ORACLE
Cloud Infrastructure

MapR Data Platform Reference Architecture for Oracle Cloud Infrastructure Deployments

ORACLE WHITE PAPER | OCTOBER 2018

ORACLE

Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Revision History

The following revisions have been made to this white paper since its initial publication:

Date	Revision
October 19, 2018	Initial publication

You can find the most recent versions of the Oracle Cloud Infrastructure white papers at https://cloud.oracle.com/iaas/technical-resources.

Table of Contents

Overview

This document details best practices for running MapR Data Platform (MapR) on Oracle Cloud Infrastructure. Although individual use cases and requirements might vary and require different approaches, the practices described here represent the ideal configuration for both performance and security for implementation on Oracle Cloud Infrastructure. Topics covered in this document include installation automation, automated configuration and tuning, and best practices for deployment and topology to support security and high availability.

Customers of MapR and Oracle Cloud Infrastructure can now run MapR deployments in the cloud, leveraging the power of Oracle Cloud Infrastructure bare metal and virtual machine instances to drive flexible, easily scalable, and performant MapR clusters in an automated fashion by using Terraform.

This cloud reference architecture represents best practices for sizing and deployment on Oracle Cloud Infrastructure. See the Appendix for links to the latest MapR documentation as supplemental reference material.

Terminology

If you are unfamiliar with Oracle Cloud Infrastructure, see the reference section in the Appendix for definitions of the basic components.

Infrastructure Guidance

All MapR deployments on Oracle Cloud Infrastructure leverage either bare metal or virtual machine instances for cost-efficient, highly performant, fast cloud infrastructure. The choice of which instances to use is up to you and is configurable as part of the deployment, but this section outlines some best practices to follow.

Terraform templates available on GitHub are preconfigured with the recommended instance types. Changing the instance types as part of the deployment might result in an unsupported cluster configuration with less than the required memory and disk capacity, so be aware of this before making changes. The preconfigured instance types comply with MapR's requirements.

Compute Considerations

When you are choosing the architecture for your MapR deployment on Oracle Cloud Infrastructure, there are many options to consider. This section provides some guidelines about the instances that MapR supports.

Bare Metal Compute

Bare metal DenseIO instances on Oracle Cloud Infrastructure provide fast, reliable compute power and NVME-based local storage. Two profiles are supported for bare metal instances running MapR as data nodes; these differ based on compute, memory, and storage density.

- BM.DenseIO1.36

  This instance provides 36 OCPUs (72 vCores), 512 GB of memory, and 28.8 TB local NVMe storage. Additional block storage can be attached, up to 1 PB per host.

- BM.DenseIO2.52

  This instance provides 52 OCPUs (104 vCores), 768 GB of memory, and 51.2 TB local NVMe storage. Additional block storage can be attached, up to 1 PB per host.

For more information about these compute profiles, including performance-related metrics, see the High Performance X7 Compute Service Review and Analysis blog post.

Virtual Machine Compute

MapR can be deployed on VMs using block storage for MapR-FS. When you use VMs, be sure to consider IOPS and bandwidth constraints when configuring your deployment.

VM.Standard and BM.Standard Instances

It is possible to configure MapR deployment to leverage BM.Standard and VM.Standard instances as data nodes. Before you assign an instance to a particular role, review the minimum required shapes per role in the Terraform Template section.

Storage Considerations

Oracle Cloud Infrastructure has several offerings to consider when choosing which storage to use for MapR-FS, or for other purposes in your MapR deployment.

Bare Metal NVME Storage

Oracle Cloud Infrastructure's bare metal NVME storage provides the fastest MapR-FS option for MapR on Oracle Cloud Infrastructure. This model uses bare metal instances, which have local NVME-based storage as the underlying capacity for MapR-FS. Bare metal NVME storage is the highest-performant storage option for MapR on Oracle Cloud Infrastructure, and we recommend it for production deployments.

