartment gdd
Allow group 'Default' / 'gdd-infrastructure-admins' to READ exascale-db-
storage-vaults in compartment gdd
```

```
Allow group 'Default' / 'gdd-infrastructure-admins' to READ distributed-
database-workrequest in compartment gdd
```

gdd-users

- Description: Compartment-level privileges for group gdd-users
- Compartment: tenant/gdd
- Statements:

```
Allow group 'Default' / 'gdd-users' to MANAGE database-family in
compartment gdd
Allow service database to manage recovery-service-family in compartment
gdd
Allow service rcs to manage recovery-service-family in compartment gdd
Allow group 'Default' / 'gdd-users' to manage objects in compartment gdd
Allow group 'Default' / 'gdd-users' to read buckets in compartment gdd
Allow group 'Default' / 'gdd-users' to USE exadb-vm-clusters in
compartment gdd
Allow group 'Default' / 'gdd-users' to MANAGE instance-family in
compartment gdd
Allow group 'Default' / 'gdd-users' to MANAGE distributed-database in
compartment gdd
Allow group 'Default' / 'gdd-users' to MANAGE tags in compartment gdd
Allow group 'Default' / 'gdd-users' to READ dns-records in compartment gdd
Allow group 'Default' / 'gdd-users' to READ dns-zone in compartment gdd
Allow group 'Default' / 'gdd-users' to READ keys in compartment gdd
Allow group 'Default' / 'gdd-users' to READ distributed-database-
workrequest in compartment gdd
Allow group 'Default' / 'gdd-users' to READ vaults in compartment gdd
Allow group 'Default' / 'gdd-users' to READ vcns in compartment gdd
Allow group 'Default' / 'gdd-users' to USE network-security-groups in
compartment gdd
Allow group 'Default' / 'gdd-users' to USE private-ips in compartment gdd
Allow group 'Default' / 'gdd-users' to USE subnets in compartment gdd
Allow group 'Default' / 'gdd-users' to USE vnics in compartment gdd
Allow group 'Default' / 'gdd-users' to USE volumes in compartment gdd
```

gdd-dg-cas

- Description: Compartment-level privileges for dynamic group gdd-cas-dg
- Compartment: tenant/gdd
- Statements:

```
Allow dynamic-group 'Default' / 'gdd-cas-dg' to MANAGE objects in
compartment gdd
Allow dynamic-group 'Default' / 'gdd-cas-dg' to USE keys in compartment gdd
```

gdd-dg-clusters

- Description: Compartment-level privileges for dynamic group gdd-clusters-dg
- Compartment: tenant/gdd

- Statements:

```
Allow dynamic-group 'Default' / 'gdd-clusters-dg' to MANAGE keys in
compartment gdd_certs_vaults_keys
Allow dynamic-group 'Default' / 'gdd-clusters-dg' to READ vaults in
compartment gdd_certs_vaults_keys
```

gdd-kms

- Description: Compartment-level privileges for Key Management Service
- Compartment: tenant/gdd
- Statements:

```
Allow service keymanagementservice to MANAGE vaults in compartment
gdd_certs_vaults_keys
```

Task 4. Configure Network Resources

As the infrastructure administrator, create the network resources and enable the connectivity needed by the distributed database.

Example resources are named throughout these instructions to simplify tracking and relationships. For example, the name "gdd_iad" refers to the VCN created in the Ashburn (IAD) region.

- Common Network Resources
- Additional Network Resources Based on Your Topology

Common Network Resources

All Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) implementations require a VCN, subnet, and a private endpoint in the Ashburn (IAD) region.

As the infrastructure administrator, create the resources as described in the following table.

Resource	Instructions
Virtual Cloud Network (VCN)+ subnet	In Ashburn (IAD), create VCN gdd_iad and subnet gdd_subnet. This VCN and subnet are required to enable connectivity between the Distributed ExaDB-XS service and databases in the topology. Use the following values: • Compartment = gdd / gdd_databases • Region = Ashburn (IAD) • Subnet name = gdd_subnet • Subnet Type = Regional The subnet must be regional, spanning all availability domains See VCNs and Subnets for steps to create them.

Resource	Instructions
Private Endpoint	Create a private endpoint in the Ashburn (IAD) region to enable connectivity between the Distributed ExaDB-XS service and the databases in the topology. 1. Open the navigation menu, click**Oracle Database**, then click **Globally Distributed Exadata Database on Exascale Infrastructure**. 2. Click **Private Endpoints** in the navigation pane. 3. Click **Create private endpoint**. 4. Enter the following information. • **Name:** For example gdd_pe • **Compartment:** gdd/gdd_databases This should be the compartment containing the Ashburn region subnet you created above. • **Subnet:** gdd_subnet If you don't see the subnet listed, verify that it was created as a **Regional** subnet. • **Virtual cloud network:** gdd_iad • **Add tags (optional):** you can select tags for this resource by clicking Show Tagging Options. See Create and Manage Private Endpoints For more information about this resource.

Additional Network Resources Based on Your Topology

Depending on your Oracle Globally Distributed Exadata Database on Exascale Infrastructure topology, create additional network resources as described below.

Note that databases for the topology include the catalog, shards, and optional Oracle Data Guard standby database for the catalog.

All network resources should be created in the gdd/gdd_databases compartment.

Use Case	Network Resources	Peering and Connectivity
All databases are placed in the Ashburn (IAD) region	Create a subnet and service gateway in Ashburn (IAD) region for your Cloud Exadata Database VM Clusters. • In region Ashburn (IAD), create subnet osd-databases-subnet-iad in VCN gdd_iad. • In region Ashburn (IAD), create service gateway gdd_sgw_iad	Required Peering None Required Connectivity Unrestricted connectivity with subnet gdd_subnet
All databases are placed in a single region, R1, that is not Ashburn (IAD)*	Create a subnet and service gateway in the region for your Cloud Exadata Database VM Clusters. • In region R1, create VCN gdd_R1 with subnet osd-database-subnet-R1 • In region R1, create service gateway gdd_sgw_R1	Required Peering gdd_iad ↔ gdd_R1 Required Connectivity Unrestricted between gdd_iad.gdd_subnet and gdd_R1.osd-database-subnet-R1

Use Case	Network Resources	Peering and Connectivity
Databases are placed in multiple regions R1, R2, ..., RN	Create subnets and service gateways in each region for your Cloud Exadata Database VM Clusters. Subnet: • In region R1, create VCN gdd_R1 with subnet osd-database-subnet-R1 • In region R2, create VCN gdd_R2 with subnet osd-database-subnet-R2 ... • In region Rn, create VCN gdd_Rn with subnet osd-database-subnet-Rn Service gateways: • In region R1, create service Gateway gdd_sgw_R1 • In region R2, create Service gateway gdd_sgw_R2 ... • In region Rn, create service Gateway gdd_sgw_Rn	Required Peering gdd_iad ↔ gdd_R1 gdd_iad ↔ gdd_R2 gdd_iad ↔ gdd_Rn gdd_R1 ↔ gdd_R2 gdd_R1 ↔ gdd_Rn gdd_R2 ↔ gdd_Rn Required Connectivity Unrestricted and bi-directional between gdd_iad.gdd_subnet and gdd_R1.osd-database-subnet-R1 gdd_R2.osd-database-subnet-R2 gdd_Rn.osd-database-subnet-Rn Unrestricted and bi-directional between gdd_R1.osd-database-subnet-R1 and gdd_R2.osd-database-subnet-R2 gdd_Rn.osd-database-subnet-Rn Unrestricted and bi-directional between gdd_R2.osd-database-subnet-R2 and gdd_Rn.osd-database-subnet-Rn

*The Oracle Globally Distributed Exadata Database on Exascale Infrastructure service control plane exists only in the Ashburn (IAD) region. The private endpoint your created in a previous step in the Ashburn (IAD) region is used to communicate with the distributed database resources in their respective regions.

Instructions for creating the resources are available at:

- VCNs and Subnets
- Creating a Service Gateway
- Peering VCNs in different regions through a DRG

Task 5. Configure Security Resources

All security resources are created in the gdd/gdd_certs_vaults_keys compartment.

> ⚠ **Caution**
>
> After creating a distributed database that references a key, you cannot move the vault or keys to a new compartment without also restarting the container databases that reference the moved vault or key.

- Create a Vault
- Create a TDE Key

Create a Vault

Create a vault in the gdd/gdd_certs_vaults_keys compartment for the Transparent Data Encryption (TDE) master encryption keys in the region where the shard databases will reside.

