aws

Oracle Database@AWS

Oracle Database@AWS: Oracle Database@AWS User Guide

Table of Contents

What is Oracle Database@AWS?

Oracle Database@AWS is an offering that enables you to access Oracle Exadata infrastructure managed by Oracle Cloud Infrastructure (OCI) inside AWS data centers. You can migrate your Oracle Exadata workloads, establish low-latency connectivity with applications running on AWS, and integrate with AWS services. You get a single invoice through AWS Marketplace, which counts towards AWS commitments and Oracle Support rewards.

The following diagram shows a high-level overview of an OCI region tied to an AWS data center that hosts Oracle Exadata infrastructure. Within an AWS Availability Zone (AZ), you can peer an Amazon VPC to a private network that is tied to the data center. By peering these networks, application servers in the VPC can access Oracle databases running on the Oracle Exadata infrastructure.

Features of Oracle Database@AWS

With Oracle Database@AWS, you benefit from the following features:

Reduced application latency

You can establish low-latency connectivity between Oracle Exadata and applications running on AWS. Proximity to applications hosted in AWS ensures minimal network delays and improved performance.

Cost optimization

You reduce on-premises Oracle Exadata costs by moving to a cloud-based solution. You also minimize data transfer fees by keeping databases and applications in the same AWS environment. You also get a single invoice through AWS Marketplace, which counts towards AWS commitments as well as Oracle Support rewards. You don't get a separate bill from OCI.

Easy data migration

You can migrate Oracle Exadata workloads into AWS with minimal application or license changes. You can use standard Oracle database migration tools such as Recovery Manager (RMAN), Oracle Data Guard, transportable tablespaces, Oracle Data Pump, Oracle GoldenGate, AWS DMS, and Oracle Zero Downtime Migration.

Flexible sizing

You can choose from different Exadata system sizes to match your workload requirements.

Seamless integration with AWS services

You can integrate with other AWS services and applications running in the same environment. For example, Oracle Database@AWS integrates with Amazon EC2, Amazon VPC, and IAM. You can also integrate with managed services such as Amazon SageMaker AI, and Amazon Bedrock.

Standard AWS interfaces

You use familiar AWS tools and interfaces to purchase, provision, and manage your Oracle Database@AWS resources. You can provision and manage your resources using AWS APIs, CLI, or SDKs. The AWS APIs call the corresponding OCI APIs necessary to provision and manage the resources.

Related AWS services

Oracle Database@AWS works with the following services to improve the availability and scalability of your Oracle database applications:

- **Amazon EC2** — Provides virtual servers that function as Oracle application servers. You can configure your load balancer to route traffic to your EC2 application servers. For more information, see the Amazon EC2 User Guide.

- **Amazon Virtual Private Cloud (VPC)** — Enables you to launch AWS resources in a logically isolated virtual network that you've defined. Oracle Exadata infrastructure resides in a special network called the ODB network that you can peer to a VPC. You can then run application servers in your VPC and access your Exadata databases. For more information, see the Amazon VPC User Guide.

- **Amazon CloudWatch** — Provides a monitoring service for Oracle Database@AWS. OCI gathers metric data about your Oracle Exadata system and sends it to CloudWatch. For more information, see Monitoring Oracle Database@AWS with Amazon CloudWatch.

- **AWS Identity and Access Management (IAM)** — Helps you securely control access to Oracle Database@AWS resources for your users. Use IAM to control who can use your AWS resources (authentication) and what resources users can use in which ways (authorization). For more information, see Identity and access management for Oracle Database@AWS.

- **AWS analytics services** — Provide a broad and cost-effective set of analytics services to help you gain insights faster from your Exadata database. Each service is purpose-built for a wide range of analytics use cases such as interactive analysis, big data processing, data warehousing, real-time analytics, operational analytics, dashboards, and visualizations. For more information, see Analytics on AWS.

- **AWS artificial intelligence (AI) services** — Integrate with your Oracle applications to address common use cases such as personalized recommendations, modernizing your contact center, improving safety and security, and increasing customer engagement. For more information, see AI Services.

Accessing Oracle Database@AWS

You can create, access, and manage Oracle Database@AWS using the AWS Management Console. It provides a web interface that you can use to access Oracle Database@AWS.

Pricing for Oracle Database@AWS

You can purchase Oracle Database@AWS offerings from AWS Marketplace. You first contact an Oracle sales representative. Oracle then makes the offer available to you in AWS Marketplace based on the private pricing agreement. Your AWS bill shows charges based on your usage.

There are no data transfer charges when your Oracle application and Oracle database are hosted in the same Availability Zone (AZ). Standard data transfer charges apply for communication between AZs.

What's next?

You're now ready to begin creating your Oracle Database@AWS resources.

1. Learn about how Oracle Database@AWS works. For more information, see How Oracle Database@AWS works.

   > **ⓘ Note**
   >
   > If you're familiar with AWS and Oracle Exadata and want to get started right away, skip this step.

2. Request a private offer for Oracle Database@AWS through the AWS Management Console, and then accept the offer. For more information, see Request a private offer for Oracle Database@AWS.

   > **ⓘ Note**
   >
   > To request a private offer in this preview, you must contact AWS to get your AWS account added to an allow list.

3. Create your ODB network, Oracle Exadata infrastructure, and Exadata VM clusters using the AWS console. Create your Exadata databases using OCI tools. For more information, see Getting started with Oracle Database@AWS.

How Oracle Database@AWS works

Oracle Database@AWS integrates Oracle Cloud Infrastructure (OCI) with the AWS Cloud. In the following sections, you can learn about the key components of this multicloud architecture.

Oracle Exadata Database Service on Dedicated Infrastructure is an OCI service that provides Exadata Database Machine. Oracle Exadata Database Machine is an integrated, preconfigured, and pretested full-stack platform for use in enterprise data centers. You create the Oracle Exadata infrastructure in an AWS Availability Zone (AZ) using the AWS console. Then you use OCI APIs to create and manage the Oracle Exadata databases. These Exadata databases are accessible to Amazon EC2 application servers running in an Amazon VPC.

The following diagram shows the Oracle Database@AWS architecture.

OCI child sites

Oracle Cloud Infrastructure is hosted in OCI regions and availability domains. An *OCI region* consists of *OCI availability domains (ADs)*, which are isolated data center clusters within an OCI

region. An *OCI child site* is a data center that extends an OCI availability domain to an Availability Zone (AZ) in an AWS Region. The Exadata infrastructure logically resides in an OCI region and physically resides in an AWS Region.

The OCI child site for Oracle Database@AWS physically resides in an AWS data center. AWS hosts the Exadata infrastructure, and OCI provisions and maintains the Exadata infrastructure hardware inside the data center. You configure the Exadata infrastructure, private network, and VM clusters using the AWS console. Because the infrastructure exists within the AWS Cloud, you can use services such as Amazon EC2 and Amazon VPC to allow application access.

Oracle Exadata infrastructure

The Oracle Exadata infrastructure is the underlying architecture of database servers, storage servers, and networking that runs Oracle Exadata databases. You create VM clusters on Exadata infrastructure using AWS APIs.

The Oracle Exadata infrastructure is distributed on physical machines called database servers. These servers provide the compute resources, analogous to Amazon EC2 dedicated servers.

Each database server hosts one or more virtual machines (VMs) running on a hypervisor. For architectural diagrams, see [Exadata Database Service on Dedicated Infrastructure Technical Architecture](#).

When you create Exadata infrastructure, you specify information such as the following:

- The total number of database servers

- The total number of storage servers

- The Availability Zone (AZ) that hosts the infrastructure

> ⓘ **Note**
>
> You can only deploy Oracle Database@AWS in the AWS Region US East (N. Virginia). The only supported physical AZs within this Region have the physical IDs `use1-az4` and `use1-az6`. To find the logical zone names in your account that map to the physical zone IDs, run the following command.
>
> ```
> aws ec2 describe-availability-zones --region us-east-1 --query
>   "AvailabilityZones[*].{ZoneName:ZoneName, ZoneId:ZoneId}"
> ```

- The Exadata system model, either X9M or X11M

To learn how to create Oracle Exadata infrastructure, see [Step 2: Create an Oracle Exadata infrastructure in Oracle Database@AWS](#).

ODB network

An *ODB network* is a private isolated network that hosts OCI infrastructure in an AWS Availability Zone (AZ). The ODB network consists of a CIDR range of IP addresses. The ODB network maps directly to the network that exists within the OCI child site, thus serving as the means of communication between AWS and OCI. You specify an ODB network when you create your Exadata VM clusters.

The ODB network is similar to a VPC but differs in the following ways:

- Lacks connectivity to on-premise networks or the internet
- Created using Oracle Database@AWS tools rather than Amazon EC2 APIs
- Supports creation of only Oracle Database@AWS resources
- Managed by AWS rather than customers

When you create an ODB network, you specify information such as the following:

- Availability Zone — The ODB network is specific to an AZ.

  You can only deploy Oracle Database@AWS in the AWS Region US East (N. Virginia). The only supported physical AZs within this Region have the physical IDs use1-az4 and use1-az6. To find the logical zone names in your account that map to the physical zone IDs, run the following command.
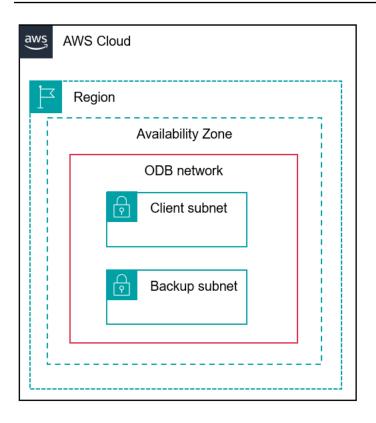
  ```
  aws ec2 describe-availability-zones --region us-east-1 --query "AvailabilityZones[*].
  {ZoneName:ZoneName, ZoneId:ZoneId}"
  ```

- Client CIDR addresses — The ODB network requires a client subnet CIDR for Exadata VM clusters.

- Backup CIDR addresses — The ODB network requires a backup subnet CIDR for managed database backups.

For more information, see Step 1: Create an ODB network in Oracle Database@AWS.

Virtual Private Cloud (VPC)

A *Virtual Private Cloud (VPC)* is a virtual network that you create in the AWS cloud. It is logically isolated from other virtual networks in the AWS cloud, providing you with complete control over the virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables and network gateways. For more information, see What is Amazon VPC?

You can launch Amazon EC2 instances into your Amazon VPC. The EC2 instances can host application servers that communicate with Oracle Exadata databases. You can manage and launch the application servers just like any other EC2 instances in your VPC. For more information, see What is Amazon EC2?
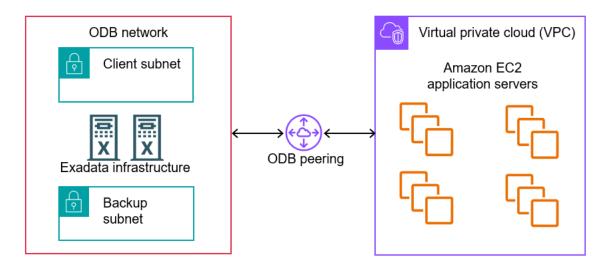
By default, the ODB network doesn't have connectivity to VPCs. To connect the ODB network to your existing AWS infrastructure, create a peering connection between the ODB network and one VPC. You can specify the VPC when you create the ODB network. For more information, see Step 1: Create an ODB network in Oracle Database@AWS.

ODB peering

ODB peering between an ODB network and a VPC enables traffic to be routed privately between one VPC and one ODB network. Because the ODB network is private, by default it does not have connectivity to VPCs, on-premise networks, or the internet. To connect to Exadata databases in the ODB network, set up an ODB peering connection.
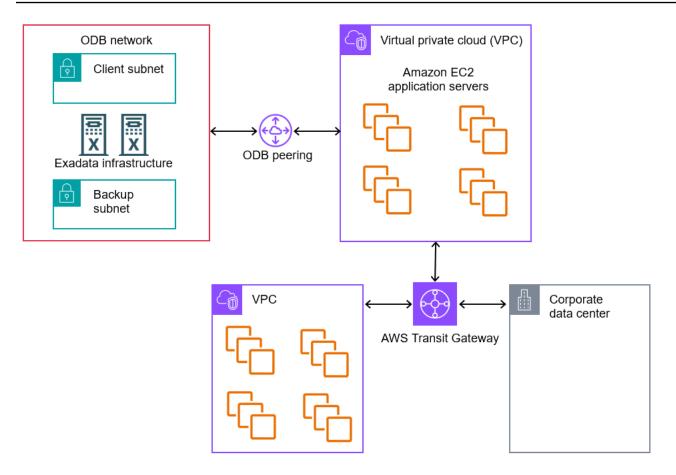
> ⓘ **Note**
>
> ODB peering is different from VPC peering, which is a peering connection between two VPCs that enables you to route traffic between them.

To set up ODB peering between an ODB network and one VPC, specify the VPC when you create or update the ODB network. Then update your VPC route tables using the Amazon EC2 command `create-route`. The ODB network route tables are automatically updated with the VPC CIDR addresses. After peering, an Amazon EC2 instance within the VPC can communicate with an Oracle Exadata database in the ODB network as if they were within the same network. For more information, see Step 1: Create an ODB network in Oracle Database@AWS and Updating an ODB network in Oracle Database@AWS.

Amazon VPC Transit Gateways is a network transit hub used to interconnect VPCs and on-premises networks. You can't create a one-to-many peering connection between an ODB network and multiple VPCs. But you can peer your ODB network to a VPC, and then attach this VPC to a transit gateway. The gateway can connect to multiple VPCs. With this configuration, you can route traffic between multiple VPC subnets to your ODB network.

Exadata VM clusters

An *Exadata VM cluster* is a set of tightly coupled Exadata VMs. Each VM has a complete Oracle database installation that includes all features of Oracle Enterprise Edition, including Oracle Real Application Clusters (Oracle RAC) and Oracle Grid Infrastructure. You can create one or more Oracle Exadata databases on a VM cluster. For diagrams that show the architecture of VMs and VM clusters, see Exadata Database Service on Dedicated Infrastructure Technical Architecture.