Block Volumes

Oracle Cloud Infrastructure Block Volumes provides a cost effective means for securely and reliably storing data while maintaining performance. Block storage volumes are completely flexible in configuration, ranging from 50 GB to 32 TB per volume in 1-GB increments. Each instance can have a maximum of 32 volumes attached.

Oracle has a guaranteed SLA on Block Volumes, ensuring 3K IOPS and 24 MB/s per 50 GB of block storage, up to a maximum of 25K IOPS and 320MB/s per volume. A block storage volume peaks at 700 GB for IOPS and throughput. This aggregates at the host level and is something that you should consider if you choose to use block storage as MapR-FS. If the aggregate volume bandwidth is not high enough, MapR-FS stability during load can become problematic for instance types with large CPU and memory capacity, or for large clusters. This is usually not a concern for smaller deployments.

Network Considerations

Oracle provides a guaranteed networking SLA for instance and block storage bandwidth. For detailed bandwidth information for each instance, see Compute Shapes.

Networking on Oracle Cloud Infrastructure uses virtual cloud networks (VCNs) as the basis for all connectivity. For basic information about VCNs, read the FAQ.

VCNs support the concept of security lists to manage security and network access. Security lists are used in combination with host-level firewalls to limit or permit access to services run on instances in Oracle Cloud Infrastructure.

VCNs are local to each region and can span multiple availability domains. Multiple subnets can exist inside a single VCN and availability domain. Subnets must have a unique CIDR inside each VCN.

Instances have virtual network interface cards (VNICs), which are attached to specific subnets inside the availability domain to which they are attached. Instances and VNICs can only be a part of the same availability domain.

- BM.DenseIO1.36 instances support 10 Gbps, with a maximum of 16 VNICs per instance.
- BM.DenseIO2.52 instances support dual 25 Gbps, with a maximum of 24 VNICs per instance (12 per physical NIC).

On-Premises Connectivity

Oracle Cloud Infrastructure supports private connectivity across your on-premises and cloud networks. You can extend your IT infrastructure with connectivity services that offer predictable and consistent performance, isolation, and availability. This feature enables you to leverage a hybrid deployment model, which allows for versatile uses of Oracle Cloud Infrastructure as part of your Big Data ecosystem.

For more information about this connectivity, see the Fast Connect FAQ.

MapR on Oracle Cloud Infrastructure Deployment Recommendations

Cluster Architecture

MapR cluster architecture on Oracle Cloud Infrastructure follows supported reference architecture from MapR. A basic cluster consists of a minimum of five data nodes for production and three data nodes for development, which also run core services required for cluster operation. In addition, a bastion host is leveraged for access to the cluster.

In Oracle Cloud Infrastructure, a bastion host is considered the same as an edge host; the terminology is simply different. A bastion host is where edge services are configured and interaction with the cluster occurs. The bastion host should have a public IP address so that it can be accessed outside the VCN, and access should be restricted though security lists.

Data nodes run MapR-FS and YARN, as well as core cluster services, and are the target for all jobs inside the cluster. Data nodes should be deployed on a private network, not directly accessible from the internet. Access to UI elements on data nodes should be done in a secure manner (either VPN or SSH passthrough on the bastion host).

The following tables shows service roles and host types:

Service	Data Nodes (1-3)	Data Nodes (4+)	Bastion Hosts
MapR Core	CLDB, Data Host	Data Host	Client
YARN	Job History Server, Resource Manager, Node Manager	Node Manager	Client
Spark	History Server		
ZooKeeper	ZooKeeper Service		

- **Data Nodes (1-3)** run cluster service daemons for the Container Location Database (CLDB), Resource Manager, and ZooKeeper. These nodes also run an instance of the MapR Control System (MCS) console for cluster UI interaction, along with Node Manager for job execution, and act as MapR hosts for file, table, and stream storage and replication.
- **Data Nodes (4+)** run Node Manager for job execution, and act as MapR hosts for file, table, and stream storage and replication.
- **Bastion Hosts** act as edge nodes for user interaction and job submission for the cluster. These hosts are also where third-party software should be installed for use with the MapR cluster.