For example, in region R1, create vault gdd_vault_R1

For details about creating a vault, see Creating a Vault.

Create a TDE Key

Create the master encryption key to access the database.

For example, create master encryption key gdd_TDE_key-oraspace in vault gdd_vault_R1 with the following attributes.

- Protection Mode = Software
- Key Shape: Algorithm = AES
- Length = 256

For details about creating master encryption keys, see Create a Master Encryption Key.

Task 6. Create Exadata Resources

As the infrastructure administrator, configure the Oracle Globally Distributed Exadata Database on Exascale Infrastructure topology in the following steps.

- Import Oracle-ApplicationName Tag Namespace
- Create Exadata VM Clusters on Exascale Infrastructure

Import Oracle-ApplicationName Tag Namespace

Import the Oracle-ApplicationName tag namespace in the root compartment of your tenancy.

1. From the Cloud console navigation menu, select **Governance & Administration**, then **Tag Namespaces** (under the Tenancy Management category).
2. In the Tag Namespaces panel, check if the Oracle-ApplicationName namespace exists in the root compartment of your tenancy.

   Make sure the root compartment of your tenancy is selected under **List Scope**.
3. If you don't see Oracle-ApplicationName in the list, do the following:
   a. Click **Import Standard Tags** (located above the list).
   b. Select the checkbox next to the Oracle-ApplicationName namespace and click **Import**.

Create Exadata VM Clusters on Exascale Infrastructure

Create a VM cluster using Exadata Database Service on Exascale Infrastructure service for the catalog database, optional Data Guard standby catalog datbase, and each shard database, you plan to deploy in the Distributed ExaDB-XS topology.

While creating the VM clusters, use the following requirements and recommendations:

- It is recommended that you use **one VM cluster per database** (shard, catalog, catalog standby).

  One shard database and a catalog database can be co-located on a given VM cluster. However, using a common VM cluster for both the catalog and shard has the potential to cause a processing bottleneck.
- Create **single node** clusters. Only single node clusters are supported.
- **Compartment:** Create the VM clusters in your tenancy's gdd/gdd_clusters compartment.
- **ECPUs enabled per VM:** Enable 8 ECPUs for the VM clusters intended for shards.
- **Database storage:** You can use same vault for all VM clusters (catalog and shards), as long as you configure a minimum of 500GB in storage capacity for every database in the topology.

  For example, if you have 3 shards and 1 catalog then the total minimum storage needed is 500GB x 4 = 2000GB. In this case you create a single vault with a minimum of 2000GB storage capacity.
- **Time zone:** Set all of the VM clusters in the topology to the same time zone. This setting is under Advanced options.
- **Tags:**Add the defined tag `Oracle-ApplicationName.Other_Oracle_Application: Sharding`. This setting is under Advanced options.

  Before you can add the tag, you must import the tag's namespace as described in Import Oracle-ApplicationName Tag Namespace.

> **Note**
>
> Once you tag a VM cluster for use in a distributed database, it will continue to bill for that SKU until the VM cluster is deleted.

See Manage VM Clusters for steps to create the clusters.

(Optional) Create API Key and User Constraints

Create an OCI API key pair if you intend to directly use the Globally Distributed Database REST API, OCI Software Development Kits, and Command Line Interface.

Follow the instructions in Required Keys and OCIDs.

If you want to set user controls on the APIs see Permissions for Globally Distributed Database APIs.

Interfaces for Oracle Globally Distributed Exadata Database on Exascale Infrastructure

You can use the Oracle Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) service through the Oracle Cloud Interface Console (a browser based interface), REST APIs, or Oracle Cloud Infrastructure Software Development Kits and Command Line Interface.

Using the Console

To access Distributed ExaDB-XS using the Console:

1. Use a supported browser to access the Console.

   See Signing In to the Console for details.
2. Enter your cloud tenant, user name, and password, when prompted.
3. Click **Sign in**.
4. In the upper-right corner of the window, select a region that offers the Distributed ExaDB-XS service enabled; for example, **US East (Ashburn)**.
5. From the navigation menu, select **Oracle Database**, then **Globally Distributed Exadata Database on Exascale Infrastructure**.

   The home page for Globally Distributed Exadata Database on Exascale Infrastructure is displayed.

Using Globally Distributed Database APIs

You can find the complete Globally Distributed Database REST API reference at https://docs.oracle.com/iaas/api/#/en/globally-distributed-database/latest/

See REST APIs and Software Development Kits and Command Line Interface for more information about using REST APIs and the OCI Software Development Kits and Command Line Interface.

3

Create and Manage Oracle Globally Distributed Exadata Database on Exascale Infrastructure

You create a Globally Distributed Exadata Database on Exascale Infrastructure configuration, which is used as a blueprint for the service to procure VMs, deploy the Globally Distributed Exadata Database on Exascale Infrastructure software components on systems you designate in the configuration and start required services. You can then monitor and perform life cycle operations on the database.

The topics that follow explain how to configure, deploy, and perform operations on Globally Distributed Exadata Database on Exascale Infrastructure.

- Creation and Deployment Workflow
  To get started with Oracle Globally Distributed Exadata Database on Exascale Infrastructure) (Distributed ExaDB-XS, you must create the configuration, validate it, and then deploy the configuration.
- Creating a Globally Distributed Exadata Database on Exascale Infrastructure Resource
  A Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) resource contains the connectivity and configuration details of the shards and shard catalog databases.
- Retrying Creation of Distributed Database Resources
  Before the Configure Sharding operation, in some cases, you can retry the creation of Globally Distributed Database resources rather than deleting the distributed database and starting over.
- Listing Globally Distributed Databases
- Viewing Details for Distributed ExaDB-XS
  You view Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) configuration, backup, and maintenance information by going to its details page.
- Adding Catalog Data Guard Replication
  In Globally Distributed Exadata Database on Exascale Infrastructure you can add an Oracle Data Guard standby database to the catalog for high availability and disaster prevention.
- Deploying Globally Distributed Exadata Database on Exascale Infrastructure
  You deploy a Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) when you are ready to start running your completed configuration or any time you make changes to the configuration, such as adding a shard.
- Connecting to the Database
  You need the connection string to connect to the databases in the Distributed ExaDB-XS.
- Adding Shards
  Add shards to scale out your Globally Distributed Database.
- Modifying Shards
  You can modify a shard before you deploy the Distributed ExaDB-XS configuration.

- Terminating (Deleting) a Shard
  Terminating a shard in a Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) configuration permanently deletes it and removes all automatic backups.
- Stopping a Globally Distributed Exadata Database on Exascale Infrastructure
- Starting a Globally Distributed Exadata Database on Exascale Infrastructure
- Terminating (Deleting) a Globally Distributed Exadata Database on Exascale Infrastructure
  Terminating Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) permanently deletes it and removes all automatic backups.
- Managing Raft Replication
  To run Raft replication operations on your Globally Distributed Database you must create a node where you can run GDSCTL commands.
- Moving Globally Distributed Exadata Database on Exascale Infrastructure Resources
  You can move a Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) from one compartment to another.
- Updating the Display Name
  You can change the display name of a Globally Distributed Exadata Database on Exascale Infrastructure from its details page.
- Managing Tags
  Tags help you locate resources within your tenancy.
- Globally Distributed Database REST APIs
  The following REST APIs are used to interact with the Globally Distributed Database (distributed-database) resource.

Creation and Deployment Workflow

To get started with Oracle Globally Distributed Exadata Database on Exascale Infrastructure) (Distributed ExaDB-XS, you must create the configuration, validate it, and then deploy the configuration.

Task	Description	More Information
Create the Distributed ExaDB-XS configuration	Configure the connectivity, security, and topology details of the shards and shard catalog databases.	Creating a Globally Distributed Exadata Database on Exascale Infrastructure Resource

Task	Description	More Information
Deploy the Distributed ExaDB-XS	Deploy the configuration and start the services.	Deploying Globally Distributed Exadata Database on Exascale Infrastructure

Note

Deployment must take place within 7 days of completing the creation operation, or you must terminate the resources and start again.

Creating a Globally Distributed Exadata Database on Exascale Infrastructure Resource

A Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) resource contains the connectivity and configuration details of the shards and shard catalog databases.

You create the resource in the Globally Distributed Exadata Database on Exascale Infrastructure list page. If you need help finding the list page, see Listing Globally Distributed Exadata Database on Exascale Infrastructure Resources.