When you create a *VM cluster*, you specify information that includes the following:

- An ODB network

- An Oracle Exadata infrastructure

- The database servers on which to place the VMs in the cluster

- The total amount of usable Exadata storage

You can configure the CPU cores, memory, and local storage for each VM in a VM cluster. For more information, see Step 3: Create an Exadata VM cluster or Autonomous VM cluster in Oracle Database@AWS.

Autonomous VM clusters

Autonomous VM clusters are fully managed databases that automate key management tasks using machine learning and AI. Unlike traditional databases, autonomous databases automatically provision, secure, update, backup, and tune the database with no human intervention required.

You can configure the ECPU core count per VM, database memory per CPU, database storage, and maximum number of autonomous container database. For more information, see Step 3: Create an Exadata VM cluster or Autonomous VM cluster in Oracle Database@AWS.

Oracle Exadata databases

Oracle Exadata is an engineered system that provide a high-performance platform for running Oracle databases. With Oracle Database@AWS, you use the AWS console to create the Oracle Exadata infrastructure and VM clusters that host the Exadata databases. You then use OCI APIs to create and manage the Oracle databases. For more information, see Step 4: Create Oracle Exadata databases in Oracle Cloud Infrastructure.

Onboarding to Oracle Database@AWS

Before you can begin using Oracle Database@AWS, make sure you're signed up for AWS and create necessary users. Then you can purchase Oracle Database@AWS from AWS Marketplace by accepting a private offer from Oracle.

Sign up for an AWS account

If you do not have an AWS account, complete the following steps to create one.

To sign up for an AWS account

1. Open https://portal.aws.amazon.com/billing/signup.

2. Follow the online instructions.

   Part of the sign-up procedure involves receiving a phone call and entering a verification code on the phone keypad.

   When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign administrative access to a user, and use only the root user to perform tasks that require root user access.

AWS sends you a confirmation email after the sign-up process is complete. At any time, you can view your current account activity and manage your account by going to https://aws.amazon.com/ and choosing **My Account**.

Create a user with administrative access

After you sign up for an AWS account, secure your AWS account root user, enable AWS IAM Identity Center, and create an administrative user so that you don't use the root user for everyday tasks.

Secure your AWS account root user

1. Sign in to the AWS Management Console as the account owner by choosing **Root user** and entering your AWS account email address. On the next page, enter your password.

For help signing in by using root user, see [Signing in as the root user](#) in the *AWS Sign-In User Guide*.

2. Turn on multi-factor authentication (MFA) for your root user.

   For instructions, see [Enable a virtual MFA device for your AWS account root user (console)](#) in the *IAM User Guide*.

Create a user with administrative access

1. Enable IAM Identity Center.

   For instructions, see [Enabling AWS IAM Identity Center](#) in the *AWS IAM Identity Center User Guide*.

2. In IAM Identity Center, grant administrative access to a user.

   For a tutorial about using the IAM Identity Center directory as your identity source, see [Configure user access with the default IAM Identity Center directory](#) in the *AWS IAM Identity Center User Guide*.

Sign in as the user with administrative access

- To sign in with your IAM Identity Center user, use the sign-in URL that was sent to your email address when you created the IAM Identity Center user.

  For help signing in using an IAM Identity Center user, see [Signing in to the AWS access portal](#) in the *AWS Sign-In User Guide*.

Assign access to additional users

1. In IAM Identity Center, create a permission set that follows the best practice of applying least-privilege permissions.

   For instructions, see [Create a permission set](#) in the *AWS IAM Identity Center User Guide*.

2. Assign users to a group, and then assign single sign-on access to the group.

   For instructions, see [Add groups](#) in the *AWS IAM Identity Center User Guide*.

Request a private offer for Oracle Database@AWS

The AWS Marketplace seller private offer feature enables you to request and receive Oracle Database@AWS pricing and EULA terms from Oracle. You negotiate pricing and terms with Oracle, and then Oracle creates a private offer for the AWS account that you designate. You accept the private offer and receive the negotiated price and terms of use. When the private offer agreement reaches its expiration date, you're either moved automatically to the product's public pricing or unsubscribed from Oracle Database@AWS. For more information about private offers, see Private offers in AWS Marketplace.

To request and accept a private offer for Oracle Database@AWS

1. Sign in to the AWS Management Console.

2. Search for and then choose Oracle Database@AWS.

3. Choose **Request private offer**.

> ⓘ **Note**
>
> The dashboard isn't available until after you have accepted a private offer.

4. On the OCI site, specify details such as the region and your contact information.

5. Wait for an OCI representative to contact you and make a private offer available.

6. In the AWS Management Console, choose **View private offer**.

7. Choose the offer and then choose **View offer**.

8. Choose **Create contract** and respond to the subsequent prompts to accept the private offer.

9. After accepting the private offer, you'll need to activate your Oracle Cloud Infrastructure (OCI) account. You can access the Oracle activation links directly from AWS Management Console.

   1. In the console, navigate to the **Get started** section.

   2. Click on the Oracle activation link provided in the console. Alternatively, you can also use the activation link sent to you via email.

   3. On the Oracle activation page, choose whether to create a new Oracle cloud account or add to an existing account.

   4. Complete the activation process by following the on-screen instructions.

5. After submitting your activation request, you'll see an **Activation in progress** status in the AWS Management Console, and the dashboard will be temporarily disabled with a reason displayed.

6. Once activation is complete, the dashboard will become available, allowing you to manage your resources.

10. In the AWS Management Console, choose **Dashboard**.

Getting started with Oracle Database@AWS

To begin using Oracle Database@AWS, you create an ODB network, Oracle Exadata infrastructure, and VM cluster. You can then create Oracle Exadata databases using the Oracle Cloud Infrastructure console or APIs.

Prerequisites for setting up Oracle Database@AWS

Before configuring your Oracle Exadata infrastructure, make sure that you perform the steps in Onboarding to Oracle Database@AWS. You must have accepted a private offer to use Oracle Database@AWS.

Limitations for Oracle Database@AWS

- You can only deploy Oracle Database@AWS in the AWS Region US East (N. Virginia). The only supported physical AZs within this Region have the physical IDs `use1-az4` and `use1-az6`. To find the logical zone names in your account that map to the physical zone IDs, run the following command.

  ```
  aws ec2 describe-availability-zones --region us-east-1 --query "AvailabilityZones[*].
  {ZoneName:ZoneName, ZoneId:ZoneId}"
  ```

- You can deploy a VM cluster only into the AZ where you created your ODB network and Oracle Exadata infrastructure.
- The subnet in the VPC or applications connecting to Oracle Database@AWS must be in AZ `use1-az6` or `use1-az4`.
- Your AWS account has the following limitations for Oracle Database@AWS:
  - By default, you can create only 5 ODB networks in an AWS account.
  - By default, you can create only 10 ODB peering connections in an AWS account.
  - You can't share Oracle Exadata infrastructure across AWS accounts.
  - VM clusters must be in the same AWS account as the Oracle Exadata infrastructure.
- You can deploy only VM clusters in your ODB network. No other resources are permitted.
- You can't create an ODB peering connection between a VPC and multiple ODB networks or between an ODB network and multiple VPCs in the same AWS account. There is a 1:1 relationship between a VPC and an ODB network.

- IPv6 isn't supported for Oracle Exadata X9M.
- You can't change the storage allocation after you create a VM cluster.

Planning IP address space in Oracle Database@AWS

Plan carefully for IP address space in Oracle Database@AWS. Consider the IP address consumption based on the number of VM clusters, including the number of VMs per cluster that you can provision into the ODB network.

Topics

- [Restrictions for IP addresses in the ODB network](#)
- [Client subnet CIDR requirements for the ODB network](#)
- [Backup subnet CIDR requirements for the ODB network](#)
- [IP consumption scenarios for the ODB network](#)

Restrictions for IP addresses in the ODB network

Note the following restrictions regarding CIDR ranges in the ODB network:

- You can't modify the client or backup subnet CIDR range for the ODB network after you create it.
- You can't use the VPC CIDR ranges in the **Restricted associations** column in the table in [IPv4 CIDR block association restrictions](#).
- For Exadata X9M, IP addresses 100.106.0.0/16 and 100.107.0.0/16 are reserved for the cluster interconnect by OCI automation, so you can't do the following:
  - Assign these ranges to the client or backup CIDR range of the ODB network.
  - Use these ranges for a VPC CIDR that is used to connect to the ODB network.
- The following CIDR ranges are reserved for Oracle Cloud Infrastructure and can't be used for the ODB network:
  - Oracle Cloud reserved range CIDR 169.254.0.0/16
  - Reserved Class D 224.0.0.0 — 239.255.255.255
  - Reserved Class E 240.0.0.0 — 255.255.255.255
- You can't overlap the IP address CIDR ranges for the client and backup subnets.
- You can't overlap the IP address CIDR ranges allocated for the client and backup subnets with the VPC CIDR ranges used to connect to the ODB network.

- You can't provision VMs in a VM cluster into different ODB networks. The network is a property of the VM cluster, which means you can only provision the VMs in the VM cluster into the same ODB network.

Client subnet CIDR requirements for the ODB network

In the following table, you can find the number of IP addresses consumed by the service and infrastructure for the client subnet CIDR.

Number of IP addresses	Consumed by	Notes
5	Oracle Database@AWS	These IP addresses are reserved regardless of how many VM clusters you provision in the ODB network. Oracle Database@AWS consumes the following: • 2 IP addresses at the beginning of the CIDR range • 1 IP address at the end of the CIDR range • 2 IP addresses from anywhere in the CIDR range
3	Each VM cluster	These IP addresses are reserved for Single Client Access Names (SCANs) regardless of how many VMs are present in each VM cluster.
4	Each VM	These IP addresses depend solely on the number of VMs in the infrastructure.

Backup subnet CIDR requirements for the ODB network

In the following table, you can find the number of IP addresses consumed by the service and infrastructure for the backup subnet CIDR.

Number of IP addresses	Consumed by	Notes
3	Oracle Database@AWS	These IP addresses are reserved regardless of how many VM clusters you provision in the ODB network. Oracle Database@AWS consumes the following: • 2 IP addresses at the beginning of the CIDR range • 1 IP address at the end of the CIDR range
3	Each VM	These IP addresses depend solely on the number of VMs in the infrastructure.

IP consumption scenarios for the ODB network

In the following table, you can see the IP addresses consumed in the ODB network for different configurations of VM clusters. Whereas /28 is the minimum CIDR range for the client subnet CIDR to deploy 1 VM cluster with 2 VMs, we recommend that you use at least a /27 CIDR range. In this case, the IP range isn't fully consumed by the VM clusters and permits allocation of additional IP addresses.

Configuration	Client IPs consumed	Client IPs minimum	Backup IPs consumed	Backup IPs minimum
1 VM cluster with 2 VMs	16 (5 service + 3 cluster + 4*2)	16 (/28 CIDR range)	9 (3 service + 3*2)	16 (/28 CIDR range)
1 VM cluster with 3 VMs	20 (5 service + 3 cluster + 4*3)	32 (/27 CIDR range)	12 (3 service + 3*3)	16 (/28 CIDR range)
1 VM cluster with 4 VMs	24 (5 service + 3 cluster + 4*4)	32 (/27 CIDR range)	15 (3 service + 3*4)	16 (/28 CIDR range)
1 VM cluster with 8 VMs	40 (5 service + 3 cluster + 4*8)	64 (/26 CIDR range)	27 (3 service + 3*8)	32 (/27 CIDR range)

The following table shows how many instances of each configuration are possible given a specific client CIDR range. For example, 1 VM cluster with 4 VMs consumes 24 IP addresses in the client subnet. If the CIDR range is /25, 128 IP addresses are available. Thus, you can provision 4 VM clusters in the subnet.

VM cluster configuration	Number with /27 (32 IPs)	Number with /26 (64 IPs)	Number with /25 (128 IPs)	Number with /24 (256 IPs)	Number when /23 (512 IPs)	Number when /22 (1024 IPs)
1 VM cluster with 2 VMs (16 IPs)	2	5	11	22	46	92
1 VM cluster with 3 VMs (20 IPs)	1	3	8	16	33	67
1 VM cluster with 4 VMs (24 IPs)	1	3	6	13	26	53
2 VM clusters with 2 VMs each (27 IPs)	1	2	5	11	23	46
2 VM clusters with 3 VMs each (35 IPs)	0	1	4	8	16	33
2 VM clusters with 4 VMs each (43 IPs)	0	1	3	6	13	26

Step 1: Create an ODB network in Oracle Database@AWS

An ODB network is a private isolated network that hosts OCI infrastructure in an Availability Zone (AZ). An ODB network and an Oracle Exadata infrastructure are preconditions for provisioning VM clusters and creating Exadata databases. You can create the ODB network and Oracle Exadata infrastructure in either order. For more information, see ODB network and ODB peering.

This task assumes that you have read Planning IP address space in Oracle Database@AWS. To modify or delete the ODB network later, see Managing Oracle Database@AWS.

To create an ODB network

1. Sign in to the AWS Management Console and open the Oracle Database@AWS console at https://console.aws.amazon.com/odb/.

2. From the left pane, choose **ODB networks**.

3. Choose **Create ODB network**.

4. For **ODB network name**, enter a network name. The name must be 1–255 characters and begin with an alphabetic character or underscore. It can't contain consecutive hyphens.

5. For **Availability Zone**, choose the AZ name associated with physical ID use1-az4 or use1-az6.

6. For **Client subnet CIDR**, specify a CIDR range for the client connections. For more information, see Client subnet CIDR requirements for the ODB network.

7. For **Backup subnet CIDR**, specify a CIDR range for the backup connections. To isolate the backup traffic and improve resiliency, we recommend that you don't overlap the backup CIDR and the client CIDR. For more information, see Backup subnet CIDR requirements for the ODB network.