Network Architecture

The recommended network architecture for MapR deployment on Oracle Cloud Infrastructure consists of a VCN containing three separate subnets, which are duplicated across all availability domains in a target region. This architecture enables you to deploy an MapR cluster in any availability domain in the region and have the same topology and security lists associated with each network.

- **Bastion network**

  The bastion network is used as an edge network, has direct access to the internet, and is where the bastion hosts are deployed. Instances in this network have both a public and a private IP address. This network acts as an entry point for accessing cluster resources, while not exposing those services directly to the internet.

- **Public network**

  The public network is secondary to the bastion network. It has direct access to the internet and public and private IP addresses for each instance associated with it. You can deploy additional hosts to this network to segregate the management of internet-facing hosts, and it's useful for deploying third-party applications that interact with the MapR cluster.

- **Private network**

  The private network should have only private IP addresses for all instances associated with it. This network is more secure because the instances on it can't be accessed directly from the internet. Data nodes are deployed on this network, which provides additional security for services and data on those instances.

Access to all of these networks is controlled by security lists. Security lists are whitelists that allow network connectivity between the internet and subnets, along with subnet interaction inside a VCN. For more information, see Security Lists.

By design, there is no deny rule for network traffic on Oracle Cloud Infrastructure in a VCN because the default behavior is to deny. The only way for traffic to route is to create a security list rule that allows the traffic, whether it's allowing the entire network segment internal access between subnets in the VCN or allowing a specific host IP address or network access to the bastion host.

Automation deployment for MapR on Oracle Cloud Infrastructure using Terraform creates the MapR VCN and associated subnets automatically. The network CIDR used for the VCN is an entire Class B 10-net, and each subnet is programmatically set as a unique Class C network member. SSH access to hosts with public IP addresses is enabled by default, and a few specific ports are enabled with global access via security lists for ease of access. These configurations are completely customizable after deployment, and we recommend that you review the rules and adjust them to meet your network security requirements.

Network Topology

The recommended network topology for a MapR deployment consists of a single VCN in the region of the customer's choice. This VCN should contain nine subnets, three in each availability domain, for the bastion, public, and private networks. This model allows for granular control of hosts deployed in each subnet by using security lists. This network model is illustrated in the following diagram, with host associations at the subnet level and showing a single availability domain.