1. On the Globally Distributed Exadata Database on Exascale Infrastructure list page, select **Create Distributed ExaDB-XS**.
2. Choose which creation method you want to use:

- **Quick create:** Requires that Exadata VM clusters are precreated before you begin creating the distributed database resource. Once Exadata VM clusters are available, proceed with Creating Distributed ExaDB-XS Using Quick Create.
- **Custom create:** Lets you configure the clusters and vault storage as part of the distributed database creation process. Proceed with Creating Distributed ExaDB-XS Using Custom Create.

Creating Distributed ExaDB-XS Using Quick Create

1. Provide the following basic information.

Setting	Description and Notes
Display name	Enter a user-friendly description or other information that helps you easily identify the Distributed ExaDB-XS. Avoid entering confidential information. You can modify this name after resource creation.
Database name prefix	This prefix is appended to all of the database names in the configuration for ease of use.
Compartment	Select a compartment to host the Distributed ExaDB-XS resource
Database version	Oracle Database release 26ai is supported at this time.

2. Optionally, add tags to the Distributed ExaDB-XS resource. These can also be added after creation.
3. Enter the following information:

Setting	Description
Data distribution	**Automated** - Data is automatically distributed across shards using partitioning by consistent hash. The partitioning algorithm evenly and randomly distributes data across shards. **User managed** - not currently supported
Replication type	**Raft** replication creates replication units consisting of sets of chunks and distributes them automatically among the shards to handle chunk assignment, chunk movement, workload distribution, and balancing upon scaling.
Replication factor	In Raft replication, the replication factor is the number of replicas in a replication unit. This number includes the primary (**leader**) member of the unit and its replicas (**followers**).
Shard count	The number of shards is determined by the replication factor.

In **Shards configuration** you can configure shards using the map view or list view.

- On the **map view**, select the region where you want the database shards to be deployed, then select **Configure Shards** to enter the settings.

- In the **list view**, the settings are presented in the Create Globally Distributed Exadata Database on Exascale Infrastructure page.

Shards List

The shards in the list view and map view are pre-populated with shard names in the home tenancy, and you must edit the shards to add more required information.

To configure a shard's VM cluster placement:

- In the list view, select the **Edit** action in the action menu (…).
- In the map view, you can select one or more regions where clusters are available, then select **View/Edit** or **Configure Shards**.

If all of your shards will be in the same region, availability domain, and subnet, enable **Apply same settings to all shards**.

You can select **Add Shard** to have up to 10 shards in the list. You can also remove shards, but note that the shard count must be greater than the Raft replication factor.

4. Configure the shard catalog in **Catalog configuration**.

   You can choose to use the same configuration that is applied to the shards, or slide the **Same as Shard's configuration** toggle and make selections that apply only to the catalog database.

   Note that Raft replication does not apply to the catalog. Data protection for the catalog is configured after the Globally Distributed Database is created. See Adding Catalog Data Guard Replication.

   Select **Edit** in the Actions menu to configure the catalog settings described below.

Setting	Description
Region	Select the region to host the catalog If Data Guard is enabled this is the primary catalog database region.
VM cluster	Choose a VM cluster to host the catalog database. **Note** It is recommended that you use one VM cluster per database (shard, catalog, or catalog standby).

5. Configure the remaining settings.

Setting	Description and Notes
Create administrator credentials	Set the SYS user password. This user will be able to access all of the shard databases and catalog databases in the configuration.
Encryption key	Select the vault and master encryption key that were configured in Task 5. Configure Security Resources. All shards will have the same encryption vault and encryption key.

Setting	Description and Notes
Select private endpoint	Select the private endpoint that was created for this Distributed ExaDB-XS in Common Network Resources.
Select character sets	Select the Character set and National character set that will be used in all of the shard and shard catalog databases. The AL32UTF8 character set is recommended by default for character set and the AL16UTF16 character set is recommended by default for National character set.
Select ports	Enter the **GSM listener port**, **ONS port (local)**, and **ONS port (remote)** . **Note** The **ONS port (remote)** number must be unique to each Globally Distributed Database. Do not reuse a port number used in another Globally Distributed Database unless a delete operation is fully processed on the original.
Advanced options: Shard configuration - Chunks	Under Advanced Options you can optionally configure the number of chunks per shard.
Advanced options: Shard configuration - Replication unit	Displays the number of Raft replication units that will be created. A distributed database with Raft replication contains multiple **replication units**. A replication unit (RU) is a set of chunks that have the same replication topology. Each RU consists of a leader and replicas and those are placed on different shards.
Advanced options: Configure database backups	Under Advanced Options you can enable and schedule automated database backups. See Exadata Database Service on Exascale Infrastructure documentation for information about the settings.

6. Select **Validate** to let Distributed ExaDB-XS run validation tests against the configuration.
7. Once any validation errors are addressed and validation is successful, select **Create**.

   Now the Distributed ExaDB-XS display name appears in the list while the creation operation runs.

   Creation can take a while, because several tasks are performed as part of the create operation, including host procurement, installing software, and generating certificates for the shard directors (GSMs).

   You can monitor the operation status in the State column and select the Distributed ExaDB-XS display name to track progress in the **Work requests** tab.

   When the status of all of the shards on the **Shards** tab is Available, Distributed ExaDB-XS creation is complete and successful.

> ⚠ **Caution**
>
> After creating a Distributed ExaDB-XS, do not move any of its vaults or keys or the Distributed ExaDB-XS will not work.

Creating Distributed ExaDB-XS Using Custom Create

1. Provide the following basic information.

Setting	Description and Notes
Compartment	Select a compartment to host the Distributed ExaDB-XS resource
Display name	Enter a user-friendly description or other information that helps you easily identify the Distributed ExaDB-XS. Avoid entering confidential information. You can modify this name after resource creation.
Database name prefix	This prefix is appended to all of the database names in the configuration for ease of use.
Database version	Oracle AI Database release 26ai is supported at this time.

2. Select your **License type**.
3. Optionally, add tags to the Distributed ExaDB-XS resource. These can also be added after creation.
4. Enter the following information:

Setting	Description
Data distribution	**Automated** - Data is automatically distributed across shards using partitioning by consistent hash. The partitioning algorithm evenly and randomly distributes data across shards. **User managed** - not currently supported
Replication type	**Raft** replication creates replication units consisting of sets of chunks and distributes them automatically among the shards to handle chunk assignment, chunk movement, workload distribution, and balancing upon scaling.
Replication factor	In Raft replication, the replication factor is the number of replicas in a replication unit. This number includes the primary (**leader**) member of the unit and its replicas (**followers**).
Shard count	The number of shards is determined by the replication factor.

Setting	Description
ECPU count	Enter the number of ECPU cores to enable for each VM cluster. Specify the number of ECPUs as an integer. Available cores are subject to your tenancy's service limits. You must enter a minimum of 8 ECPUs. ECPUs are based on the number of cores, elastically allocated, from the shared pool of Exadata database servers and storage servers. Aggregated ECPU consumption on a given cluster is 1.5 times the ECPU count. Note that a number of ECPUs are consumed in overhead and are not available to the shards. See Oracle Cloud Infrastructure Documentation for more information.
Vault storage capacity	GB of storage to allocate to the shard (database) You can use same vault for all VM clusters (catalog and shards), as long as you configure a minimum of 500GB in storage capacity for every database in the topology. For example, if you have 3 shards and 1 catalog then the total minimum storage needed is 500GB x 4 = 2000GB. In this case you create a single vault with a minimum of 2000GB storage capacity.

In **Shards configuration** you can configure shards using the map view or list view.

- On the **map view**, select the region where you want the database shards to be deployed, then select **Configure Shards** to enter the settings.
- In the **list view**, the settings are presented in the Create Globally Distributed Exadata Database on Exascale Infrastructure page.

Shards List

The shards in the list view and map view are pre-populated with shard names in the home tenancy, and you must edit the shards to add more required information.

To configure a shard's region placement, availability domain, and subnet:

- In the list view, select the **Edit** action in the action menu (…).
- In the map view, you can select one or more regions, then select **View/Edit** or **Configure Shards**.

If all of your shards will be in the same region, availability domain, and subnet, enable **Apply same settings to all shards**.

You can select **Add Shard** to have up to 10 shards in the list. You can also remove shards, but note that the shard count must be greater than the Raft replication factor.

5. Configure the shard catalog in **Catalog configuration**.

   You can choose to use the same configuration that is applied to the shards using **Same as Shard's configuration**, or make modifications that apply only to the catalog database.