8. (Optional) For **VPC ID**, specify the ID of a VPC to use for ODB peering to the ODB network. When you create an ODB peering connection, application servers running in the specified VPC can access Exadata databases that you create on VM clusters.

   > ⚠️ **Important**
   >
   > After you create your ODB network, update your VPC route tables with the destination CIDR in the ODB network. For more information, see Configuring VPC route tables for ODB peering.

9. (Optional) For **Domain name prefix**, enter a name to use as a prefix to your domain. The domain name is fixed as **oraclevcn.com**. For example, if you enter **myhost**, the fully qualified domain name is **myhost.oraclevcn.com.**

10. (Optional) For **Tags**, enter up to 50 tags for the network. A tag is a key-value pair that you can use to organize and track your resources.

11. Choose **Create ODB network**.

Step 2: Create an Oracle Exadata infrastructure in Oracle Database@AWS

The Oracle Exadata infrastructure is the underlying architecture of database servers, storage servers, and networking that run Oracle Exadata databases. Choose either Exadata X9M or X11M as the system model. You can then create VM clusters on Exadata infrastructure using the AWS console.

You can create the Oracle Exadata infrastructure and the ODB network in either order. You don't need to specify networking information when you create the infrastructure.

You can't modify an Oracle Exadata infrastructure after you create it. To delete an Exadata infrastructure, see [Deleting an Oracle Exadata infrastructure in Oracle Database@AWS](#).

To create an Exadata infrastructure

1. Sign in to the AWS Management Console and open the Oracle Database@AWS console at https://console.aws.amazon.com/odb/.
2. From the left pane, choose **Exadata infrastructures**.
3. Choose **Create Exadata infrastructure**.
4. For **Exadata infrastructure name**, enter a name. The name must be 1–255 characters and begin with an alphabetic character or underscore. It can't contain consecutive hyphens.
5. For **Availability Zone**, leave the default. Then choose **Next**.
6. For **Exadata system model**, choose either **Exadata.X9M** or **Exadata.X11M**. For **Exadata.X11M**, also choose the following server types:

   - For **Database server type**, choose the database server model type of your Exadata infrastructure. Currently, the only choice is **X11M**.
   - For **Storage server type**, choose the storage server model type of your Exadata infrastructure. Currently, the only choice is **X11M-HC**.

7. For **Database servers**, leave the default of 2 or move the slider to choose up to 32 servers. To specify more than 2, request a limit increase from OCI.

   Each database server supports 126 compute units, which are measured in ECPUs for Exadata X11M and OCPUs for Exadata X9M. The total compute count changes as you change the number of servers. For more information about OCPUs and ECPUs, see [Compute Models in Autonomous Database](#) in the Oracle documentation.

8. For **Storage servers**, leave the default of 3 or move the slider to choose up to 64 servers. To specify more than 3, request a limit increase from OCI. Each storage server provides 64 TB. The total TB of storage changes as you change the number of servers. Then choose **Next**.

9. For **Maintenance window**, configure when system maintenance can occur:

   a. For **Scheduling preference**, select one of the following options:

      - **Oracle-managed schedule** - Oracle determines the optimal time for maintenance activities.

      - **Customer-managed schedule** - You specify when maintenance activities can occur.

   b. For **Patching mode**, select one of the following options:

      - **Rolling** - Updates are applied to one node at a time, allowing the database to remain available during patching.

      - **Non-rolling** - Updates are applied to all nodes simultaneously, which may require downtime.

   c. If you selected **Customer-managed schedule**, configure the following additional settings:

      - For **Maintenance months**, select the months when maintenance can be performed.

      - For **Week of the month**, select which week of the month maintenance can be performed (First, Second, Third, Fourth, or Last).

      - For **Day of week**, select the day when maintenance can be performed (Monday through Sunday).

      - For **Start hour**, select the hour when the maintenance window begins. The time is in UTC.

      - For **Notification lead time**, select how many days in advance you want to be notified about upcoming maintenance.

   > ⓘ **Note**
   >
   > Oracle Cloud Infrastructure performs system maintenance during this window. During maintenance, your Exadata infrastructure remains available, but you might experience brief periods of higher latency.

10. (Optional) For **OCI maintenance notification contacts**, enter up to 10 email addresses. AWS forwards these email addresses to OCI. When updates occur, OCI mails notifications to the listed addresses.

11. (Optional) For **Tags**, enter up to 50 tags for the infrastructure. A tag is a key-value pair that you can use to organize and track your resources.

12. Choose **Next** and review your infrastructure settings.

13. Choose **Create Exadata infrastructure**.

Step 3: Create an Exadata VM cluster or Autonomous VM cluster in Oracle Database@AWS

An Exadata VM cluster is a set of VMs on which you can install Oracle Exadata databases. You create VM clusters on Exadata infrastructure. You can deploy multiple VM clusters with different Oracle Exadata infrastructures in the same ODB network.

You can also create an Autonomous database to deploy on your Exadata infrastructure.

> ⚠️ **Important**
>
> The creation process can take over 6 hours, depending on the size of the VM cluster.

Exadata VM cluster

To create an Exadata VM cluster

1. Sign in to the AWS Management Console and open the Oracle Database@AWS console at https://console.aws.amazon.com/odb/.

2. From the left pane, choose **Exadata VM clusters**.

3. Choose **Create VM cluster**.

4. For **VM cluster name**, enter a name. The name must be 1–255 characters and begin with an alphabetic character or underscore. It can't contain consecutive hyphens.

5. (Optional) For **Grid Infrastructure cluster name**, enter a name. The name must be 1–11 characters and can't contain hyphens.

6. For **Time zone**, enter a time zone.

7. For **License options**, choose **Bring Your Own License (BYOL)** or **License Included**, and then choose **Next**. This license is the OCI license provided by Oracle, not a license provided by AWS.

8. Configure Exadata infrastructure settings as follows:

   a. For **Infrastructure**, choose the following:

      - For **Exadata infrastructure name**, choose the infrastructure to use for this VM cluster.

      - For **Grid Infrastructure version**, choose the version to use for this VM cluster.

      - For **Exadata image version**, choose the version to use for this VM cluster. We recommend that you choose the version shown, which is the highest version available.

   b. For **Database servers**, select one or more database servers to host your VM cluster.

   c. For **Configuration**, do the following:

      - Choose the **CPU core count**, **Memory**, and **Local storage** for each VM, or accept the defaults.

      - Choose the total amount of **Exadata storage** for the VM cluster, or accept the default.

   d. (Optional) For **Storage allocation**, select any of the following options:

      - **Enable storage allocation for Exadata sparse snapshots**

      - **Enable storage allocation for local backups**

      You can't change this storage allocation later. Review your selection, and then choose **Next**.

9. Configure connectivity as follows:

   a. For **ODB network**, choose an existing ODB network.

   b. For **Host name prefix**, enter a prefix for the VM cluster. Make sure not to include the domain name. The prefix forms the first portion of the Oracle Exadata VM cluster host name.

> ⓘ **Note**
>
> The **Host domain name** is fixed as **oraclevcn.com**.

    c.   For **SCAN listener port (TCP/IP)**, enter a port number that for TCP access to the single client access name (SCAN) listener. The default port is **1521**. Or you can enter a custom SCAN port in the range **1024–8999**, excluding the following port numbers: **2484**, **6100**, **6200**, **7060**, **7070**, **7085**, and **7879**. Then choose **Next**.

    d.   For **SSH key pairs**, enter the public key portion of one or more key pairs used for SSH access to the VM cluster. Then choose **Next**.

10. (Optional) Choose diagnostics and tags as follows:

    a.   Choose whether to enable diagnostic collection for **Diagnostic events**, **Health monitor**, and **Incident logs and trace collections**. Oracle can use this diagnostic information to identify, track, and resolve issues.

    b.   For **Tags**, enter up to 50 tags for the VM cluster. A tag is a key-value pair that you can use to organize and track your resources. Then choose **Next**.

11. Review your settings. Then choose **Create VM cluster**.

Autonomous VM cluster

To create an Autonomous VM cluster

1. Sign in to the AWS Management Console and open the Oracle Database@AWS console at https://console.aws.amazon.com/odb/.

2. From the left pane, choose **Autonomous VM clusters**.

3. Choose **Create Autonomous VM cluster**.

4. For **VM cluster name**, enter a name. The name must be 1–255 characters and begin with an alphabetic character or underscore. It can't contain consecutive hyphens.

5. For **Time zone**, enter a time zone.

6. For **License options**, choose **Bring Your Own License (BYOL)** or **License Included**, and then choose **Next**. This license is the OCI license provided by Oracle, not a license provided by AWS.

7. Configure Exadata infrastructure settings as follows:

a. For **Exadata infrastructure name**, choose the infrastructure to use for this Autonomous VM cluster.

b. For **Database servers**, select one or more database servers to host your Autonomous VM cluster.

c. For **Configuration**, do the following:

   - Choose the **ECPU core count per VM**, **Database memory per CPU**, **Database storage**, and **Maximum number of Autonomous Container Database** or accept the defaults.

   - Choose the total amount of **Exadata storage** for the Autonomous VM cluster, or accept the default.

8. Configure connectivity as follows:

   a. For **ODB network**, choose an existing ODB network.

   b. For **SCAN listener port (TCP/IP)**, enter a port number for Port (non-TLS). The default port is **1521**. Or you can enter a Port(TLS) in the range **1024–8999**, excluding the following port numbers: **2484**, **6100**, **6200**, **7060**, **7070**, **7085**, and **7879**. Then choose **Next**.

      Select **Enable mutual TLS (mTLS) authentication** to allow mutual TLS authentication.

9. (Optional) Choose diagnostics and tags as follows:

   a. Choose whether to schedule modification configuration to **Oracle-managed schedule** or **Customer-managed schedule**. If you choose **Customer-managed schedule**, set the **Maintenance months**, **Weeks of the month**, **Day of the week**, and **Start hour (UTC)**.

   b. For **Tags**, enter up to 50 tags for the Autonomous VM cluster. A tag is a key-value pair that you can use to organize and track your resources. Then choose **Next**.

10. Review your settings. Then choose **Create Autonomous VM cluster**.

Step 4: Create Oracle Exadata databases in Oracle Cloud Infrastructure

You create and manage ODB networks, Oracle Exadata infrastructure, and Exadata VM clusters in Oracle Database@AWS. You create and manage Oracle Exadata databases in Oracle Cloud Infrastructure.

To create Oracle Exadata databases

1. Sign in to the AWS Management Console and open the Oracle Database@AWS console at
    https://console.aws.amazon.com/odb/.

2. From the left pane, choose **Exadata VM clusters**.

3. Choose an Exadata VM cluster to see the details page.

4. Choose **Manage in OCI** to be redirected to the Oracle Cloud Infrastructure console.

> ⓘ **Note**
>
> In this preview, the **Manage in OCI** button is supported only in the Exadata VM cluster
> details page, not in the Exadata VM clusters list page.

5. Create your Exadata databases in OCI.

Configuring the network in Oracle Database@AWS

If you specified a VPC for ODB peering to your ODB network, make sure to update your VPC route tables and configure DNS resolution. For more information about ODB peering, see ODB peering.

Configuring VPC route tables for ODB peering

A *route table* contains a set of rules, called *routes*, that determine where network traffic from your subnet or gateway is directed. The destination CIDR in a route table is a range of IP addresses where you want traffic to go. If you specified a VPC for ODB peering to your ODB network, update your VPC route table with the destination IP range in your ODB network. For more information about ODB peering, see ODB peering.

To update a route table, use the AWS CLI command `ec2 create-route` as follows:

```
aws ec2 create-route
   --route-table-id route-table-id
   --destination-cidr-block cidr-block
   --odb-network-arn odb-network-arn
```

The ODB network route tables are automatically updated with the VPC CIDRs. To allow access to the ODB network for only specific subnet CIDRs rather than all CIDRs in the VPC, update the ODB network by adding or removing peered CIDR ranges. For more information, see Updating an ODB network in Oracle Database@AWS.

For more information about VPC route tables, see Subnet route tables in the *Amazon Virtual Private Cloud User Guide* and ec2 create-route in the *AWS CLI Command Reference*.

Configuring DNS for Oracle Database@AWS

Amazon Route 53 is a highly available and scalable Domain Name System (DNS) web service that you can use for DNS routing. When you create an ODB peering connection between your ODB network and a VPC, you need a mechanism to resolve DNS queries for ODB network resources from within the VPC. You can use Amazon Route 53 to configure the following resources:

- An outbound endpoint

The endpoint is required to send DNS queries to the ODB network.

- A resolver rule

  This rule specifies the domain name of the DNS queries that the Route 53 Resolver forwards to the DNS for the ODB network.

How DNS works in Oracle Database@AWS

Oracle Database@AWS manages Domain Name System (DNS) configuration for the ODB network automatically. The domain name for the ODB network is fixed as `oraclevcn.com`. You can specify a custom domain name prefix when you create the ODB network. For more information, see Step 1: Create an ODB network in Oracle Database@AWS.

When Oracle Database@AWS provisions an ODB network, it creates the following resources:

- An Oracle Cloud Infrastructure (OCI) virtual cloud network (VCN) with the same CIDR blocks as the ODB network

  This VCN resides in the customer's linked OCI tenancy. There is a 1:1 mapping between an ODB network and an OCI VCN. Every ODB network is associated with an OCI VCN.

- A private DNS resolver within the OCI VCN

  This DNS resolver handles DNS queries within the OCI VCN. OCI automation creates records for the VM cluster. Scans use the `*.oraclevcn.com` fully qualified domain name (FQDN).

- A DNS listening endpoint within the OCI VCN for the private DNS resolver

  You can find the DNS listening endpoint in the ODB network details page on the Oracle Database@AWS console.