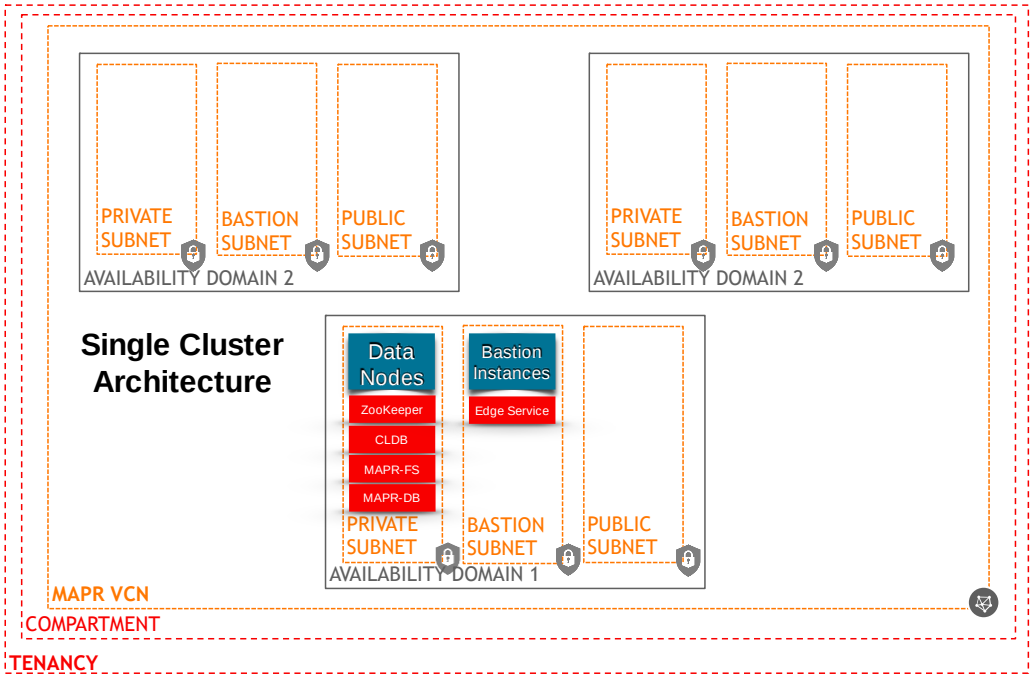

Connectivity and Security

Connectivity between hosts inside the VCN is controlled by a combination of security lists and local firewalls. Any connection between hosts is required to exist both in a security list and the local firewall on the hosts where the connection is needed. Security list rules are global in the sense that they allow a particular port or port range across all hosts inside the subnets associated with the security list. There is no host-level control at the security list level; host-level control is applied only at the local firewall level. This setup makes it important to manage security lists in a manner that is most restrictive to allowed traffic into subnets that are publicly addressable.

For this reason, we recommend keeping host-level firewalls in place across all deployed hosts. Many Hadoop vendors suggest disabling the local firewall for connectivity, but that security model is appropriate only for noncloud deployments. Connectivity at the host level can be whitelisted for internal networks in a broad manner, and fine-grained control for external access can also be applied with iptables (EL6) or firewalld (EL7). Information about how to leverage connectivity using these firewalls can be found readily online.

For more information about security best practices with automation, see the readme.

Automated Cluster Deployment with Terraform and the Oracle Cloud Infrastructure Provider

Deploying MapR on Oracle Cloud Infrastructure by using Terraform and the Oracle Cloud Infrastructure Provider is automated, fast, and flexible. The template is available from the Oracle Cloud Infrastructure Cloud Partners GitHub repository. Provisioning a fully ready cluster typically takes about 45 minutes, requiring minimal user interaction after you set a few configuration values in the Terraform template.

Detailed steps for deploying MapR on Oracle Cloud Infrastructure are located in the readme file in the GitHub repository. The deployment templates leverage Terraform by HashiCorp. Detailed setup instructions for Terraform are located on the HashiCorp website, and complementary information is located in the Oracle Cloud Infrastructure Provider GitHub repository.

Installation Model Overview

At a high level, the deployment process leverages the Terraform deployer to invoke Oracle Cloud Infrastructure API calls, which provision infrastructure inside the customer tenancy. A compartment is targeted for the deployment, where a VCN is set up with nine subnets, which are duplicated across each availability domain to allow deployment to any availability domain in the region. A bastion subnet is set up for the bastion hosts, a public subnet for additional public hosts and services, and a private subnet for data nodes. Hosts are then provisioned in these subnets in the target availability domain.

After all infrastructure provisioning is completed, the following steps are performed. These steps are illustrated in the diagram that follows.

1. An automated setup script is triggered to run on the bastion host.
2. Hosts are bootstrapped.
3. The script mirrors the MapR repository on the bastion host to facilitate the deployment of software dependencies and MapR cluster software on data nodes in the private subnet.
4. MapR setup is triggered, which sets up the cluster using a MapR Advanced Stanza template that is generated dynamically.