   Select **Edit** in the actions menu (…) to configure the catalog availability domain and subnet. Note that the catalog region will be US East (Ashburn).

   Note that Raft replication does not apply to the catalog. Data protection for the catalog is configured after the Globally Distributed Database is created. See Adding Catalog Data Guard Replication.

6. Configure the remaining settings.

Setting	Description and Notes
Create administrator credentials	Set the SYS user password to access all of the shard databases and catalog databases in the configuration.
SSH Keys	Generate, upload, or paste the SSH keys, and save these keys for later database connection.
Encryption key	Select the vault and master encryption key that were configured in Task 5. Configure Security Resources.
Select private endpoint	Select the private endpoint that was created for this Distributed ExaDB-XS in Common Network Resources.
Advanced options: Shard configuration - Chunks	Under Advanced Options you can optionally configure the number of chunks per shard.
Advanced options: Shard configuration - Replication unit	Displays the number of Raft replication units that will be created. A distributed database with Raft replication contains multiple **replication units**. A replication unit (RU) is a set of chunks that have the same replication topology. Each RU consists of a leader and replicas and those are placed on different shards.
Advanced options: Extended VM cluster capacity	Configure extended capacity in ECPUs per VM.
Advanced options: VM file system storage	Specify file system storage capacity per VM.
Advanced options: Smart flash cache	Specify percentage of storage capacity to add as smart flash cache.
Advanced options: Select character sets	Select the Character set and National character set that will be used in all of the shard and shard catalog databases. The AL32UTF8 character set is recommended by default for character set and the AL16UTF16 character set is recommended by default for National character set.
Select ports	Enter the **GSM listener port**, **ONS port (local)**, **ONS port (remote)**, and **SCAN listener port**. **Note** The **ONS port (remote)** number must be unique to each Globally Distributed Database. Do not reuse a port number used in another Globally Distributed Database unless a delete operation is fully processed on the original.
Advanced options: Diagnostics collection	Enable diagnostic events, health monitoring, and logs.
Advanced options: Database backups	Enable and schedule automated database backups. See Exadata Database Service on Exascale Infrastructure documentation for information about the settings.

7. Select **Validate** to let Distributed ExaDB-XS run validation tests against the configuration.
8. Once any validation errors are addressed and validation is successful, select **Create** to create the resources, VM clusters, and so on.

   Now the Distributed ExaDB-XS display name appears in the list while the creation operation runs.

   Creation can take a while, because several tasks are performed as part of the create operation, including host procurement, VM deployment, installing software, and generating certificates for the shard directors (GSMs).

   You can monitor the operation status in the State column and select the Distributed ExaDB-XS display name to track progress in the **Work requests** tab.

   When the status of all of the shards on the **Shards** tab is Available, Distributed ExaDB-XS creation is complete and successful.

> ⚠ **Caution**
>
> After creating a Distributed ExaDB-XS, do not move any of its vaults or keys or the Distributed ExaDB-XS will not work.

Retrying Creation of Distributed Database Resources

Before the Configure Sharding operation, in some cases, you can retry the creation of Globally Distributed Database resources rather than deleting the distributed database and starting over.

You can retry creating these resources under these conditions:

- Configure Sharding operation has not yet been run
- The Globally Distributed Database is in a FAILED or INACTIVE state
- A shard, catalog, or GSM is in FAILED state

Retrying Shards and Catalog Creation

- In the **Shards** or **Catalog** tab of the distributed database details page, select **Retry create** in the actions menu (three dots) for any shard or catalog in a failed state.

Retrying GSM Creation

1. Open the **Actions** menu at the top of the distributed database details page, and select **More actions** then **Configure GSM**.
2. In the GSM Configuration panel, select **Retry create** in the actions menu (three dots) for any GSM in a failed state. Confirm the operation in the dialog.

Listing Globally Distributed Databases

- Listing Globally Distributed Exadata Database on Exascale Infrastructure Resources

Listing Globally Distributed Exadata Database on Exascale Infrastructure Resources

Open the **navigation menu** and select **Oracle Database**. Then select **Globally Distributed Exadata Database on Exascale Infrastructure**.

> **Note**
>
> The **navigation menu** is the main menu located in the upper-left corner of the Oracle Cloud Console. Use the menu to navigate to OCI services, dashboards, and marketplace.

The list of distributed databases is shown by default.

Viewing Details for Distributed ExaDB-XS

You view Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) configuration, backup, and maintenance information by going to its details page.

On the Globally Distributed Exadata Database on Exascale Infrastructure list page, select the Distributed ExaDB-XS you want to view.

On the details page you can view the following information:

- **Distributed database information** shows a summary of the configuration.
- **Shards** displays a list of all of the shards with their configuration settings.

  Select **Additional shard details** in the action menu for more details.

  Select **View Replication Units** to see the status of the replication unit leaders and followers on each shard.
- **Catalog** displays the configuration settings for the catalog database. In the catalog action menu:

  Additional catalog details shows more details.

  View Data Guard shows details about the Data Guard standby, if configured.
- **Resource map** displays a visual representation of the shards and catalog configuration. Hover over the numbers to see some configuration information, including availability domain (AD), and click the information icon for more details.
- **Replication details** displays a list of the replication units by ID number and each member is labeled with the shard it resides on.
- **Latest backups** displays backups associated with the shards.
- **Work requests** displays the status of ongoing operations on the databases.
- **Tags** displays any tags added to this distributed database.

Adding Catalog Data Guard Replication

In Globally Distributed Exadata Database on Exascale Infrastructure you can add an Oracle Data Guard standby database to the catalog for high availability and disaster prevention.

1. In the details page for the Distributed Database, select the **Catalog** tab, and select **Add standby**.
2. Select a configuration type:
   - **Quick add:** Use an existing cluster for the standby database.
   - **Custom add:** The service will deploy a new cluster for the standby database.
3. Enter the following details:

Setting	Description
Peer VM Cluster	Select the VM cluster to use for the standby database. **Note** It is recommended that you use one VM cluster per database (shard, catalog, or catalog standby).
Standby database password	Enter the standby database ADMIN password. This must be the same as the primary database ADMIN password.

Deploying Globally Distributed Exadata Database on Exascale Infrastructure

You deploy a Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) when you are ready to start running your completed configuration or any time you make changes to the configuration, such as adding a shard.

Note

Deployment must take place within 7 days of completing the operation in Creating a Globally Distributed Exadata Database on Exascale Infrastructure Resource or Adding Shards, or you must terminate the resources and start again.

1. On the Globally Distributed Exadata Database on Exascale Infrastructure list page, select the Distributed ExaDB-XS for which you want to complete the deployment.
2. On the details page, select **Configure Sharding**.
3. Optionally, select **Rebalance** to automatically redistribute data among the shards, if this is not the initial deployment of this Distributed ExaDB-XS.
   This is typically done after adding or removing shards from the configuration in case of Automated Sharding type.

4. Select **Configure Sharding** to start the deployment.

Connecting to the Database

You need the connection string to connect to the databases in the Distributed ExaDB-XS.

Oracle client credentials (wallet files) are downloaded from Distributed ExaDB-XS by a service administrator. If you are not a Distributed ExaDB-XS administrator, your administrator can provide you with the client credentials.

1. On the Globally Distributed Exadata Database on Exascale Infrastructure list page, select the Distributed ExaDB-XS.
2. On the details page select **Database connection**.

   Connection strings are displayed in the list.

Adding Shards

Add shards to scale out your Globally Distributed Database.

You can add shards when:

- You have completed Creating a Globally Distributed Exadata Database on Exascale Infrastructure Resource, but have not yet completed Deploying Globally Distributed Exadata Database on Exascale Infrastructure.
- You have completed Deploying Globally Distributed Exadata Database on Exascale Infrastructure and want to scale up your Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) with more shards.

1. On the Globally Distributed Exadata Database on Exascale Infrastructure list page select a Distributed ExaDB-XS.
2. On the details page, under **Shards**, select **Add Shard**.
3. On the **Add shard** pane configure the new shard.
   - **Shard count** - enter the number of shards you want to add.

     You can add up to 10 shards in each set to deploy, and then add more after deployment if needed.
   - **Shard** - Shows the display name for each shard or shardspace in the configuration. Once you select a region the name is populated.
   - **Region** - Select the region where you would like to host the shard.