Configuring an outbound endpoint in an ODB network in Oracle Database@AWS

An outbound endpoint allows DNS queries to be sent from your VPC to a network or IP address. The endpoint specifies the IP addresses from which queries originate. To forward DNS queries from your VPC to your ODB network, create an outbound endpoint using the Route 53 console. For more information, see Forwarding outbound DNS queries to your network.

To configure an outbound endpoint in an ODB network

1. Sign in to the AWS Management Console and open the Route 53 console at https://console.aws.amazon.com/route53/.

2. From the left pane, choose **Outbound endpoints**.

3. On the navigation bar, choose the **Region** for the VPC where you want to create the outbound endpoint.

4. Choose **Create outbound endpoint**.

5. Complete the **General settings for outbound endpoint** section as follows:

   a. Choose a **Security group** that allows outbound TCP and UDP connectivity to the following:

      - IP addresses that the resolvers use for DNS queries on your ODB network

      - Ports that the resolvers use for DNS queries on your ODB network

   b. For **Endpoint Type**, choose **IPv4**.

   c. For **Protocols for this endpoint**, choose **Do53**.

6. In **IP addresses**, provide the following information:

   - Either specify IP addresses or let the Route 53 Resolver choose IP addresses for you from the available addresses in the subnet. Choose a minimum of 2 up to a maximum of 6 IP addresses for DNS queries. We recommend that you choose IP addresses in at least two different Availability Zones.

   - For **Subnet**, choose subnets that have the following:

      - Route tables that include routes to the IP addresses of the DNS listener on ODB network

      - Network access control lists (ACLs) that allow UDP and TCP traffic to the IP addresses and the ports that the resolvers use for DNS queries on ODB network

      - Network ACLs that allow traffic from resolvers on destination port range 1024-65535

7. (Optional) For **Tags**, specify tags for the endpoint.

8. Choose **Submit**.

Configuring a resolver rule in Oracle Database@AWS

A resolver rule is a set of criteria that determines how to route DNS queries. Either reuse or create a rule that specifies the domain name of the DNS queries that the resolver forwards to the DNS for the ODB network.

Using an existing resolver rule

To use an existing resolver rule, your action depends on the type of rule:

A rule for the same domain in the same AWS Region as the VPC in your AWS account

Associate the rule with your VPC instead of creating a new rule. Choose the rule from the rule dashboard and associate it with the applicable VPCs in the AWS Region.

A rule for the same domain in the same Region as your VPC but in a different account

Use AWS Resource Access Manager to share the rule from the remote account to your account. When you share a rule, you also share the corresponding outbound endpoint. After you share the rule with your account, choose the rule from the rule dashboard and associate it with the VPCs in your account. For more information, see Managing forwarding rules.

Creating a new resolver rule

If you can't reuse an existing resolver rule, create a new rule using the Amazon Route 53 console.

To create a new resolver rule

1. Sign in to the AWS Management Console and open the Route 53 console at https://console.aws.amazon.com/route53/.
2. From the left pane, choose **Rules**.
3. On the navigation bar, choose the **Region** for the VPC where the outbound endpoint exists.
4. Choose **Create rule**.
5. Complete the **Rule for outbound traffic** sections as follows:

   a. For **Rule type**, choose **Forward rule**.

   b. For **Domain name**, specify the full domain name from ODB network.

   c. For **VPCs that use this rule**, associate it with the VPC from where DNS queries are forwarded to your ODB network.

    d.   For **Outbound endpoint**, choose the outbound endpoint that you created in [Configuring an outbound endpoint in an ODB network in Oracle Database@AWS](#).

> ⓘ **Note**
>
> The VPC associated with this rule doesn't need to be the same VPC where you created the outbound endpoint.

6.   Complete the **Target IP addresses** section as follows:

    a.   For **IP address**, specify the IP address of the DNS listener IP on your ODB network.

    b.   For **Port**, specify **53**. This is the port that the resolver use for DNS queries.

> ⓘ **Note**
>
> The Route 53 Resolver forwards DNS queries that match this rule and originate from a VPC associated with this rule to the referenced outbound endpoint. These queries are forwarded to the target IP addresses that you specify in the **Target IP addresses**.

    c.   For **Transmission protocol**, choose **Do53**.

7.   (Optional) For **Tags**, specify tags for the rule.

8.   Choose **Submit**.

Testing your DNS configuration in Oracle Database@AWS

After you have creating your outbound endpoint and resolver rule, test to make sure that the DNS resolves correctly. Using an Amazon EC2 instance in your application VPC, perform a DNS resolution as follows:

For Linux or MacOS

Use a command of the form `dig` *`record-name record-type`*.

For Windows
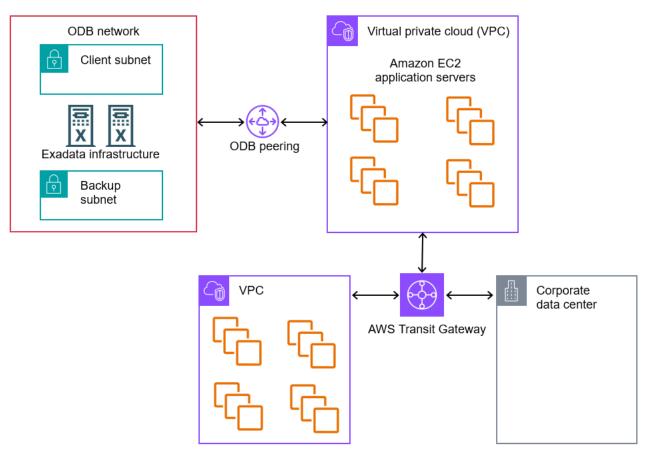
Use a command of the form `nslookup -type=`*`record-name record-type`*.

Configuring Amazon VPC Transit Gateways for Oracle Database@AWS

Amazon VPC Transit Gateways is a network transit hub that interconnects virtual private clouds (VPCs) and on-premises networks. Each VPC in the hub-and-spoke architecture can connect to the transit gateway to gain access to other connected VPCs. AWS Transit Gateway supports traffic for both IPv4 and IPv6.

In Oracle Database@AWS, an ODB network supports a peering connection to only one VPC. If you connect a transit gateway to a VPC that is peered to an ODB network, you can connect multiple VPCs to this gateway. Applications running in these VPCs can access an Exadata VM cluster running in your ODB network.

The following diagram shows a transit gateway that is connected to two VPCs and one on-premises network.

In the preceding diagram, one VPC is peered to an ODB network. In this configuration, the ODB network can route traffic to all VPCs attached to the transit gateway. The route table for each VPC

includes both the local route and routes that send traffic destined for the ODB network to the transit gateway.

Note the following limitations of Amazon VPC Transit Gateways for Oracle Database@AWS:

- Amazon VPC Transit Gateways doesn't offer native integration to use an ODB network as an attachment. Therefore, VPC features such as the following aren't available:
  - Resolution of public DNS hostnames to private IP addresses
  - Event notification for changes in the ODB network topology, routing, and connection status
- Multicast traffic to the ODB network isn't supported.

In AWS Transit Gateway, you're charged for the number of connections that you make to the transit gateway per hour and the amount of traffic that flows through AWS Transit Gateway. For cost information, see AWS Transit Gateway pricing.

To configure a transit gateway for Oracle Database@AWS

1. Add CIDR ranges to your ODB network for the VPCs and on-premises networks that you plan to attach to your transit gateway. For more information, see Updating an ODB network in Oracle Database@AWS.
2. Follow the steps in Get started with using Amazon VPC Transit Gateways.

Managing Oracle Database@AWS

You can modify and delete some Oracle Database@AWS resources after you create them.

Updating an ODB network in Oracle Database@AWS

You can update the following ODB network resources:

- The ODB network name
- The Amazon VPC to use for establishing an ODB peering connection to the ODB network
- The VPC CIDR ranges that can access Exadata resources in the ODB network

> ⓘ **Note**
>
> By specifying CIDR ranges, you limit connectivity to the necessary VPC subnets instead of making the entire VPC available to the ODB network.

This section assumes that you have already created an ODB network in Step 1: Create an ODB network in Oracle Database@AWS.

To update an ODB network

1. Sign in to the AWS Management Console and open the Oracle Database@AWS console at https://console.aws.amazon.com/odb/.
2. From the left pane, choose **ODB networks**.
3. Select the network that you want to modify.
4. Choose **Modify**.
5. (Optional) For **ODB network name**, enter a new network name. The name must be 1–255 characters and begin with an alphabetic character or underscore. It can't contain consecutive hyphens.
6. (Optional) For **Peered CIDRs**, specify CIDR ranges from the peered VPC that need connectivity to the ODB network. To limit access, we recommend that you specify the minimum required CIDR ranges.
7. (Optional) For **VPC ID**, specify the ID of a VPC to use for ODB peering.

8. Choose **Continue**, and then choose **Modify**.

Deleting an ODB network in Oracle Database@AWS

You can delete an ODB network. This section assumes that you have already created an ODB network in Step 1: Create an ODB network in Oracle Database@AWS. You can't delete an ODB network that is currently in use by a VM cluster.

To delete an ODB network

1. Sign in to the AWS Management Console and open the Oracle Database@AWS console at https://console.aws.amazon.com/odb/.

2. From the left pane, choose **ODB networks**.

3. Select the network that you want to delete.

4. Choose **Delete**.

5. (Optional) Choose **Delete associated OCI resources** to delete the OCI resources that were created along with the ODB network.

6. In the text box, enter `delete me`.

7. Choose **Delete**.

Deleting an Exadata VM cluster in Oracle Database@AWS

You can delete an Exadata VM cluster. This section assumes that you have already created an Exadata VM cluster in Step 3: Create an Exadata VM cluster or Autonomous VM cluster in Oracle Database@AWS.

To delete an Exadata VM cluster

1. Sign in to the AWS Management Console and open the Oracle Database@AWS console at https://console.aws.amazon.com/odb/.

2. From the left pane, choose **Exadata VM clusters**.

3. Choose an Exadata VM cluster to delete.

4. Choose **Delete**.

5. When prompted, enter `delete me` and then choose **Delete**.

Deleting an Oracle Exadata infrastructure in Oracle Database@AWS

You can delete an Oracle Exadata infrastructure. This section assumes that you have already created an Oracle Exadata infrastructure in [Step 2: Create an Oracle Exadata infrastructure in Oracle Database@AWS](#). You can't delete an Exadata infrastructure that is currently in use by a VM cluster.

To delete an Oracle Exadata infrastructure

1. Sign in to the AWS Management Console and open the Oracle Database@AWS console at https://console.aws.amazon.com/odb/.
2. From the left pane, choose **Exadata infrastructures**.
3. Choose an Exadata infrastructure to delete.
4. Choose **Delete**.
5. When prompted, enter **delete me** and then choose **Delete**.

Security in Oracle Database@AWS

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from data centers and network architectures that are built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS, OCI, and you. The shared responsibility model describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the AWS Compliance Programs.

- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors, including the sensitivity of your data, your organization's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Oracle Database@AWS. You also learn how to use other AWS services that help you to monitor and secure your Oracle Database@AWS resources.

You can manage access to your Oracle Database@AWS resources. The method you use to manage access depends on what type of task you need to perform with Oracle Database@AWS:

- Use AWS Identity and Access Management (IAM) policies to assign permissions that determine who is allowed to manage Oracle Database@AWS resources. For example, you can use IAM to determine who is allowed to create, describe, modify, and delete Exadata infrastucture, VM clusters or tag resources.

- Use the security features of your Oracle database engine to control who can log in to the databases on a DB instance. These features work just as if the database was on your local network.

- Use Secure Socket Layers (SSL) or Transport Layer Security (TLS) connections with Exadata databases. For more information, see Prepare for TLS Walletless Connections.

- Oracle Database@AWS isn't immediately accessible from the internet and deployed on private subnets in AWS only.

- Oracle Database@AWS uses many default Transmission Control Protocol (TCP) ports for various operations. For the full list of ports, see Default port assignments.

- To store and manage keys by using Transparent Data Encryption (TDE), which is enabled by default, Oracle Database@AWS uses OCI vaults or Oracle Key Vault. Oracle Database@AWS doesn't support AWS Key Management Service.

- By default, the database is configured by using Oracle-managed encryption keys. The database also supports customer-managed keys.

- To enhance data protection, use Oracle Data Safe with Oracle Database@AWS.

The following topics show you how to configure Oracle Database@AWS to meet your security and compliance objectives.

Topics

- Data protection in Oracle Database@AWS

- Identity and access management for Oracle Database@AWS

- Compliance validation for Oracle Database@AWS

- Resilience in Oracle Database@AWS

- Using service-linked roles for Oracle Database@AWS

- AWS managed policies for Oracle Database@AWS

- Oracle Database@AWS updates to AWS managed policies

Data protection in Oracle Database@AWS

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.

- Use SSL/TLS to communicate with AWS resources. We require TLS 1.2 and recommend TLS 1.3.

- Set up API and user activity logging with AWS CloudTrail. For information about using CloudTrail trails to capture AWS activities, see Working with CloudTrail trails in the *AWS CloudTrail User Guide*.

- Use AWS encryption solutions, along with all default security controls within AWS services.

- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.

- If you require FIPS 140-3 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see Federal Information Processing Standard (FIPS) 140-3.

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with Oracle Database@AWS or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Data encryption

Exadata databases use Oracle Transparent Data Encryption (TDE) to encrypt your data. Your data is also protected in temporary tablespaces, undo segments, redo logs and during internal database operations such as JOIN and SORT. For more information, see Data Security.

Encryption in transit

Exadata databases use native Oracle Net Services encryption and integrity capabilities to secure connections to the database. For more information, see Security of data in transit.

Key management

Transparent Data Encryption includes a keystore to securely store master encryption keys, and a management framework to securely and efficiently manage the keystore and perform key maintenance operations. For more information, see To administer Vault encryption keys.

Identity and access management for Oracle Database@AWS

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. IAM administrators control who can be *authenticated* (signed in)

and *authorized* (have permissions) to use Oracle Database@AWS resources. IAM is an AWS service that you can use with no additional charge.