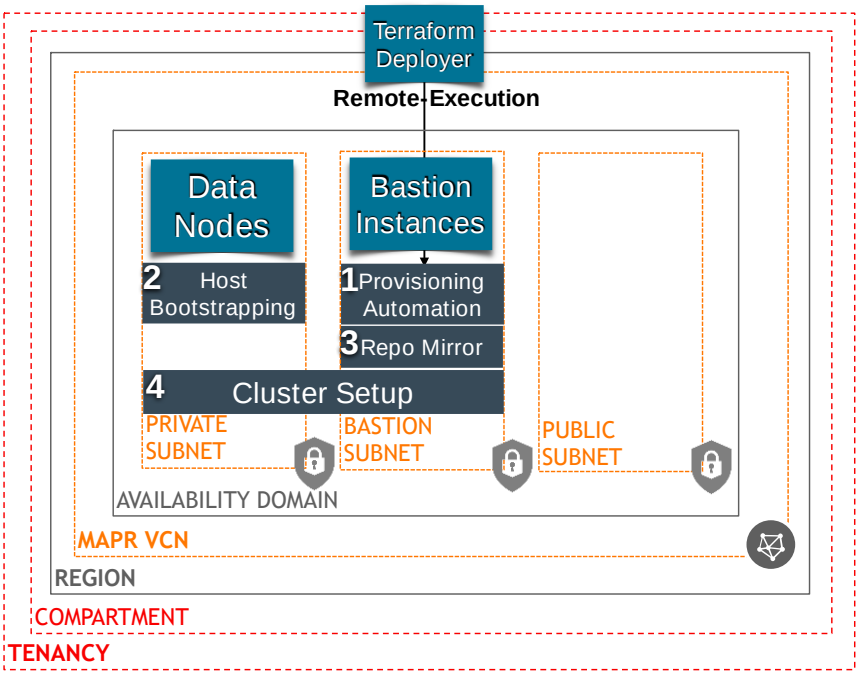

Single Availability Domain Deployment Model

Hosts are deployed and configured for MapR in a single availability domain. This is the only vendor-supported architecture; availability-domain spanning is not supported. If you want redundancy inside a region, consider using fault domains (a feature on the Oracle Cloud Infrastructure roadmap) or deploying a separate MapR cluster to another availability domain in the target region and using volume mirroring between the two clusters. This same architecture also applies at the regional redundancy level.