   - **VM cluster** - Select a cluster available in the selected region.

     > **Note**
     >
     > It is recommended that you use one VM cluster per database (shard or catalog).

4. In **Create administrator credentials**, set the password for the shard database ADMIN user.
5. Select the **Encryption key** details for the new shards.

Select the compartment, vault, and master encryption key that were configured in Task 5. Configure Security Resources.

6. Click **Validate** to run checks to make sure the new shards are valid.
7. Once any validation errors are addressed and validation is successful, click **Add Shards** to deploy the new shards.

> **Note**
>
> There is a time limit for deploying new shards.
>
> - When scaling up a deployed Distributed ExaDB-XS, you must complete Deploying Globally Distributed Exadata Database on Exascale Infrastructure within 7 days of completing this procedure or you will get an error and must terminate the new shard resources and start again.
> - When adding shards to an undeployed Distributed ExaDB-XS, you have 7 days from completing Creating a Globally Distributed Exadata Database on Exascale Infrastructure Resource to add any shards and complete Deploying Globally Distributed Exadata Database on Exascale Infrastructure.

For more information about the concepts and considerations of adding shards to a Distributed ExaDB-XS see Shard Management in *Using Oracle Globally Distributed Database*.

Modifying Shards

You can modify a shard before you deploy the Distributed ExaDB-XS configuration.

You can modify shards in a Distributed ExaDB-XS from its **Details** page.

1. Go to the **Details** page of the Distributed ExaDB-XS in which you want to modify a shard.
2. In the **Shards** section of the details page, select **Edit** from the Actions (three dots) menu for the shard you want to make changes to.

   On the **Edit Shard** pane you can configure the Region and VM cluster settings.
3. Click **Edit** to save the changes to the shard.

Terminating (Deleting) a Shard

Terminating a shard in a Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) configuration permanently deletes it and removes all automatic backups.

You cannot recover a terminated shard.

For more information about the concepts and considerations of removing shards see Shard Management in *Using Oracle Sharding*.

1. On the Globally Distributed Exadata Database on Exascale Infrastructure list page, select a Distributed ExaDB-XS.
2. On the Details page, in the **Shards** section, select the shard, and then select **Terminate**.
3. On the **Terminate shard** dialog enter the Distributed ExaDB-XS name to confirm that you want to remove the shard.

4. Click **Remove**.

Stopping a Globally Distributed Exadata Database on Exascale Infrastructure

> **Note**
>
> When you stop a Distributed ExaDB-XS, the following details apply:
>
> - Tools are no longer able to connect to the database.
> - In-flight database transactions and queries are stopped.
> - ECPU billing is halted.

1. On the Globally Distributed Exadata Database on Exascale Infrastructure list page, select the Distributed ExaDB-XS.
2. On the details page, select **More actions** and then select **Stop**.
3. Click **Stop** to confirm.

Starting a Globally Distributed Exadata Database on Exascale Infrastructure

> **Note**
>
> When you start Distributed ExaDB-XS, Exascale Infrastructure billing is initiated. See Metering and Billing for details.

1. On the Globally Distributed Exadata Database on Exascale Infrastructure list page, select the Distributed ExaDB-XS to start.
2. On the details page, select **More actions** and then select **Start**.

   Start is only enabled for a stopped Distributed ExaDB-XS that is successfully deployed.
3. Click **Start** to confirm.

Terminating (Deleting) a Globally Distributed Exadata Database on Exascale Infrastructure

Terminating Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) permanently deletes it and removes all automatic backups.

You cannot recover a terminated Distributed ExaDB-XS.

1. There are two ways to initiate a Distributed ExaDB-XS delete:

- On the Globally Distributed Exadata Database on Exascale Infrastructure list page, open the actions menu (…) for the Distributed ExaDB-XS and select **Terminate**.
- On the Globally Distributed Exadata Database on Exascale Infrastructure list page, select a Distributed ExaDB-XS, and on the details page, select **More actions** and then select **Terminate**.

2. On the Terminate page enter the Distributed ExaDB-XS name to confirm that you want to terminate that Distributed ExaDB-XS.
3. Click **Terminate**.

Managing Raft Replication

To run Raft replication operations on your Globally Distributed Database you must create a node where you can run GDSCTL commands.

More information about Raft replication operations can be found in the *Oracle Database Globally Distributed Database Guide* at Raft Replication Operations.

- Create a GDSCTL Node
  For many Raft replication life cycle management operations, you will need to create a node to run GDSCTL commands for Raft operations.
- Moving Replication Units
  You can move follower members of replication units from one shard to another.

Create a GDSCTL Node

For many Raft replication life cycle management operations, you will need to create a node to run GDSCTL commands for Raft operations.

1. Open the Actions menu in the Oracle Globally Distributed Exadata Database on Exascale Infrastructure details page, and select **More actions**, then **Create Distributed Database CLI Shell**.
2. Upload your public key for this node.

   If needed, you can generate keys in this panel.
3. If required, specify the subnet under **Advanced Options**.
4. Click **Create**.

Moving Replication Units

You can move follower members of replication units from one shard to another.

In the **Shards** tab or **Replication details** tab of the Globally Distributed Database details page, select **Move Replication unit** in the actions menu.

- **Source shard:** shard from which you want to move the replication unit follower.
- **Destination shard:** shard to which you want to move the replication unit follower.
- **Replication unit:** the replication unit numbers listed here identify which replication units have followers on the source shard that can be moved.

Moving Globally Distributed Exadata Database on Exascale Infrastructure Resources

You can move a Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) from one compartment to another.

> **⚠ Caution**
>
> If you need to move a Distributed ExaDB-XS resource, please contact Oracle customer support first. There may be unintended consequences to moving any resource within the Distributed ExaDB-XS configuration. See Moving Resources to a Different Compartment for more information.

> **ⓘ Note**
>
> - Move resource is not allowed if any GSM, shard, or catalog is in a failed state.
> - As soon as you move the Distributed ExaDB-XS to a different compartment, the policies that govern the new compartment apply immediately and affect access to the database. Therefore, your access to the database may change, depending on the policies governing your Oracle Cloud user account's access to resources.
>
> After the Distributed ExaDB-XS move to a new compartment is successful, any work request logs associated with the Distributed ExaDB-XS from the original compartment are no longer available.

To move Distributed ExaDB-XS you must have the right to manage Distributed ExaDB-XS in its current compartment and in the compartment you are moving it to.

1. On the Globally Distributed Exadata Database on Exascale Infrastructure list page, select a Distributed ExaDB-XS you want to move.
2. Select the **Move resource** action.
3. In the **Move to a different compartment** dialog, select the compartment to move the Distributed ExaDB-XS to.
4. Click **Move** to start the operation.

Updating the Display Name

You can change the display name of a Globally Distributed Exadata Database on Exascale Infrastructure from its details page.

1. On the Globally Distributed Exadata Database on Exascale Infrastructure list page, select a Distributed ExaDB-XS.
2. On the details page, select **Update display name** from the actions.
3. Enter the new display name in the **New display name** field.
4. Enter the current name in the field below to confirm the name change.
5. Click **Update display name**.

Managing Tags

Tags help you locate resources within your tenancy.

You can add and view tags from the Globally Distributed Exadata Database on Exascale Infrastructure (Distributed ExaDB-XS) list page and details page.

For directions to get to the details page, see Viewing Details for Distributed ExaDB-XS.

Adding Tags

On the list page, in the Actions (three dots) menu, select **Add Tags**.

On the details page, select **Add Tags** from the **Actions** menu.

Editing and Deleting Tags

In the Distributed ExaDB-XS details page, open the **Tags** section to view, edit, and delete tags.

More Information About Tags

See Managing Tags and Tag Namespaces to learn more about tagging.

Globally Distributed Database REST APIs

The following REST APIs are used to interact with the Globally Distributed Database (distributed-database) resource.

REST API	Description
`AddDistributedDatabaseGdsControlNode`	Adds a new Global Data Services control node for running GDSCTL commands on the distributed database
`ChangeDistributedDatabaseCompartment`	Moves the distributed database and its dependent resources to the specified compartment
`ChangeDistributedDbBackupConfig`	Change the backup configuration for the specified distributed database
`ConfigureDistributedDatabaseGsms`	Configure GSMs for the distributed database
`ConfigureDistributedDatabaseSharding`	Initiates the distributed database configuration deployment
`CreateDistributedDatabase`	Creates a new distributed database resource
`DeleteDistributedDatabase`	Deletes the distributed database
`DownloadDistributedDatabaseGsmCertificateSigningRequest`	Generates the common certificate signing request for the distributed database GSM instances.