Topics

- [Audience](#)

- [Authenticating with identities](#)

- [Managing access using policies](#)

- [How Oracle Database@AWS works with IAM](#)

- [Identity-based policy examples for Oracle Database@AWS](#)

- [Troubleshooting Oracle Database@AWS identity and access](#)

Audience

How you use AWS Identity and Access Management (IAM) differs, depending on the work that you do in Oracle Database@AWS.

Service user – If you use the Oracle Database@AWS service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more Oracle Database@AWS features to do your work, you might need additional permissions. Understanding how access is managed can help you request the right permissions from your administrator. If you cannot access a feature in Oracle Database@AWS, see [Troubleshooting Oracle Database@AWS identity and access](#).

Service administrator – If you're in charge of Oracle Database@AWS resources at your company, you probably have full access to Oracle Database@AWS. It's your job to determine which Oracle Database@AWS features and resources your service users should access. You must then submit requests to your IAM administrator to change the permissions of your service users. Review the information on this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with Oracle Database@AWS, see [How Oracle Database@AWS works with IAM](#).

IAM administrator – If you're an IAM administrator, you might want to learn details about how you can write policies to manage access to Oracle Database@AWS. To view example Oracle Database@AWS identity-based policies that you can use in IAM, see [Identity-based policy examples for Oracle Database@AWS](#).

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. You must be *authenticated* (signed in to AWS) as the AWS account root user, as an IAM user, or by assuming an IAM role.

You can sign in to AWS as a federated identity by using credentials provided through an identity source. AWS IAM Identity Center (IAM Identity Center) users, your company's single sign-on authentication, and your Google or Facebook credentials are examples of federated identities. When you sign in as a federated identity, your administrator previously set up identity federation using IAM roles. When you access AWS by using federation, you are indirectly assuming a role.

Depending on the type of user you are, you can sign in to the AWS Management Console or the AWS access portal. For more information about signing in to AWS, see How to sign in to your AWS account in the *AWS Sign-In User Guide*.

If you access AWS programmatically, AWS provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use AWS tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see AWS Signature Version 4 for API requests in the *IAM User Guide*.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, AWS recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see Multi-factor authentication in the *AWS IAM Identity Center User Guide* and AWS Multi-factor authentication in IAM in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see Tasks that require root user credentials in the *IAM User Guide*.

Federated identity

As a best practice, require human users, including users that require administrator access, to use federation with an identity provider to access AWS services by using temporary credentials.

A *federated identity* is a user from your enterprise user directory, a web identity provider, the AWS Directory Service, the Identity Center directory, or any user that accesses AWS services by using credentials provided through an identity source. When federated identities access AWS accounts, they assume roles, and the roles provide temporary credentials.

For centralized access management, we recommend that you use AWS IAM Identity Center. You can create users and groups in IAM Identity Center, or you can connect and synchronize to a set of users and groups in your own identity source for use across all your AWS accounts and applications. For information about IAM Identity Center, see What is IAM Identity Center? in the *AWS IAM Identity Center User Guide*.

IAM users and groups

An *IAM user* is an identity within your AWS account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see Rotate access keys regularly for use cases that require long-term credentials in the *IAM User Guide*.

An *IAM group* is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see Use cases for IAM users in the *IAM User Guide.*

IAM roles

An *IAM role* is an identity within your AWS account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. To temporarily assume an IAM role in the AWS Management Console, you can switch from a user to an IAM role (console). You can assume a

role by calling an AWS CLI or AWS API operation or by using a custom URL. For more information about methods for using roles, see Methods to assume a role in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

- **Federated user access** – To assign permissions to a federated identity, you create a role and define permissions for the role. When a federated identity authenticates, the identity is associated with the role and is granted the permissions that are defined by the role. For information about roles for federation, see Create a role for a third-party identity provider (federation) in the *IAM User Guide*. If you use IAM Identity Center, you configure a permission set. To control what your identities can access after they authenticate, IAM Identity Center correlates the permission set to a role in IAM. For information about permissions sets, see Permission sets in the *AWS IAM Identity Center User Guide*.

- **Temporary IAM user permissions** – An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.

- **Cross-account access** – You can use an IAM role to allow someone (a trusted principal) in a different account to access resources in your account. Roles are the primary way to grant cross-account access. However, with some AWS services, you can attach a policy directly to a resource (instead of using a role as a proxy). To learn the difference between roles and resource-based policies for cross-account access, see Cross account resource access in IAM in the *IAM User Guide*.

- **Cross-service access** – Some AWS services use features in other AWS services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.

  - **Forward access sessions (FAS)** – When you use an IAM user or role to perform actions in AWS, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see Forward access sessions.

  - **Service role** – A service role is an IAM role that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see Create a role to delegate permissions to an AWS service in the *IAM User Guide*.

- **Service-linked role** – A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

- **Applications running on Amazon EC2** – You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making AWS CLI or AWS API requests. This is preferable to storing access keys within the EC2 instance. To assign an AWS role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the EC2 instance to get temporary credentials. For more information, see Use an IAM role to grant permissions to applications running on Amazon EC2 instances in the *IAM User Guide*.

Managing access using policies

You control access in AWS by creating policies and attaching them to AWS identities or resources. A policy is an object in AWS that, when associated with an identity or resource, defines their permissions. AWS evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in AWS as JSON documents. For more information about the structure and contents of JSON policy documents, see Overview of JSON policies in the *IAM User Guide*.

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the `iam:GetRole` action. A user with that policy can get role information from the AWS Management Console, the AWS CLI, or the AWS API.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can

perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see Define custom IAM permissions with customer managed policies in the *IAM User Guide*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your AWS account. Managed policies include AWS managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see Choose between managed policies and inline policies in the *IAM User Guide*.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must specify a principal in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

Resource-based policies are inline policies that are located in that service. You can't use AWS managed policies from IAM in a resource-based policy.

Access control lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, AWS WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see Access control list (ACL) overview in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

AWS supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- **Permissions boundaries** – A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user

or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the `Principal` field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see [Permissions boundaries for IAM entities](#) in the *IAM User Guide*.

- **Service control policies (SCPs)** – SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in AWS Organizations. AWS Organizations is a service for grouping and centrally managing multiple AWS accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each AWS account root user. For more information about Organizations and SCPs, see [Service control policies](#) in the *AWS Organizations User Guide*.

- **Resource control policies (RCPs)** – RCPs are JSON policies that you can use to set the maximum available permissions for resources in your accounts without updating the IAM policies attached to each resource that you own. The RCP limits permissions for resources in member accounts and can impact the effective permissions for identities, including the AWS account root user, regardless of whether they belong to your organization. For more information about Organizations and RCPs, including a list of AWS services that support RCPs, see [Resource control policies (RCPs)](#) in the *AWS Organizations User Guide*.

- **Session policies** – Session policies are advanced policies that you pass as a parameter when you programmatically create a temporary session for a role or federated user. The resulting session's permissions are the intersection of the user or role's identity-based policies and the session policies. Permissions can also come from a resource-based policy. An explicit deny in any of these policies overrides the allow. For more information, see [Session policies](#) in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how AWS determines whether to allow a request when multiple policy types are involved, see [Policy evaluation logic](#) in the *IAM User Guide*.

How Oracle Database@AWS works with IAM

Before you use IAM to manage access to Oracle Database@AWS, learn what IAM features are available to use with Oracle Database@AWS.

IAM feature	Oracle Database@AWS support
Identity-based policies	Yes
Resource-based policies	No
Policy actions	Yes
Policy resources	Yes
Policy condition keys	Yes
ACLs	No
ABAC (tags in policies)	Partial
Temporary credentials	Yes
Principal permissions	Yes
Service roles	No
Service-linked roles	Yes

To get a high-level view of how Oracle Database@AWS and other AWS services work with most IAM features, see AWS services that work with IAM in the *IAM User Guide*.

Identity-based policies for Oracle Database@AWS

Supports identity-based policies: Yes

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see Define custom IAM permissions with customer managed policies in the *IAM User Guide*.

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. You can't specify the principal in an identity-based policy because it applies to the user or role to which it is attached. To learn about all

of the elements that you can use in a JSON policy, see [IAM JSON policy elements reference](#) in the *IAM User Guide*.

Identity-based policy examples for Oracle Database@AWS

To view examples of Oracle Database@AWS identity-based policies, see [Identity-based policy examples for Oracle Database@AWS](#).

Resource-based policies within Oracle Database@AWS

Supports resource-based policies: No

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must [specify a principal](#) in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

To enable cross-account access, you can specify an entire account or IAM entities in another account as the principal in a resource-based policy. Adding a cross-account principal to a resource-based policy is only half of establishing the trust relationship. When the principal and the resource are in different AWS accounts, an IAM administrator in the trusted account must also grant the principal entity (user or role) permission to access the resource. They grant permission by attaching an identity-based policy to the entity. However, if a resource-based policy grants access to a principal in the same account, no additional identity-based policy is required. For more information, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Policy actions for Oracle Database@AWS

Supports policy actions: Yes

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Action` element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Policy actions usually have the same name as the associated AWS API operation. There are some exceptions, such as *permission-only actions* that don't have a matching API

operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Include actions in a policy to grant permissions to perform the associated operation.

To see a list of Oracle Database@AWS actions, see Actions Defined by Oracle Database@AWS in the *Service Authorization Reference*.

Policy actions in Oracle Database@AWS use the following prefix before the action:

```
odb
```

To specify multiple actions in a single statement, separate them with commas.

```
"Action": [
      "odb:action1",
      "odb:action2"
         ]
```

To view examples of Oracle Database@AWS identity-based policies, see Identity-based policy examples for Oracle Database@AWS.

Policy resources for Oracle Database@AWS

Supports policy resources: Yes

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Resource` JSON policy element specifies the object or objects to which the action applies. Statements must include either a `Resource` or a `NotResource` element. As a best practice, specify a resource using its Amazon Resource Name (ARN). You can do this for actions that support a specific resource type, known as *resource-level permissions*.

For actions that don't support resource-level permissions, such as listing operations, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*"
```

To see a list of Oracle Database@AWS resource types and their ARNs, see Resources Defined by Oracle Database@AWS in the *Service Authorization Reference*. To learn with which actions you can specify the ARN of each resource, see Actions Defined by Oracle Database@AWS.

To view examples of Oracle Database@AWS identity-based policies, see Identity-based policy examples for Oracle Database@AWS.

Policy condition keys for Oracle Database@AWS

Supports service-specific policy condition keys: Yes

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Condition` element (or `Condition` *block*) lets you specify conditions in which a statement is in effect. The `Condition` element is optional. You can create conditional expressions that use condition operators, such as equals or less than, to match the condition in the policy with values in the request.

If you specify multiple `Condition` elements in a statement, or multiple keys in a single `Condition` element, AWS evaluates them using a logical AND operation. If you specify multiple values for a single condition key, AWS evaluates the condition using a logical OR operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see IAM policy elements: variables and tags in the *IAM User Guide*.

AWS supports global condition keys and service-specific condition keys. To see all AWS global condition keys, see AWS global condition context keys in the *IAM User Guide*.

To see a list of Oracle Database@AWS condition keys, see Condition Keys for Oracle Database@AWS in the *Service Authorization Reference*. To learn with which actions and resources you can use a condition key, see Actions Defined by Oracle Database@AWS.

To view examples of Oracle Database@AWS identity-based policies, see Identity-based policy examples for Oracle Database@AWS.

ACLs in Oracle Database@AWS

Supports ACLs: No

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

ABAC with Oracle Database@AWS

Supports ABAC (tags in policies): Partial

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes. In AWS, these attributes are called *tags*. You can attach tags to IAM entities (users or roles) and to many AWS resources. Tagging entities and resources is the first step of ABAC. Then you design ABAC policies to allow operations when the principal's tag matches the tag on the resource that they are trying to access.

ABAC is helpful in environments that are growing rapidly and helps with situations where policy management becomes cumbersome.

To control access based on tags, you provide tag information in the [condition element](#) of a policy using the aws:ResourceTag/*key-name*, aws:RequestTag/*key-name*, or aws:TagKeys condition keys.

If a service supports all three condition keys for every resource type, then the value is **Yes** for the service. If a service supports all three condition keys for only some resource types, then the value is **Partial**.

For more information about ABAC, see [Define permissions with ABAC authorization](#) in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see [Use attribute-based access control (ABAC)](#) in the *IAM User Guide*.

Using temporary credentials with Oracle Database@AWS

Supports temporary credentials: Yes

Some AWS services don't work when you sign in using temporary credentials. For additional information, including which AWS services work with temporary credentials, see [AWS services that work with IAM](#) in the *IAM User Guide*.

You are using temporary credentials if you sign in to the AWS Management Console using any method except a user name and password. For example, when you access AWS using your company's single sign-on (SSO) link, that process automatically creates temporary credentials. You also automatically create temporary credentials when you sign in to the console as a user and then

switch roles. For more information about switching roles, see [Switch from a user to an IAM role (console)](#) in the *IAM User Guide*.

You can manually create temporary credentials using the AWS CLI or AWS API. You can then use those temporary credentials to access AWS. AWS recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see [Temporary security credentials in IAM](#).

Cross-service principal permissions for Oracle Database@AWS

Supports forward access sessions (FAS): Yes

When you use an IAM user or role to perform actions in AWS, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see [Forward access sessions](#).

Service roles for Oracle Database@AWS

Supports service roles: No

A service role is an [IAM role](#) that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see [Create a role to delegate permissions to an AWS service](#) in the *IAM User Guide*.

> ⚠️ **Warning**
>
> Changing the permissions for a service role might break Oracle Database@AWS functionality. Edit service roles only when Oracle Database@AWS provides guidance to do so.

Service-linked roles for Oracle Database@AWS

Supports service-linked roles: Yes

A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS

account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

For details about creating or managing service-linked roles, see AWS services that work with IAM. Find a service in the table that includes a Yes in the **Service-linked role** column. Choose the **Yes** link to view the service-linked role documentation for that service.

Identity-based policy examples for Oracle Database@AWS

By default, users and roles don't have permission to create or modify Oracle Database@AWS resources. They also can't perform tasks by using the AWS Management Console, AWS Command Line Interface (AWS CLI), or AWS API. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

To learn how to create an IAM identity-based policy by using these example JSON policy documents, see Create IAM policies (console) in the *IAM User Guide*.

For details about actions and resource types defined by Oracle Database@AWS, including the format of the ARNs for each of the resource types, see Actions, Resources, and Condition Keys for Oracle Database@AWS in the *Service Authorization Reference*.

Topics

- Policy best practices
- Using the Oracle Database@AWS console
- Allow users to provision Oracle Database@AWS resources
- Allow users to view their own permissions

Policy best practices

Identity-based policies determine whether someone can create, access, or delete Oracle Database@AWS resources in your account. These actions can incur costs for your AWS account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with AWS managed policies and move toward least-privilege permissions** – To get started granting permissions to your users and workloads, use the *AWS managed policies* that grant permissions for many common use cases. They are available in your AWS account. We recommend that you reduce permissions further by defining AWS customer managed policies

that are specific to your use cases. For more information, see AWS managed policies or AWS managed policies for job functions in the *IAM User Guide*.

- **Apply least-privilege permissions** – When you set permissions with IAM policies, grant only the permissions required to perform a task. You do this by defining the actions that can be taken on specific resources under specific conditions, also known as *least-privilege permissions*. For more information about using IAM to apply permissions, see Policies and permissions in IAM in the *IAM User Guide*.

- **Use conditions in IAM policies to further restrict access** – You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific AWS service, such as AWS CloudFormation. For more information, see IAM JSON policy elements: Condition in the *IAM User Guide*.

- **Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions** – IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy checks and actionable recommendations to help you author secure and functional policies. For more information, see Validate policies with IAM Access Analyzer in the *IAM User Guide*.

- **Require multi-factor authentication (MFA)** – If you have a scenario that requires IAM users or a root user in your AWS account, turn on MFA for additional security. To require MFA when API operations are called, add MFA conditions to your policies. For more information, see Secure API access with MFA in the *IAM User Guide*.

For more information about best practices in IAM, see Security best practices in IAM in the *IAM User Guide*.

Using the Oracle Database@AWS console

To access the Oracle Database@AWS console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Oracle Database@AWS resources in your AWS account. If you create an identity-based policy that is more restrictive than the minimum required permissions, the console won't function as intended for entities (users or roles) with that policy.

You don't need to allow minimum console permissions for users that are making calls only to the AWS CLI or the AWS API. Instead, allow access to only the actions that match the API operation that they're trying to perform.

Allow users to provision Oracle Database@AWS resources

This policy allows users full access to provision Oracle Database@AWS resources. To setup DNS resolution from your VPC, you need to create an outbound Route 53 resolver and add rules to forward DNS traffic with the OCI domain name to OCI DNS listener IP.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowODBAndEC2Actions",
            "Action": [
                "odb:*",
                "ec2:DescribeVpcs",
                "ec2:CreateOdbNetworkPeering",
                "ec2:DeleteOdbNetworkPeering",
                "ec2:DescribeAvailabilityZones",
                "ec2:DescribeRouteTables",
                "ec2:CreateRoute",
                "route53resolver:CreateResolverEndpoint",
                "route53resolver:CreateResolverRule",
                "route53resolver:AssociateResolverRule"
            ],
            "Effect": "Allow",
            "Resource": "*"
        },
        {
            "Sid": "AllowSLRActions",
            "Effect": "Allow",
            "Action": [
                "iam:CreateServiceLinkedRole"
            ],
            "Resource": "*",
            "Condition": {
                "StringEquals": {
                    "iam:AWSServiceName": [
                        "odb.amazonaws.com"
                    ]
                }
            }
        }
    ]
}
```