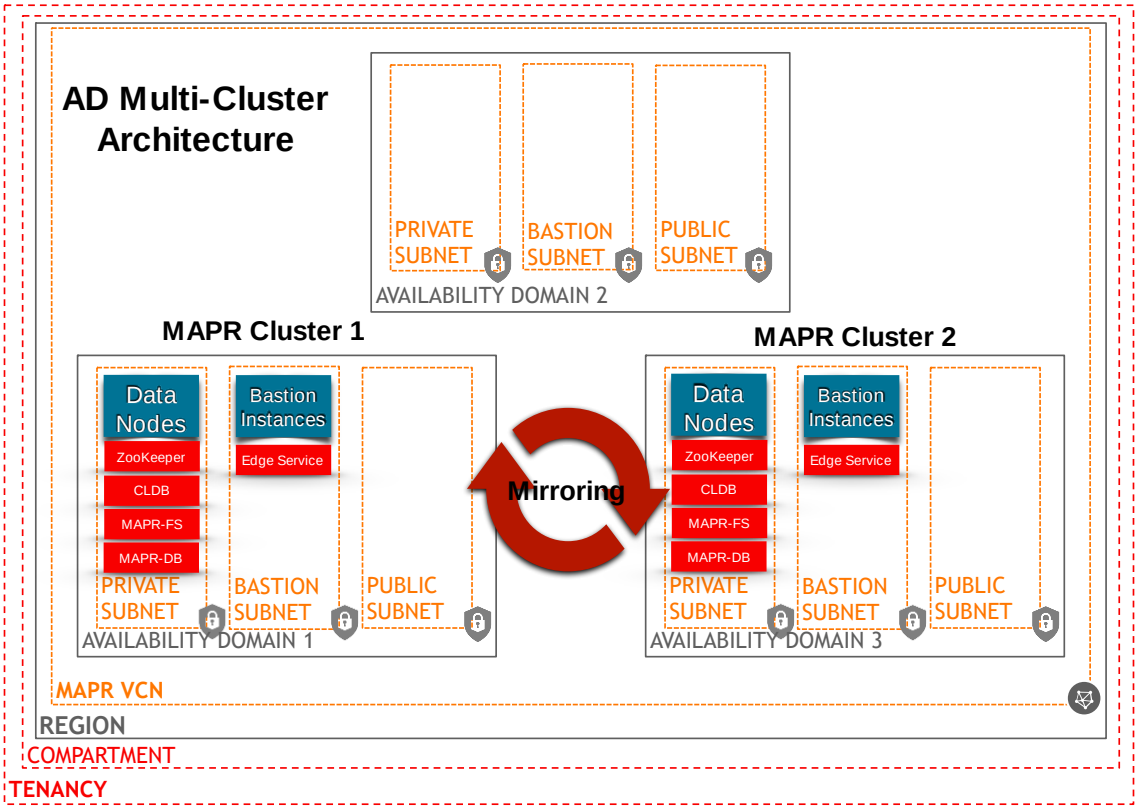

Terraform Template

The Terraform template that is available to automatically deploy a MapR cluster on Oracle Cloud Infrastructure is an N-Node template, which allows for a dynamic number of data nodes to be deployed with MapR.

Oracle Cloud Infrastructure supports N-Node MapR implementations for customers whose needs might exceed the performance or capacity limitations of the largest preset cluster configuration. Contact Oracle Cloud Infrastructure for more information. We will work with you to provide guidance on the optimal cluster deployment for your needs, and have an automated solution to support dynamic cluster sizes scaling into the thousands of nodes.

- **Minimum data node shape:** BM.DenseIO1.36
- **Suggested data node shape:** BM.DenseIO2.52
- **Minimum bastion shape:** VM.Standard1.4
- **Suggested bastion shape:** VM.Standard2.4

MapR Configuration Recommendations

MapR-FS

We recommend that you configure MapR-FS volumes with a replication factor of three for Bare metal MapR clusters. Because these hosts use local NVME storage for MapR-FS, redundancy should be built in to the MapR-FS topology to ensure high availability and failure tolerance.

ZooKeeper

ZooKeeper is set up by default on data nodes (1-3). We recommend maintaining an odd number of ZooKeeper instances for efficiently establishing a zookeeper quorum.

CLDB

For high availability, we recommend provisioning multiple instances of CLDB as part of the MapR deployment. This typically consists of three CLDBs on data nodes (1-3). When you are building larger clusters (hundreds or thousands of data nodes), we recommend scaling this to five CLDBs.

Summary

Automated deployment with Terraform provides a flexible, highly scalable framework for MapR on Oracle Cloud Infrastructure. Combined with bare metal performance on Oracle Cloud Infrastructure's fast network, this is an excellent solution for customers who want to explore MapR on the Oracle Cloud Infrastructure platform, leverage cloud for a low-cost alternative to on-premises deployments, or even offload entire Hadoop ecosystems to the cloud.

MapR on Oracle Cloud Infrastructure delivers a cost-effective, performant means to enable customer Big Data workloads in the cloud.

MapR on Oracle Cloud Infrastructure

MapR on Oracle Cloud Infrastructure is a joint solution that combines the power of Oracle Cloud Infrastructure with the performance of MapR. This joint solution allows for large, scalable data management using MapR, deployed by leveraging the flexibility and performance of Oracle Cloud Infrastructure. This solution provides a powerful, cost-efficient, easy-to-manage platform for running diverse Big Data workloads in the cloud.

Appendix

Related Links

- MapR website
- MapR documentation
- Oracle Cloud Infrastructure documentation
- Oracle Cloud Infrastructure Provider GitHub
- Terraform template for MapR on Oracle Cloud Infrastructure
- Bare metal and VM shape reference

Terminology Reference

This section provides definitions for some Oracle Cloud Infrastructure components.