`GenerateDistributedDatabaseGsmCertificateSigningRequest`	Generates the certificate signing request for distributed database GSM instances. Once the certificate signing request is generated, you can download it using `DownloadGsmCertificateSigningRequest`.
`GenerateDistributedDatabaseWallet`	Generates the wallet for application connections with a distributed database.
`GetDistributedDatabase`	Gets the details of a distributed database resource

REST API	Description
`PatchDistributedDatabase`	Lets you add, remove, or update shards in the distributed database topology. You can add, remove, or update multiple shards in a single patch operation; however, combinations of inserts, updates, and removes in a single operation are not allowed. This operation is followed by `ConfigureDistributedDatabaseSharding` to deploy the updated distributed database configuration.
`RotateDistributedDatabasePasswords`	Rotate passwords for different components of the distributed database
`StartDistributedDatabase`	Starts the specified distributed database
`StopDistributedDatabase`	Stops the specified distributed database
`UpdateDistributedDatabase`	Lets you change the display name and edit tags associated with a distributed database resource.
`UploadDistributedDatabaseSignedCertificateAndGenerateWallet`	Uploads the CA signed certificate to the distributed database GSM instances, and generate wallets for the GSM instances.
`ValidateDistributedDatabaseNetwork`	Validates the network connectivity between components of the distributed database
`ListDistributedDatabases`	Gets a list of distributed database resources

See Private Endpoint REST APIs for descriptions of those APIs.

See Globally Distributed Database Policies for API permissions and policy guidelines.

4

Create and Manage Private Endpoints

A private endpoint is required in the Ashburn region to connect Oracle Cloud databases running in a customer VCN to the Globally Distributed Database services.

You create the private endpoint as part of setting up your network resources in Task 4. Configure Network Resources. For general information about private endpoints, see About Private Endpoints.

The topics that follow describe the steps for creating a private endpoint for a Globally Distributed Database and the life cycle operations on an existing private endpoint.

- Creating a Private Endpoint
- Listing Private Endpoints
- Viewing Private Endpoint Details
- Editing Private Endpoints
- Moving Private Endpoints
- Private Endpoint REST APIs

Creating a Private Endpoint

You create a private endpoint in the Private Endpoints list page. To find the Private Endpoints list page, see Listing Private Endpoints.

1. In the **Private Endpoints** list page select **Create private endpoint**.
2. In the Create private endpoint panel, enter the following information.
   - **Name:** Enter a name.
   - **Description:** Optionally, enter a description.
   - **Choose compartment:** Choose the compartment containing the Ashburn region subnet that you created in Task 4. Configure Network Resources.
   - **Subnet in *compartment*:** Choose the subnet you created in Task 4. Configure Network Resources.
   - **Virtual cloud network in *compartment*:** Select a VCN
3. Optionally, you can select tags for this resource by clicking **Show Tagging Options**.

Listing Private Endpoints

- Listing Private Endpoints for Globally Distributed Exadata Database on Exascale Infrastructure

Listing Private Endpoints for Globally Distributed Exadata Database on Exascale Infrastructure

1. Open the **navigation menu** and select **Oracle Database**. Then select **Globally Distributed Exadata Database on Exascale Infrastructure**.

   > **Note**
   >
   > The **navigation menu** is the main menu located in the upper-left corner of the Oracle Cloud Console. Use the menu to navigate to OCI services, dashboards, and marketplace.

2. On the left side of the screen, select **Private Endpoints**.

   A list of existing private endpoints is displayed.

Viewing Private Endpoint Details

To find a private endpoint's details, go the the Private Endpoints list page and select a private endpoint from the list. To find the Private Endpoints list page, see Listing Private Endpoints.

You can find information about private endpoints, run operations, and make changes on the Private Endpoint Details page for each private endpoint resource.

At the top of the details page there are buttons to run operations on the private endpoint, such as update the display name, move resource, add tags, and terminate. On this page there are also sections (tabs) which show configuration information and tags.

The details page also lets you view private endpoint-related related Work Requests and any Distributed Databases that use this private endpoint.

Editing Private Endpoints

You can edit a private endpoint in the Private Endpoints list page. To find the Private Endpoints list page, see Listing Private Endpoints.

In the list, select **Edit private endpoint** from the Actions (three dots) menu for the private endpoint you want to make changes to.

You can change the name and description of the private endpoint.

Moving Private Endpoints

You can move a private endpoint resource from one compartment to another.

1. In the Private Endpoints list page, select **Move Resource** from the Actions (three dots) menu for the private endpoint you want to move.

   To find the Private Endpoints list page, see Listing Private Endpoints.
   You can also select **Move Resource** on the private endpoint's details page.

2. In the **Move resource** dialog, select the compartment to move the private endpoint to from the dropdown.
3. Click **Move Resource**.

After you move the private endpoint to the new compartment, inherent policies apply immediately and may affect access to the private endpoint through the Console. For more information, see Managing Compartments.

Private Endpoint REST APIs

The following REST APIs are used to interact with the Distributed Database Private Endpoint resource.

These APIs are documented in the Globally Distributed Database REST API reference at https://docs.oracle.com/iaas/api/#/en/globally-distributed-database/latest/PrivateEndpoint/.

REST API	Description
`ChangeDistributedDatabasePrivateEndpointCompartment`	Moves the private endpoint to the specified compartment.
`CreateDistributedDatabasePrivateEndpoint`	Creates a private endpoint.
`DeleteDistributedDatabasePrivateEndpoint`	Deletes a private endpoint.
`GetDistributedDatabasePrivateEndpoint`	Gets a private endpoint.
`ReinstateProxyInstance`	Reinstates the proxy instance associated with the private endpoint
`UpdateDistributedDatabasePrivateEndpoint`	Updates private endpoint configuration details.
`ListDistributedDatabasePrivateEndpoints`	Lists private endpoints.

See Globally Distributed Database Policies for API permissions and policy guidelines.

5

Monitoring a Globally Distributed Database

Monitoring Work Requests

Globally Distributed Databases use their own APIs for Work Requests.

Using the Console:

Work request status is displayed in a Globally Distributed Database's details page.

From the Globally Distributed Database list page, click any database name and go to its details page. To find the Globally Distributed Database list page, see Listing Globally Distributed Databases.

The **Work requests** section displays the status of ongoing operations.

Using the REST APIs

You can use the `GetWorkRequest` and `ListWorkRequests` APIs to get work request status.

See Work Request Reference for details.

Metrics for Globally Distributed Exadata Database on Exascale Infrastructure

Because Globally Distributed Exadata Database on Exascale Infrastructure is a collection of database instances and services, you monitor metrics for those resources which make up the Globally Distributed Exadata Database on Exascale Infrastructure topology.

Refer to the *Exadata Database Service on Exascale Infrastructure* documentation:

- Monitor Metrics for VM Cluster Resources
- Metrics for Oracle Exadata Database Service on Dedicated Infrastructure in the Monitoring Service

Events

A Globally Distributed Database emits events in Oracle Cloud Infrastructure (OCI), which are structured messages that indicate changes in the distributed database resource.

You can define rules in the OCI Event Service to get notified of events happening in an OCI native service and use the Notification Service (ONS) to send emails or other notifications from these events.