Allow users to view their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the AWS CLI or AWS API.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ViewOwnUserInfo",
            "Effect": "Allow",
            "Action": [
                "iam:GetUserPolicy",
                "iam:ListGroupsForUser",
                "iam:ListAttachedUserPolicies",
                "iam:ListUserPolicies",
                "iam:GetUser"
            ],
            "Resource": ["arn:aws:iam::*:user/${aws:username}"]
        },
        {
            "Sid": "NavigateInConsole",
            "Effect": "Allow",
            "Action": [
                "iam:GetGroupPolicy",
                "iam:GetPolicyVersion",
                "iam:GetPolicy",
                "iam:ListAttachedGroupPolicies",
                "iam:ListGroupPolicies",
                "iam:ListPolicyVersions",
                "iam:ListPolicies",
                "iam:ListUsers"
            ],
            "Resource": "*"
        }
    ]
}
```

Troubleshooting Oracle Database@AWS identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with Oracle Database@AWS and IAM.

Topics

- [I am not authorized to perform an action in Oracle Database@AWS](#)
- [I am not authorized to perform iam:PassRole](#)
- [I want to allow people outside of my AWS account to access my Oracle Database@AWS resources](#)

I am not authorized to perform an action in Oracle Database@AWS

If you receive an error that you're not authorized to perform an action, your policies must be updated to allow you to perform the action.

The following example error occurs when the `mateojackson` IAM user tries to use the console to view details about a fictional *my-example-widget* resource but doesn't have the fictional `odb:`*GetWidget* permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform:
  odb:GetWidget on resource: my-example-widget
```

In this case, the policy for the `mateojackson` user must be updated to allow access to the *my-example-widget* resource by using the `odb:`*GetWidget* action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the `iam:PassRole` action, your policies must be updated to allow you to pass a role to Oracle Database@AWS.

Some AWS services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named `marymajor` tries to use the console to perform an action in Oracle Database@AWS. However, the action requires the service to have

permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws:iam::123456789012:user/marymajor is not authorized to perform:
 iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the `iam:PassRole` action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my AWS account to access my Oracle Database@AWS resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether Oracle Database@AWS supports these features, see How Oracle Database@AWS works with IAM.
- To learn how to provide access to your resources across AWS accounts that you own, see Providing access to an IAM user in another AWS account that you own in the *IAM User Guide*.
- To learn how to provide access to your resources to third-party AWS accounts, see Providing access to AWS accounts owned by third parties in the *IAM User Guide*.
- To learn how to provide access through identity federation, see Providing access to externally authenticated users (identity federation) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see Cross account resource access in IAM in the *IAM User Guide*.

Compliance validation for Oracle Database@AWS

To learn whether an AWS service is within the scope of specific compliance programs, see AWS services in Scope by Compliance Program and choose the compliance program that you are interested in. For general information, see AWS Compliance Programs.

You can download third-party audit reports using AWS Artifact. For more information, see
Downloading Reports in AWS Artifact.

Your compliance responsibility when using AWS services is determined by the sensitivity of your
data, your company's compliance objectives, and applicable laws and regulations. AWS provides the
following resources to help with compliance:

- Security Compliance & Governance – These solution implementation guides discuss architectural
  considerations and provide steps for deploying security and compliance features.
- HIPAA Eligible Services Reference – Lists HIPAA eligible services. Not all AWS services are HIPAA
  eligible.
- AWS Compliance Resources – This collection of workbooks and guides might apply to your
  industry and location.
- AWS Customer Compliance Guides – Understand the shared responsibility model through the
  lens of compliance. The guides summarize the best practices for securing AWS services and map
  the guidance to security controls across multiple frameworks (including National Institute of
  Standards and Technology (NIST), Payment Card Industry Security Standards Council (PCI), and
  International Organization for Standardization (ISO)).
- Evaluating Resources with Rules in the *AWS Config Developer Guide* – The AWS Config service
  assesses how well your resource configurations comply with internal practices, industry
  guidelines, and regulations.
- AWS Security Hub – This AWS service provides a comprehensive view of your security state within
  AWS. Security Hub uses security controls to evaluate your AWS resources and to check your
  compliance against security industry standards and best practices. For a list of supported services
  and controls, see Security Hub controls reference.
- Amazon GuardDuty – This AWS service detects potential threats to your AWS accounts,
  workloads, containers, and data by monitoring your environment for suspicious and malicious
  activities. GuardDuty can help you address various compliance requirements, like PCI DSS, by
  meeting intrusion detection requirements mandated by certain compliance frameworks.
- AWS Audit Manager – This AWS service helps you continuously audit your AWS usage to simplify
  how you manage risk and compliance with regulations and industry standards.

Resilience in Oracle Database@AWS

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions
provide multiple physically separated and isolated Availability Zones, which are connected with

low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about AWS Regions and Availability Zones, see [AWS Global Infrastructure](#).

In addition to the AWS global infrastructure, Oracle Database@AWS offers several features to help support your data resiliency and backup needs.

Using service-linked roles for Oracle Database@AWS

Oracle Database@AWS uses AWS Identity and Access Management (IAM) [service-linked roles](#). A service-linked role is a unique type of IAM role that is linked directly to Oracle Database@AWS. Service-linked roles are predefined by Oracle Database@AWS and include all the permissions that the service requires to call other AWS services on your behalf.

A service-linked role makes using Oracle Database@AWS easier because you don't have to manually add the necessary permissions. Oracle Database@AWS defines the permissions of its service-linked roles, and unless defined otherwise, only Oracle Database@AWS can assume its roles. The defined permissions include the trust policy and the permissions policy, and that permissions policy cannot be attached to any other IAM entity.

You can delete the roles only after first deleting their related resources. This protects your Oracle Database@AWS resources because you can't inadvertently remove permission to access the resources.

Service-linked role permissions for Oracle Database@AWS

Oracle Database@AWS uses the service-linked role named AWSServiceRoleForODB to allow Oracle Database@AWS to call AWS services on behalf of your resources.

The AWSServiceRoleForODB service-linked role trusts the following services to assume the role:

- `odb.amazonaws.com`

This service-linked role has a permissions policy attached to it called `AmazonODBServiceRolePolicy` that grants it permissions to operate in your account. For more information, see [AmazonODBServiceRolePolicy](#).

> ⓘ **Note**
>
> You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. If you encounter the following error message: **Unable to create the resource. Verify that you have permission to create service-linked role. Otherwise wait and try again later.**
> Make sure you have the following permissions enabled:
>
> ```
> {
>     "Action": "iam:CreateServiceLinkedRole",
>     "Effect": "Allow",
>     "Resource": "arn:aws:iam::*:role/aws-service-role/odb.amazonaws.com/
> AWSServiceRoleForODB",
>     "Condition": {
>         "StringLike": {
>             "iam:AWSServiceName":"odb.amazonaws.com"
>         }
>     }
> }
> ```
>
> For more information, see [Service-linked role permissions](#) in the *IAM User Guide*.

Creating a service-linked role for Oracle Database@AWS

You don't need to manually create a service-linked role. When you create an Exadata database, Oracle Database@AWS creates the service-linked role for you.

If you delete this service-linked role, and then need to create it again, you can use the same process to recreate the role in your account. When you create an Exadata database, Oracle Database@AWS creates the service-linked role for you again.

Editing a service-linked role for Oracle Database@AWS

Oracle Database@AWS does not allow you to edit the AWSServiceRoleForODB service-linked role. After you create a service-linked role, you cannot change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM For more information, see [Editing a service-linked role](#) in the *IAM User Guide*.

Deleting a service-linked role for Oracle Database@AWS

If you no longer need to use a feature or service that requires a service-linked role, we recommend that you delete that role. That way you don't have an unused entity that is not actively monitored or maintained. However, you must delete all of your resources before you can delete the service-linked role.

Cleaning up a service-linked role for Oracle Database@AWS

Before you can use IAM to delete a service-linked role, you must first confirm that the role has no active sessions and remove any resources used by the role.

To check whether the service-linked role has an active session in the IAM console

1. Sign in to the AWS Management Console and open the IAM console at https://console.aws.amazon.com/iam/.

2. In the navigation pane of the IAM console, choose **Roles**. Then choose the name (not the check box) of the AWSServiceRoleForODB role.

3. On the **Summary** page for the chosen role, choose the **Access Advisor** tab.

4. On the **Access Advisor** tab, review recent activity for the service-linked role.

> ⓘ **Note**
>
> If you're unsure whether Oracle Database@AWS is using the AWSServiceRoleForODB role, you can try to delete the role. If the service is using the role, then the deletion fails and you can view the AWS Regions where the role is being used. If the role is being used, then you must wait for the session to end before you can delete the role. You cannot revoke the session for a service-linked role.

If you want to remove the AWSServiceRoleForODB role, you must first delete all of your Oracle Database@AWS resources.

Supported Regions for Oracle Database@AWS service-linked roles

Oracle Database@AWS supports using service-linked roles in all of the AWS Regions where the service is available. For more information, see AWS Regions and Endpoints.

AWS managed policies for Oracle Database@AWS

To add permissions to permission sets and roles, it's easier to use AWS managed policies than to write policies yourself. It takes time and expertise to create IAM customer managed policies that provide your team with only the permissions they need. To get started quickly, you can use our AWS managed policies. These policies cover common use cases and are available in your AWS account. For more information about AWS managed policies, see AWS managed policies in the *IAM User Guide*.

AWS services maintain and update AWS managed policies. You can't change the permissions in AWS managed policies. Services occasionally add additional permissions to an AWS managed policy to support new features. This type of update affects all identities (permission sets and roles) where the policy is attached. Services are most likely to update an AWS managed policy when a new feature is launched or when new operations become available. Services don't remove permissions from an AWS managed policy, so policy updates don't break your existing permissions.

Additionally, AWS supports managed policies for job functions that span multiple services. For example, the `ReadOnlyAccess` AWS managed policy provides read-only access to all AWS services and resources. When a service launches a new feature, AWS adds read-only permissions for new operations and resources. For a list and descriptions of job function policies, see AWS managed policies for job functions in the *IAM User Guide*.