Regions and Availability Domains

Oracle Cloud Infrastructure is hosted in regions and availability domains. A region is a localized geographic area, and an availability domain is one or more data centers located within a region. A region is composed of several availability domains. Most Oracle Cloud Infrastructure resources are either region-specific, such as a virtual cloud network, or availability domain–specific, such as a compute instance or block storage volume.

Oracle Cloud Infrastructure has many regions where you can deploy MapR clusters. Each cluster is localized to that specific region and targets a specific availability domain inside that region. With the release of fault domains, you can configure a “rack topology” to provide enhanced high availability for MapR deployments.

For more information, see Regions and Availability Domains.

Virtual Cloud Network

A virtual cloud network (VCN) is a customizable and private network in Oracle Cloud Infrastructure. Just like a traditional data center network, the VCN provides you with complete control over your network environment, which includes assigning your own private IP address space, creating subnets and route tables, and configuring stateful firewalls. A single tenant can have multiple VCNs, thereby providing grouping and isolation of related resources. Oracle’s new 25 Gb network infrastructure offers significantly more bandwidth and allows enterprises to cost effectively take full advantage of compute, storage, and database services.

For more information, see Overview of Networking.

Security Lists

A security list provides a virtual firewall for an instance, with ingress and egress rules that specify the types of traffic allowed in and out. Each security list is enforced at the instance level. However, you configure your security lists *at the subnet level*, which means that all instances in a given subnet are subject to the same set of rules. The security lists apply to a given instance whether it's talking with another instance in the VCN or a host outside the VCN.

For more information, see Security Lists.

Compute Service Instances

Oracle Cloud Infrastructure Compute lets you provision and manage compute hosts, or instances. You can launch instances as needed to meet your compute and application requirements. After you launch an instance, you can access it securely from your computer, restart it, attach and detach volumes, and terminate it when you're done with it. Any changes made to the instance's local drives are lost when you terminate it. Any saved changes to volumes attached to the instance are retained.

Oracle Cloud Infrastructure offers both bare metal and virtual machine instances:

- **Bare metal:** A bare metal compute instance gives you dedicated physical server access for highest performance and strong isolation.
- **Virtual machine:** A virtual machine (VM) is an independent computing environment that runs on top of physical bare metal hardware. The virtualization makes it possible to run multiple VMs that are isolated from each other. VMs are ideal for running applications that don’t require the performance and resources (CPU, memory, network bandwidth, and storage) of an entire physical machine.

An Oracle Cloud Infrastructure VM compute instance runs on the same hardware as a bare metal instance, leveraging the same cloud-optimized hardware, firmware, software stack, and networking infrastructure.

For more information, see Overview of the Compute Service.

Service Limits

When you sign up for Oracle Cloud Infrastructure, a set of service limits is configured for your tenancy. The service limit is the quota or allowance that is set on a resource. For example, your tenancy is allowed a maximum number of compute instances per availability domain. These limits

are generally established with your Oracle sales representative when you purchase Oracle Cloud Infrastructure. If you did not establish limits with your Oracle sales representative or if you signed up through the Oracle Store, default or trial limits are set for your tenancy. You can request to have a service limit raised.

For more information and a list of the default service limits, see Service Limits.

Identity and Access Management

Oracle Cloud Infrastructure Identity and Access Management (IAM) lets you control who has access to your cloud resources. You can control what type of access a group of users has and to which specific resources. You can write policies to control access to all of the services within Oracle Cloud Infrastructure.

For more information, see Overview of Identity and Access Management.

ORACLE

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US


blogs.oracle.com/oracle


facebook.com/oracle


twitter.com/oracle


oracle.com

Integrated Cloud Applications & Platform Services

MapR Data Platform Reference Architecture for Oracle Cloud Infrastructure Deployments
October 2018
Author: Zachary Smith