Table 5-1 Event Types for Globally Distributed Database

Friendly Name	Event Type
Distributed Database - Add GDSCTL Node Begin	com.oraclecloud.globaldb.adddistributeddatabase gdscontrolnode.begin
Distributed Database - Add GDSCTL Node End	com.oraclecloud.globaldb.adddistributeddatabase gdscontrolnode.end
Distributed Database - Change Compartment Begin	com.oraclecloud.globaldb.changedistributeddatab asecompartment.begin
Distributed Database - Change Compartment End	com.oraclecloud.globaldb.changedistributeddatab asecompartment.end
Distributed Database - Configure Sharding Begin	com.oraclecloud.globaldb.configuredistributedda tabasesharding.begin
Distributed Database - Configure Sharding End	com.oraclecloud.globaldb.configuredistributedda tabasesharding.end
Distributed Database - Configure GSMs Begin	com.oraclecloud.globaldb.configuredistributedda tabasegsms.begin
Distributed Database - Configure GSMs End	com.oraclecloud.globaldb.configuredistributedda tabasegsms.end
Distributed Database - Create Begin	com.oraclecloud.globaldb.createdistributeddatab ase.begin
Distributed Database - Create End	com.oraclecloud.globaldb.createdistributeddatab ase.end
Distributed Database - Delete Begin	com.oraclecloud.globaldb.deletedistributeddatab ase.begin
Distributed Database - Delete End	com.oraclecloud.globaldb.deletedistributeddatab ase.end
Distributed Database - Download GSM Certificate Signing Request	com.oraclecloud.globaldb.downloaddistributeddat abasegsmcertificatesigningrequest
Distributed Database - Fetch VM Clusters	com.oraclecloud.globaldb.fetchdistributeddataba sevmclusters
Distributed Database - Generate GSM Certificate Signing Request Begin	com.oraclecloud.globaldb.generatedistributeddat abasegsmcertificatesigningrequest.begin
Distributed Database - Generate GSM Certificate Signing Request End	com.oraclecloud.globaldb.generatedistributeddat abasegsmcertificatesigningrequest.end
Distributed Database - Generate Wallet	com.oraclecloud.globaldb.generatedistributeddat abasewallet
Distributed Database - Patch Begin	com.oraclecloud.globaldb.patchdistributeddataba se.begin
Distributed Database - Patch End	com.oraclecloud.globaldb.patchdistributeddataba se.end
Distributed Database - Prevalidate	com.oraclecloud.globaldb.prevalidatedistributed database
Distributed Database - Start Begin	com.oraclecloud.globaldb.startdistributeddataba se.begin

Table 5-1 (Cont.) Event Types for Globally Distributed Database

Friendly Name	Event Type
Distributed Database - Start End	`com.oraclecloud.globaldb.startdistributeddatabase.end`
Distributed Database - Stop Begin	`com.oraclecloud.globaldb.stopdistributeddatabase.begin`
Distributed Database - Stop End	`com.oraclecloud.globaldb.stopdistributeddatabase.end`
Distributed Database - Update	`com.oraclecloud.globaldb.updatedistributeddatabase`
Distributed Database - Upload Signed Certificate And Generate Wallet Begin	`com.oraclecloud.globaldb.uploaddistributeddatabasesignedcertificateandgeneratewallet.begin`
Distributed Database - Upload Signed Certificate And Generate Wallet End	`com.oraclecloud.globaldb.uploaddistributeddatabasesignedcertificateandgeneratewallet.end`
Distributed Database - Validate Network Begin	`com.oraclecloud.globaldb.validatedistributeddatabasenetwork.begin`
Distributed Database - Validate Network End	`com.oraclecloud.globaldb.validatedistributeddatabasenetwork.end`

Table 5-2 Event Types for Globally Distributed Database Private Endpoint

Friendly Name	Event Type
Distributed Database Private Endpoint - Change Compartment Begin	`com.oraclecloud.globaldb.changedistributeddatabaseprivateendpointcompartment.begin`
Distributed Database Private Endpoint - Change Compartment End	`com.oraclecloud.globaldb.changedistributeddatabaseprivateendpointcompartment.end`
Distributed Database Private Endpoint - Create Begin	`com.oraclecloud.globaldb.createdistributeddatabaseprivateendpoint.begin`
Distributed Database Private Endpoint - Create End	`com.oraclecloud.globaldb.createdistributeddatabaseprivateendpoint.end`
Distributed Database Private Endpoint - Delete Begin	`com.oraclecloud.globaldb.deletedistributeddatabaseprivateendpoint.begin`
Distributed Database Private Endpoint - Delete End	`com.oraclecloud.globaldb.deletedistributeddatabaseprivateendpoint.end`
Distributed Database Private Endpoint - Update	`com.oraclecloud.globaldb.updatedistributeddatabaseprivateendpoint`

6

Globally Distributed Database Policies

To control access to Globally Distributed Database resources and the type of access each user group has, you must create policies.

- Giving Permissions to Users
- Required Policies
- Using Distributed Database Management Policy Builder Templates
  Several templates specific to Globally Distributed Database are included in the OCI Identity and Security Policy Builder.
- Resource-Types
- Resource-Permissions Model
- Permissions for Globally Distributed Database APIs
- Details for Verbs + Resource-Type Combinations
- Supported Variables

Giving Permissions to Users

Use IAM policies to grant certain capabilities to a Globally Distributed Database user group.

You can configure group and group permissions so that members can manage Globally Distributed Database resources.

Create user groups to manage Globally Distributed Database resources with role-based levels of access, and then add users that require access to these resources to the groups.

Remember that only resources within the same compartment can access each other, unless the proper permissions are granted. Ensure that you have the proper permissions to view and select the appropriate VCN and subnet when creating distributed databases.

Required Policies

Several users, groups, and policies are required to set up and run a Globally Distributed Database.

See Task 3. Create User Access Constraints for complete nistructions and lists.

Using Distributed Database Management Policy Builder Templates

Several templates specific to Globally Distributed Database are included in the OCI Identity and Security Policy Builder.

The templates associated with the Distributed Database Management policy use case fall into three categories: Tenant-level templates for all platforms, templates that apply to only Globally Distributed Autonomous AI Database deployments, and templates that apply to only Globally Distributed Exadata Database on Exascale Infrastructure deployments. These categories address policies required for different platforms.

Tenant-level templates for all platforms:

- "Let Certificate Admins access required resources in Tenancy" provides tenant-level privileges to certificate administrators that create and manage keys and vaults.
- "Let Infrastructure Admins access required resources in Tenancy" provides tenant-level privileges to infrastructure administrators that create and manage cloud network and infrastructure resources.
- "Let Users access required resources in Tenancy" provides tenant-level privileges to users that create and manage Globally Distributed Database resources using the APIs and UI. Note that users need to be allowed to READ either distributed-autonomous-database or distributed-database in this policy. You can remove the statement that does not apply to your deployment.
- "Let Certificate Authority Resources to manage Objects and use Keys for both Distributed Autonomous Database and Distributed Database" is meant to provide compartment-level privileges to a dynamic group to control certificate authority resources in a designated compartment.
- "Let VM Clusters Resources to manage Keys and read Vaults for both Distributed Autonomous Database and Distributed Database" is meant to provide compartment-level privileges to a dynamic group to control VM cluster resources, and compartment-level privileges to the Key Management Service or Oracle Key Vault in a specific compartment.

Templates for Globally Distributed Autonomous AI Database:

- "Let Certificate Admins create and manage Keys and Vaults for Distributed Autonomous Database" provides compartment-level privileges to certificate administrators that create and manage keys and vaults.
- "Let Infrastructure Admins create and manage Distributed Autonomous Database" provides compartment-level privileges to infrastructure administrators that create and manage cloud network and infrastructure resources.
- "Let Users create and manage Distributed Autonomous Database" provides compartment-level privileges to users that create and manage Globally Distributed Autonomous AI Database resources using the APIs and UI.

Templates for Globally Distributed Exascale Database on Exascale Infrastructure:

- "Let Certificate Admins create and manage Keys and Vaults for Distributed Database" provides compartment-level privileges to certificate administrators that create and manage keys and vaults.
- "Let Infrastructure Admins create and manage Distributed Database" provides compartment-level privileges to infrastructure administrators that create and manage cloud network and infrastructure resources.
- "Let Users create and manage Distributed Database" provides compartment-level privileges to users that create and manage Globally Distributed Autonomous AI Database resources using the APIs and UI.

See Task 3. Create User Access Constraints for more information about creating the recommended compartments, dynamic groups, user groups, and policies for the distributed database.

See Creating a Policy for more details about using the Policy Builder.

Resource-Types

Oracle's Globally Distributed Database services offer individual resource-types for writing policies.

Resource-Type	Description
`distributed-database`	Configuration of the Globally Distributed Database, including the data distribution model and information for connecting to the shards and catalog databases.
`distributed-database-privateendpoint`	A private endpoint used to connect databases running in a customer VCN to the Globally Distributed Database services.
`distributed-database-workrequest`	Monitor for long-running operations, such as shard creation, update, or deletion.

Resource-Permissions Model

Each resource defines its own permissions model. This permissions model forms the basis of how a policy is defined to allow for authorized access to distributed database resources.

These permissions are intended to be mapped to Operations (list, get, update delete, and so on) to allow for fine grained access control.