Topics

* AWS managed policy: AmazonODBServiceRolePolicy

AWS managed policy: AmazonODBServiceRolePolicy

You can't attach the `AmazonODBServiceRolePolicy` policy to your IAM entities. This policy is attached to a service-linked role that allows Oracle Database@AWS to perform actions on your behalf. For more information, see Using service-linked roles for Oracle Database@AWS.

To view more details about the policy, including the latest version of the JSON policy document, see AmazonODBServiceRolePolicy in the *AWS Managed Policy Reference Guide*.

Oracle Database@AWS updates to AWS managed policies

View details about updates to AWS managed policies for Oracle Database@AWS since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the Oracle Database@AWS Document history page.

Change	Description	Date
AWS managed policy: AmazonODBServiceRolePolicy – New service-linked role policy	Oracle Database@AWS added the `AmazonODB ServiceRolePolicy` for the `AWSServiceRoleForO DB` service-linked role. For more information, see AWS managed policy: AmazonODB ServiceRolePolicy.	December 2, 2024
Oracle Database@AWS started tracking changes	Oracle Database@AWS started tracking changes for its AWS managed policies.	December 2, 2024

Monitoring Oracle Database@AWS

Monitoring is an important part of maintaining the reliability, availability, and performance of Oracle Database@AWS and your other AWS solutions. AWS provides the following monitoring tools to watch Oracle Database@AWS, report when something is wrong, and take automatic actions when appropriate:

- *Amazon CloudWatch* monitors your AWS resources and and the applications you run on AWS in real time. You can collect and track metrics, create customized dashboards, and set alarms that notify you or take actions when a specified metric reaches a threshold that you specify. For example, you can have CloudWatch track CPU usage or other metrics of your Amazon EC2 instances and automatically launch new instances when needed. For more information, see the Amazon CloudWatch User Guide.

- *Amazon CloudWatch Logs* enables you to monitor, store, and access your log files from Amazon EC2 instances, CloudTrail, and other sources. CloudWatch Logs can monitor information in the log files and notify you when certain thresholds are met. You can also archive your log data in highly durable storage. For more information, see the Amazon CloudWatch Logs User Guide.

- *Amazon EventBridge* can be used to automate your AWS services and respond automatically to system events, such as application availability issues or resource changes. Events from AWS services are delivered to EventBridge in near real time. You can write simple rules to indicate which events are of interest to you and which automated actions to take when an event matches a rule. For more information, see Amazon EventBridge User Guide.

- *AWS CloudTrail* captures API calls and related events made by or on behalf of your AWS account and delivers the log files to an Amazon S3 bucket that you specify. You can identify which users and accounts called AWS, the source IP address from which the calls were made, and when the calls occurred. For more information, see the AWS CloudTrail User Guide.

Monitoring Oracle Database@AWS with Amazon CloudWatch

You can monitor Oracle Database@AWS using CloudWatch, which collects raw data and processes it into readable, near real-time metrics. These statistics are kept for 15 months, so that you can access historical information and gain a better perspective on how your web application or service is performing. You can also set alarms that watch for certain thresholds, and send notifications or take actions when those thresholds are met. For more information, see the Amazon CloudWatch User Guide.

Amazon CloudWatch metrics for Oracle Database@AWS

The Oracle Database@AWS service reports metrics to Amazon CloudWatch in the AWS/ODB namespace for VM clusters, container databases, and pluggable databases.

Topics

- [Metrics for cloud VM clusters](#)
- [Metrics for container databases](#)
- [Metrics for pluggable databases](#)

Metrics for cloud VM clusters

The Oracle Database@AWS service reports the following metrics in the AWS/ODB namespace for cloud VM clusters.

Metric	Description	Units
`ASMDiskgroupUtiliz ation`	The percentage of usable space used in a Disk Group. Usable space is the space available for growth. DATA disk group stores our Oracle database files. RECO disk group contains database files for recovery such as archives and flashback logs.	Percentage
`CpuUtilization`	The percent CPU utilization.	Percentage
`FilesystemUtilizat ion`	The percent utilization of provisioned filesystem.	Percentage
`LoadAverage`	The system load average over 5 minutes.	Integer
`MemoryUtilization`	The percentage of memory available for starting new applications, without	Percentage

Metric	Description	Units
	swapping. The available memory can be obtained via the following command: `cat /proc/meminfo`	
`NodeStatus`	Indicates whether the host is reachable.	Integer
`OcpusAllocated`	The number of OCPUs allocated.	Integer
`SwapUtilization`	The percent utilization of total swap space.	Percentage

Metrics for container databases

The Oracle Database@AWS service reports the following metrics in the AWS/ODB namespace for container databases.

Metric	Description	Units
`BlockChanges`	The Average number of blocks changed per second.	Changes per second
`CpuUtilization`	The CPU utilization expressed as a percentage, aggregated across all consumer groups. The utilization percentage is reported with respect to the number of CPUs the database is allowed to use, which is two times the number of OCPUs.	Percentage
`CurrentLogons`	The number of successful logons during the selected interval.	Count

Metric	Description	Units
ExecuteCount	The number of user and recursive calls that executed SQL statements during the selected interval.	Count
ParseCount	The number of hard and soft parses during the selected interval.	Count
StorageAllocated	Total amount of storage space allocated to the database at the collection time.	GB
StorageAllocatedBy Tablespace	Total amount of storage space allocated to the tablespace at the collectio n time. In case of container database, this metric provides root container tablespaces.	GB
StorageUsed	Total amount of storage space used by the database at the collection time.	GB
StorageUsedByTable space	Total amount of storage space used by tablespace at the collection time. In case of container database, this metric provides root container tablespaces.	GB

Metric	Description	Units
StorageUtilization	The percentage of provisioned storage capacity currently in use. Represents the total allocated space for all tablespaces.	Percentage
StorageUtilization ByTablespace	This indicates the percentage of storage space utilized by the tablespace at the collection time. In case of container database, this metric provides root container tablespaces..	Percentage
TransactionCount	The combined number of user commits and user rollbacks during the selected interval.	Count
UserCalls	The combined number of logons, parses, and execute calls during the selected interval.	Count

Metrics for pluggable databases

The Oracle Database@AWS service reports the following metrics in the AWS/ODB namespace for pluggable databases.

Metric	Description	Units
AllocatedStorageUt ilizationByTablesp ace	The percentage of space used by tablespace, out of all allocated. For container databases, this metric provides data for root container tablespaces.	Percent

Metric	Description	Units
	(Statistic: Mean, Interval: 30 minutes)	
`AvgGCCRBlockReceiveTime`	The average global cache CR (consistent-read) block receive time. For RAC / cluster databases only. (Statistic: Mean, Interval: 5 minutes)	Milliseconds
`AvgGCCurrentBlockReceiveTime`	The average global cache current blocks receive time. Statistic reports the mean value. For Real Application Cluster (RAC) databases only. (Statistic: Mean, Interval: 5 minutes)	Milliseconds
`BlockChanges`	The average number of blocks changed per second. (Statistic: Mean, Interval: 1 minute)	changes per second
`BlockingSessions`	Current blocking sessions. Not applicable for container databases. (Statistic: Max, Interval: 15 minutes)	Count
`CPUTimeSeconds`	The average rate of accumulation of CPU time by foreground sessions in the database instance over the time interval. The CPU time component of Average Active Sessions. (Statistic: Mean, Interval: 1 minute)	Seconds per second

Metric	Description	Units
CpuCount	The number of CPUs during the selected interval.	Count
CpuUtilization	The CPU utilization expressed as a percentage, aggregated across all consumer groups. The utilization percentage is reported with respect to the number of CPUs the database is allowed to use, which is two times the number of OCPUs. (Statistic: Mean, Interval: 1 minute)	Percent
CurrentLogons	The number of successful logons during the selected interval. (Statistics: Sum, Interval: 1 minute)	Count
DBTimeSeconds	The average rate of accumulation of database time (CPU + Wait) by foreground sessions in the database instance over the time interval. Also known as Average Active Sessions. (Statistic: Mean, Interval: 1 minute)	Seconds per second

Metric	Description	Units
DbmgmtJobExecution sCount	The number of SQL job executions on a single managed database or a database group, and their status. Status dimension s can be the following values: "Succeeded," "Failed," "InProgress." (Statistic: Sum, Interval: 1 minute)	Count
ExecuteCount	The number of user and recursive calls that executed SQL statements during the selected interval. (Statistic: Sum, Interval: 1 minute)	Count
FRASpaceLimit	The flash recovery area space limit. Not applicable for pluggable databases. (Statisti c: Max, Interval: 15 minutes)	GB
FRAUtilization	The flash recovery area utilization. Not applicable for pluggable databases. (Statisti c: Mean, Interval: 15 minutes)	Percent
GCCRBlocksReceived	The global cache CR (consiste nt-read) blocks received per second. For RAC / cluster databases only. (Statistic: Mean, Interval: 5 minutes)	Blocks per second

Metric	Description	Units
`GCCurrentBlocksReceived`	Represents global cache current blocks received per second. Statistic reports the mean value. For Real Application Cluster (RAC) databases only. (Statistic: Mean, Interval: 5 minutes)	Blocks per second
`IOPS`	The average number of input-output operations per second. (Statistic: Mean, Interval: 1 minute)	Operations per second
`IOThroughputMB`	The average throughput in MB per second. (Statistic: Mean, Interval: 1 minute)	MB per second
`InterconnectTrafficMB`	The average internode data transfer rate. For RAC / cluster databases only. (Statistic: Mean, Interval: 5 minutes)	MB per second
`InvalidObjects`	Invalid database objects count. Not applicable for container databases. (Statistic: Max, Interval: 24 hours)	Count
`LogicalBlocksRead`	The average number of blocks read from SGA/Memory (buffer cache) per second. (Statistic: Mean, Interval: 1 minute)	Reads per second

Metric	Description	Units
MaxTablespaceSize	The maximum possible tablespace size. For container databases, this metric provides data for root container tablespaces. (Statistic: Max, Interval: 30 minutes)	GB
MemoryUsage	Memory pool total size in MB. (Statistic: Mean, Interval: 15 minutes)	MB
MonitoringStatus	The monitoring status of the resource. If a metric collectio n fails, error information is captured in this metric. (Statistic: Mean, Interval: 5 minutes)	Not applicable
NonReclaimableFRA	The Non-reclaimable fast recovery area. Not applicabl e for pluggable databases. (Statistic: Mean, Interval: 15 minutes)	Percent
OcpusAllocated	The actual number of OCPUs allocated by the service during the selected interval of time. (Statistic: Count, Interval: 1 minute)	Integer
ParseCount	The number of hard and soft parses during the selected interval. (Statistic: Sum, Interval: 1 minute)	Count

Metric	Description	Units
ParsesByType	The number of hard or soft parses per second. (Statistic: Mean, Interval: 1 minute)	Parses per second
ProblematicSchedul edDBMSJobs	The problematic scheduled database jobs count. Not applicable for container databases. (Statistic: Max, Interval: 15 minutes)	Count
ProcessLimitUtiliz ation	The process limit utilization. Not applicable for pluggable databases. (Statistic: Mean, Interval: 1 minute)	Percent
Processes	The database processes count. Not applicable for pluggable databases. (Statisti c: Max, Interval: 1 minute)	Count
ReclaimableFRA	The reclaimable fast recovery area. Not applicable for pluggable databases. (Statisti c: Mean, Interval: 15 minutes)	Percent
ReclaimableFRASpace	The flash recovery area reclaimable space. Not applicable for pluggable databases. (Statistic: Mean, Interval: 15 minutes)	GB
RedoSizeMB	The average amount of redo generated, in MB per second. (Statistic: Mean, Interval: 1 minute)	MB per second

Metric	Description	Units
`SessionLimitUtiliz ation`	The session limit utilization. Not applicable for pluggable databases. (Statistic: Mean, Interval: 1 minute)	Percent
`Sessions`	The number of sessions in the database. (Statistic: Mean, Interval: 1 minute)	Count
`StorageAllocated`	The maximum amount of space allocated by tablespac e during the interval. For container databases, this metric provides data for root container tablespaces. (Statistic: Max, Interval: 30 minutes)	GB
`StorageAllocatedBy Tablespace`	The maximum amount of space allocated by tablespac e during the interval. For container databases, this metric provides data for root container tablespaces. (Statistic: Max, Interval: 30 minutes)	GB
`StorageUsed`	The maximum amount of space used during the interval. (Statistic: Max, Interval: 30 minutes)	GB

Metric	Description	Units
StorageUsedByTable space	The maximum amount of space used by tablespac e during the interval. For container databases, this metric provides data for root container tablespaces. (Statistic: Max, Interval: 30 minutes)	GB
StorageUtilization	The percentage of provision ed storage capacity currently in use. Represents the total allocated space for all tablespaces. (Statistic: Mean, Interval: 30 minutes)	Percent
StorageUtilization ByTablespace	The percentage of the space utilized, by tablespace. For container databases, this metric provides data for root container tablespaces. (Statistic: Mean, Interval: 30 minutes)	Percent
TransactionCount	The combined number of user commits and user rollbacks during the selected interval. (Statistic: Sum, Interval: 1 minute)	Count
TransactionsByStatus	The number of committed or rolled back transactions per second. (Statistic: Mean, Interval: 1 minute)	Transactions per second

Metric	Description	Units
UnusableIndexes	Unusable indexes count in database schema. Not applicable for container databases. (Statistic: Max, Interval: 24 hours)	Count
UsableFRA	The useable fast recovery area. Not applicable for pluggable databases. (Statistic: Mean, Interval: 15 minutes)	Percent
UsedFRASpace	The flash recovery area space usage. Not applicable for pluggable databases. (Statistic: Max, Interval: 15 minutes)	GB
UserCalls	The combined number of logons, parses, and execute calls during the selected interval. (Statistic: Sum, Interval: 1 minute)	Count
WaitTimeSeconds	The average rate of accumulation of non-idle wait time by foreground sessions in the database instance over the time interval. The wait time component of Average Active Sessions. (Statistic: Mean, Interval: 5 minutes)	Seconds per second

Amazon CloudWatch dimensions for Oracle Database@AWS

You can filter Oracle Database@AWS metrics data by using any dimension in the following table.

Dimension	Filters the requested data for . . .
cloudVmClusterId	The identifier of a VM cluster.
cloudExadataInfrastructureId	The identifier of the Exadata infrastructure.
collectionName	A name of a collection.
deploymentType	The type of infrastructure.
diskgroupName	A name of a disk group
errorCode	An error code.
errorSeverity	The severity of an error.
filesystemName	The name of a file system.
hostName	The name of the host machine.
instanceName	The name of a database instance.
instanceNumber	The instance number of a database instance.
ioType	A type of I/O operation.
jobId	A unique identifier for a job.
managedDatabaseGroupId	The identifier of a Managed Database Group.
managedDatabaseId	The identifier of a Managed Database.
memoryPool	A type of memory pool.
memoryType	A type of memory.
ociCloudVmClusterId	The OCI identifier of a VM cluster.

Dimension	Filters the requested data for . . .
ociCloudExadataInfrastructureId	The OCI identifier of the Exadata infrastructure.
parseType	A type of parse.
resourceId	The identifier of a resource.
resourceId_Database	The identifier of a database.
resourceId_DbNode	The identifier of a database node.
resourceName	The name of a resource.
resourceName_Database	The name of a database.
resourceName_DbNode	The name of a database node.
resourceType	A type of database.
schemaName	The name of a schema.
status	The status of a database.
tablespaceContents	The contents of a tablespace.
tablespaceName	The name of a tablespace.
tablespaceType	A type of tablespace.
transactionStatus	The status of a transaction.
type	A type of Problematic Scheduled DMS job.
waitClass	A class of wait event.

Logging Oracle Database@AWS API calls using AWS CloudTrail

Oracle Database@AWS is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service. CloudTrail captures all API calls for Oracle Database@AWS as events. The calls captured include calls from the Oracle Database@AWS console and code calls to the Oracle Database@AWS API operations. Using the information collected by CloudTrail, you can determine the request that was made to Oracle Database@AWS, the IP address from which the request was made, when it was made, and additional details.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root user or user credentials.
- Whether the request was made on behalf of an IAM Identity Center user.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

> **ⓘ Note**
>
> Oracle Database@AWS records `GetCallerIdentity` API calls from AWS Security Token Service (STS) in your CloudTrail logs. These STS API calls verify the identity of Oracle Database@AWS when interacting with OCI on your behalf. They are a normal and secure part of AWS operations and do not expose sensitive information.

CloudTrail is active in your AWS account when you create the account and you automatically have access to the CloudTrail **Event history**. The CloudTrail **Event history** provides a viewable, searchable, downloadable, and immutable record of the past 90 days of recorded management events in an AWS Region. For more information, see Working with CloudTrail Event history in the *AWS CloudTrail User Guide*. There are no CloudTrail charges for viewing the **Event history**.

For an ongoing record of events in your AWS account past 90 days, create a trail or a CloudTrail Lake event data store.

CloudTrail trails

A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. All trails created using the AWS Management Console are multi-Region. You can create a single-Region or a multi-Region

trail by using the AWS CLI. Creating a multi-Region trail is recommended because you capture activity in all AWS Regions in your account. If you create a single-Region trail, you can view only the events logged in the trail's AWS Region. For more information about trails, see Creating a trail for your AWS account and Creating a trail for an organization in the *AWS CloudTrail User Guide*.

You can deliver one copy of your ongoing management events to your Amazon S3 bucket at no charge from CloudTrail by creating a trail, however, there are Amazon S3 storage charges. For more information about CloudTrail pricing, see AWS CloudTrail Pricing. For information about Amazon S3 pricing, see Amazon S3 Pricing.

CloudTrail Lake event data stores

CloudTrail Lake lets you run SQL-based queries on your events. CloudTrail Lake converts existing events in row-based JSON format to Apache ORC format. ORC is a columnar storage format that is optimized for fast retrieval of data. Events are aggregated into *event data stores*, which are immutable collections of events based on criteria that you select by applying advanced event selectors. The selectors that you apply to an event data store control which events persist and are available for you to query. For more information about CloudTrail Lake, see Working with AWS CloudTrail Lake in the *AWS CloudTrail User Guide*.

CloudTrail Lake event data stores and queries incur costs. When you create an event data store, you choose the pricing option you want to use for the event data store. The pricing option determines the cost for ingesting and storing events, and the default and maximum retention period for the event data store. For more information about CloudTrail pricing, see AWS CloudTrail Pricing.

Oracle Database@AWS management events in CloudTrail

Management events provide information about management operations that are performed on resources in your AWS account. These are also known as control plane operations. By default, CloudTrail logs management events.

Oracle Database@AWS logs all Oracle Database@AWS control plane operations as management events.

Oracle Database@AWS event examples

An event represents a single request from any source and includes information about the requested API operation, the date and time of the operation, request parameters, and so on. CloudTrail log

files aren't an ordered stack trace of the public API calls, so events don't appear in any specific
order.

The following example shows a CloudTrail event that demonstrates the `CreateOdbNetwork`
operation.

```
{
    "eventVersion": "1.09",
    "userIdentity": {
        "type": "AssumedRole",
        "principalId": "AKIAIOSFODNN7EXAMPLE:yourRole",
        "arn": "arn:aws:sts::123456789012:assumed-role/Admin/yourRole",
        "accountId": "123456789012",
        "accessKeyId": "AKIAI44QH8DHBEXAMPLE",
        "sessionContext": {
            "sessionIssuer": {
                "type": "Role",
                "principalId": "AKIAIOSFODNN7EXAMPLE",
                "arn": "arn:aws:iam::123456789012:role/Admin",
                "accountId": "123456789012",
                "userName": "Admin"
            },
            "attributes": {
                "creationDate": "2024-11-06T21:17:29Z",
                "mfaAuthenticated": "false"
            }
        }
    },
    "eventTime": "2024-11-06T21:17:44Z",
    "eventSource": "odb.amazonaws.com",
    "eventName": "CreateOdbNetwork",
    "awsRegion": "us-east-1",
    "sourceIPAddress": "192.0.2.0",
    "userAgent": "python-requests/2.28.2",
    "requestParameters": {
        "availabilityZoneId": "use1-az6",
        "backupSubnetCidr": "123.45.6.7/89",
        "clientSubnetCidr": "123.44.6.7/89",
        "clientToken": "testClientToken",
        "defaultDnsPrefix": "testLabel",
        "displayName": "yourOdbNetwork"
    },
    "responseElements": {
```

```
            "displayName": "yourOdbNetwork",
            "odbNetworkId": "odbnet_1234567",
            "status": "PROVISIONING"
        },
        "requestID": "daf2e3f5-96a3-4df7-a026-863f96db793e",
        "eventID": "797163d3-5726-441d-80a7-6eeb7464acd4",
        "readOnly": false,
        "eventType": "AwsApiCall",
        "managementEvent": true,
        "recipientAccountId": "123456789012",
        "eventCategory": "Management",
        "tlsDetails": {
            "tlsVersion": "TLSv1.2",
            "cipherSuite": "ECDHE-RSA-AES128-GCM-SHA256",
            "clientProvidedHostHeader": "odb.us-east-1.amazonaws.com"
        }
 }
```

For information about CloudTrail record contents, see CloudTrail record contents in the *AWS CloudTrail User Guide.*

Troubleshooting Oracle Database@AWS

Use the following sections to help troubleshoot networking issues you may encounter with Oracle Database@AWS.

Topics

- [Creation of ODB network fails](#)
- [Connectivity issues between your VPC and ODB network or VM clusters](#)
- [Unresolvable hostnames or scannames of VM clusters from VPC](#)

Creation of ODB network fails

When you can't create a ODB network, the following are common causes:

- **Restricted CIDR Ranges** – The ODB network uses specific CIDR ranges for the client and backup subnets. Ensure that the CIDR ranges you've chosen for these subnets do not overlap with any restricted or reserved IP address ranges.

  The following CIDR ranges are reserved and cannot be used for the ODB network:

  - Oracle cloud reserved range: 169.254.0.0/16

  - Reserved Class D: 224.0.0.0 - 239.255.255.255

  - Reserved Class E: 240.0.0.0 - 255.255.255.255

  - Future OCI use: 100.105.0.0/16

  Follow the EC2 rules for CIDR ranges as outlined in the VPC documentation. To learn more, see [CIDR block association restrictions](#).

  Additionally, avoid overlap between specified CIDR ranges and those used for VPC connectivity to the ODB network.

- **Overlapping VPC CIDR** – The CIDR range you've specified for the ODB network should not overlap with the CIDR ranges used by any of your existing VPCs. Overlapping CIDR ranges can cause routing conflicts and prevent the successful creation of the ODB network. Check the CIDR ranges of ODB peering VPCs and ensure the ODB network CIDR is unique and non-overlapping.

- **Ownership of VPCs** – The ODB network and the VPC you're connecting to must be owned by the same AWS account. If you're trying to peer the ODB network to a VPC owned by a different

account, the creation will fail. Verify that the ODB network and VPC are both owned by the same AWS account.

Connectivity issues between your VPC and ODB network or VM clusters

When you can't connect from your VPC to the ODB network or the VM clusters within it, the following are common causes:

- **Verifying VPC configuration** – In the Oracle Database@AWS console, locate the VPC that is peered with the ODB network. Confirm the VPC ID matches the one shown in the ODB network details.

- **Inspecting route tables** – In the Amazon VPC console, find the route table attached to the subnet where your application is running. Check for a route with a destination CIDR that matches the client subnet CIDR of the ODB network. Confirm that this route points to the correct ODB network ARN. If the route is missing, add a new one to the ODB network's client subnet CIDR.

- **Validating peered CIDRs** – Review the `Peered CIDRs` section in the ODB network details. Confirm all the relevant CIDR blocks from your VPC are listed. If a required CIDR is missing, update the peered CIDRs.

- **Checking security group rules** – In the Amazon EC2 console, locate the security groups for resources in your VPC. Review the inbound and outbound rules, updating them as needed to permit the necessary traffic.

- **Confirming Availability Zones** – In the Amazon VPC console, identify the Availability Zone (AZ) of your subnet. Verify that the ODB network is also deployed in the same AZ as your subnet.

- **Avoiding multiple ODB network peerings** – Check your VPC peering connections in the Oracle Database@AWS Console. Make sure you have only one active connection to an ODB network. If you see more than one ODB network peering, remove the extra ones.

Unresolvable hostnames or scannames of VM clusters from VPC

If the hostnames or scannames of the VM clusters are not resolvable from your VPC, you need to configure DNS forwarding on the VPC and the following resources to resolve DNS records hosted on the ODB network.

- An outbound endpoint to send DNS queries to the ODB network. For more information, see Configuring an outbound endpoint in an ODB network in Oracle Database@AWS.

- A resolver rule to specify the domain name of the DNS queries that the resolver forwards to the DNS for ODB network. For more information, see Configuring a resolver rule in Oracle Database@AWS.

Quotas for Oracle Database@AWS

Your AWS account has default quotas, formerly referred to as limits, for each AWS service. Unless otherwise noted, each quota is Region-specific. You can request increases for some quotas, and other quotas cannot be increased.

To view the quotas for Oracle Database@AWS, open the Service Quotas console. In the navigation pane, choose **AWS services** and select **Oracle Database@AWS**.

To request a quota increase, see Requesting a Quota Increase in the *Service Quotas User Guide*. If the quota is not yet available in Service Quotas, use the limit increase form.

Your AWS account has the following quotas related to Oracle Database@AWS.

Resource	Default	Description
Exadata infrastructure	1000	The maximum number of Exadata infrastructures allowed in the current Region.
VM clusters	1000	The maximum number of VM clusters allowed in the current Region.
ODB network	5	The maximum number of ODB networks allowed in the current Region.

Document history for the Oracle Database@AWS User Guide

The following table describes the documentation releases for Oracle Database@AWS.

Change	Description	Date
Oracle Database@AWS supports Autonomous VM clusters	You can now create Autonomous VM clusters on your Exadata infrastructure. Autonomous VM clusters are fully managed databases that automate key management tasks using machine learning and AI. For more information, see Step 3: Create an ExaVM or Autonomous VM cluster in Oracle Database@AWS.	May 28, 2025
Oracle Database@AWS supports customizable maintenance windows	You can now configure maintenance windows for your Exadata infrastructure with options for Oracle-ma naged or Customer-managed schedules. You can also select patching modes (Rolling or Non-rolling) and specify maintenance timing preferenc es. For more information, see Create an Oracle Exadata infrastructure in Oracle Database@AWS.	May 1, 2025
Oracle Database@AWS supports a new Availability Zone (AZ)	You can now create an ODB network in an AZ with the physical ID use1-az4	March 26, 2025

or use1-az6. For more information, see Oracle Exadata infrastructure.

Oracle Database@AWS supports Amazon VPC Transit Gateways	If you connect a transit gateway to a VPC that is peered to an ODB network, you can connect multiple VPCs to this gateway. Applications running in these VPCs can access an Exadata VM cluster running in your ODB network. For more information, see Configuring Amazon VPC Transit Gateways for Oracle Database@AWS.	March 26, 2025
Oracle Database@AWS supports database and storage server types for Exadata X11M	You can specify the database server type and storage server type when you create an infrastructure using Exadata X11M. For more informati on, see Create an Oracle Exadata infrastructure in Oracle Database@AWS.	February 4, 2025
New service-linked role policy	Oracle Database@AWS added a new policy AmazonODB ServiceRolePolicy for the AWSServiceRoleForO DB service-linked role. For more information, see Oracle Database@AWS updates to AWS managed policies.	December 2, 2024

| Initial release | Initial release of the Oracle Database@AWS User Guide | December 2, 2024 |