- **Read** (read-only)– allows the user to view resource details
- **Update** – grants View permission, plus allows the user to edit an existing resource, including move, add shard, remove shard
- **Create** – grants Update permission, plus allows the user to create new resources
- **Delete** – grants Create permission, plus allows the user to delete a resource

The following table details the permissions model for Oracle's Globally Distributed Database resources.

Resource	Permissions
`distributed-database`	• DISTRIBUTED_DB_INSPECT • DISTRIBUTED_DB_READ • DISTRIBUTED_DB_MANAGE • DISTRIBUTED_DB_MOVE • DISTRIBUTED_DB_CREATE • DISTRIBUTED_DB_DELETE

Resource	Permissions
distributed-database-privateendpoint	• DISTRIBUTED_DB_PRIVATE_ENDPOINT_INSPECT • DISTRIBUTED_DB_PRIVATE_ENDPOINT_READ • DISTRIBUTED_DB_PRIVATE_ENDPOINT_MANAGE • DISTRIBUTED_DB_PRIVATE_ENDPOINT_MOVE • DISTRIBUTED_DB_PRIVATE_ENDPOINT_CREATE • DISTRIBUTED_DB_PRIVATE_ENDPOINT_DELETE
distributed-database-work-requests	• DISTRIBUTED_DB_WORK_REQUEST_LIST • DISTRIBUTED_DB_WORK_REQUEST_READ

Permissions for Globally Distributed Database APIs

Here's a list of the API operations mapped to permissions for Globally Distributed Database, grouped by resource-type.

Distributed-database API permissions

API names and permissions for distributed-database resource-type

Table 6-1 Distributed-database API permissions

API Operation	Permission
AddDistributedDatabaseGdsControlNode	DISTRIBUTED_DB_MANAGE
ChangeDistributedDatabaseCompartment	DISTRIBUTED_DB_MOVE
ChangeDistributedDbBackupConfig	DISTRIBUTED_DB_MANAGE
ConfigureDistributedDatabaseGsms	DISTRIBUTED_DB_MANAGE
ConfigureDistributedDatabaseSharding	DISTRIBUTED_DB_MANAGE
CreateDistributedDatabase	DISTRIBUTED_DB_CREATE
DeleteDistributedDatabase	DISTRIBUTED_DB_DELETE
DownloadDistributedDatabaseGsmCertificateSigningRequest	DISTRIBUTED_DB_MANAGE
GenerateDistributedDatabaseGsmCertificateSigningRequest	DISTRIBUTED_DB_MANAGE
GenerateDistributedDatabaseWallet	DISTRIBUTED_DB_READ
GetDistributedDatabase	DISTRIBUTED_DB_READ
PatchDistributedDatabase	DISTRIBUTED_DB_MANAGE
RotateDistributedDatabasePasswords	DISTRIBUTED_DB_MANAGE

Table 6-1 (Cont.) Distributed-database API permissions

API Operation	Permission
StartDistributedDatabase	DISTRIBUTED_DB_MANAGE
StopDistributedDatabase	DISTRIBUTED_DB_MANAGE
UpdateDistributedDatabase	DISTRIBUTED_DB_MANAGE
UploadDistributedDatabaseSignedCertificateAndGenerateWallet	DISTRIBUTED_DB_MANAGE
ValidateDistributedDatabaseNetwork	DISTRIBUTED_DB_MANAGE
ListDistributedDatabases	DISTRIBUTED_DB_INSPECT

Distributed-database-privateendpoint API permissions

API names and permissions for distributed-database-privateendpoint resource-type

Table 6-2 Distributed-database-privateendpoint API permissions

API Operation	Permissions
ChangeDistributedDatabasePrivateEndpointCompartment	DISTRIBUTED_DB_PRIVATE_ENDPOINT_MOVE
CreateDistributedDatabasePrivateEndpoint	DISTRIBUTED_DB_PRIVATE_ENDPOINT_CREATE
DeleteDistributedDatabasePrivateEndpoint	DISTRIBUTED_DB_PRIVATE_ENDPOINT_DELETE
GetDistributedDatabasePrivateEndpoint	DISTRIBUTED_DB_PRIVATE_ENDPOINT_READ
ReinstateProxyInstance	DISTRIBUTED_DB_PRIVATE_ENDPOINT_MANAGE
UpdateDistributedDatabasePrivateEndpoint	DISTRIBUTED_DB_PRIVATE_ENDPOINT_MANAGE
ListDistributedDatabasePrivateEndpoints	DISTRIBUTED_DB_PRIVATE_ENDPOINT_INSPECT

Distributed-database-workrequest API permissions

API names and permissions for distributed-database-workrequest resource-type

Table 6-3 Distributed-database-workrequest API permissions

API Operation	Permission
GetWorkRequest	DISTRIBUTED_DB_WORK_REQUEST_READ
ListWorkRequests	DISTRIBUTED_DB_WORK_REQUEST_LIST
ListWorkRequestErrors	DISTRIBUTED_DB_WORK_REQUEST_READ
ListWorkRequestLogs	DISTRIBUTED_DB_WORK_REQUEST_READ

Details for Verbs + Resource-Type Combinations

There are various Oracle Cloud Infrastructure verbs and resource-types that you can use when you create a policy. The topics in this section show the permissions and API operations covered by each verb for Globally Distributed Database.

The level of access is cumulative as you go from `inspect` to `read` to `use` to `manage`.

- Distributed-database
- Distributed-database-privateendpoint
- Distributed-database-workrequest

Distributed-database

Permission	APIs Fully Covered
INSPECT	
DISTRIBUTED_DB_INSPECT	`ListDistributedDatabases`
READ	
INSPECT +	INSPECT+

DISTRIBUTED_DB_READ	`GetDistributedDatabase`
	`GenerateDistributedDatabaseWallet`
	`DownloadDistributedDatabaseGsmCertificateSigningRequest`

MANAGE	
READ +	READ +

DISTRIBUTED_DB_MANAGE	`UpdateDistributedDatabase`
	`PatchDistributedDatabase`
	`ConfigureDistributedDatabaseSharding`
	`RotateDistributedDatabasePasswords`
	`StartDistributedDatabase`
	`StopDistributedDatabase`
	`ChangeDistributedDbBackupConfig`
	`ValidateDistributedDatabaseNetwork`
	`ConfigureDistributedDatabaseGsms`
	`UploadDistributedDatabaseSignedCertificateAndGenerateWallet`
	`GenerateDistributedDatabaseGsmCertificateSigningRequest`
	`AddDistributedDatabaseGdsControlNode`

DISTRIBUTED_DB_MOVE	`ChangeDistributedDatabaseCompartment`
CREATE	
UPDATE+	UPDATE+
DISTRIBUTED_DB_CREATE	`CreateDistributedDatabase`
DELETE	
CREATE+	CREATE+
DISTRIBUTED_DB_DELETE	`DeleteDistributedDatabase`

Distributed-database-privateendpoint

Permission	APIs Fully Covered
INSPECT	
DISTRIBUTED_DB_PRIVATE_ENDPOINT_INSPECT	`ListDistributedDatabasePrivateEndpoints`
READ	
INSPECT +	INSPECT+
DISTRIBUTED_DB_PRIVATE_ENDPOINT_READ	`GetDistributedDatabasePrivateEndpoint`
MANAGE	
READ +	READ +
DISTRIBUTED_DB_PRIVATE_ENDPOINT_MANAGE	`UpdateDistributedDatabasePrivateEndpoint` `ReinstateProxyInstance`
DISTRIBUTED_DB_PRIVATE_ENDPOINT_MOVE	`ChangeDistributedDatabasePrivateEndpointCompartment`
CREATE	
UPDATE+	UPDATE+
DISTRIBUTED_DB_PRIVATE_ENDPOINT_CREATE	`CreateDistributedDatabasePrivateEndpoint`
DELETE	
CREATE+	CREATE+
DISTRIBUTED_DB_PRIVATE_ENDPOINT_DELETE	`DeleteDistributedDatabasePrivateEndpoint`

Distributed-database-workrequest

Permission	APIs Fully Covered
INSPECT	
DISTRIBUTED_DB_WORK_REQUEST_LIST	`ListWorkRequests`
READ	
INSPECT +	INSPECT+

DISTRIBUTED_DB_WORK_REQUEST_READ	`GetWorkRequest`
	`ListWorkRequestErrors`
	`ListWorkRequestLogs`

Supported Variables

When you add conditions to your policies, you can use either Globally Distributed Database general or service specific variables.

Oracle's Globally Distributed Database services support all general variables. For more information, see general variables for all requests.