ORACLE
Cloud Infrastructure

Compute

Level 100

Rohit Rahi

November 2018

ORACLE

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Safe Harbor Statement

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.

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Objectives

After completing this lesson, you should be able to:

- Describe the OCI Compute Service
- Describe compute instance shapes, lifecycle and images
- Explain advanced features such as instance configuration and pools
- Create and launch a compute instance and walk through features

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Compute: Bare Metal & Virtual Machines

Bare Metal (BM)

Direct Hardware Access – customers get the full Bare Metal server
(single-tenant model)

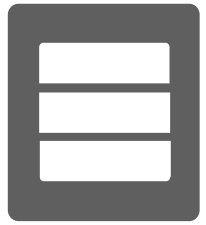

Bare Metal Server

Virtual Machine (VM)

A hypervisor to virtualize the underlying Bare Metal server into smaller VMs
(multi-tenant model)

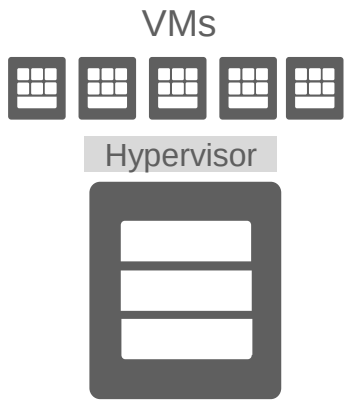

Bare Metal Server

VM compute instances runs on the same hardware as a Bare Metal instances, leveraging the same cloud-optimized hardware, firmware, software stack, and networking infrastructure

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Bare Metal

Direct Hardware Access with all the Security, Capabilities, Elasticity and Scalability of Oracle Cloud Infrastructure

Workloads that are Performance-intensive

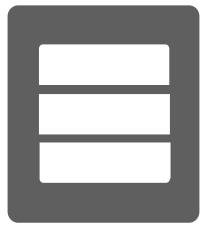

Workloads that are not virtualized

Workloads that require a specific hypervisor

Workloads that require BYO Licensing

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Bare Metal Instances

Shape	Instance type	OCPU	RAM (GB)	Local Disk (TB)	Network Bandwidth	Max vNICs
BM.Standard2.52	X7 based Standard compute	52	768	Block Storage only	2 x 25 Gbps	52
BM.DenseIO2.52	X7 based Dense I/O compute	52	768	51.2 TB NVMe SSD	2 x 25 Gbps	52
BM.Standard1.36	X5 based Standard compute	36	256	Block Storage only	10 Gbps	36
BM.DenseIO1.36	X5 based Dense I/O compute	36	512	28.8 TB NVMe SSD	10 Gbps	36
BM.Standard.E2.64	E1 AMD based Standard compute	64	512	Block Storage only	2 x 25 Gbps	150
BM.HPC2.36	X7 based - 3.7 GHz	36	384	6.7 TB NVMe SSD	1 x 100 Gbps RDMA	50

- Compute Standard E2 is based of AMD EPYC™ processor
- 2 x 25 Gbps implies two NIC cards with 25 Gbps bandwidth
- Network bandwidth is based on expected bandwidth for traffic within a VCN.
- Max vNICs numbers are for Linux. For Windows check the documentation

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Use cases for AMD EPYC based instances

- AMD EPYC Bare Metal server (64 cores, 512 GB RAM, 2 x 25 Gbps bandwidth, 75 vNICs) available at $0.03 core/hour, 66% cheaper than other options
- AMD EPYC based instances ideal for maximizing price performance
- Supported for Oracle applications, including E-Business Suite, JD Edwards, and PeopleSoft
- Certified to run Cloudera, Hortonworks, MapR, and Transwarp
- On a 10-TB full TeraSort benchmark, including TeraGen, TeraSort and TeraValidate, the AMD EPYC based instance demonstrated a 40 percent reduction in cost / OCPU v/s x86 alternatives with only a very slight increase in run times
- On a 4-node, 14M cell Fluent CFD simulation of an aircraft wing, the AMD EPYC based instance demonstrated a 30 percent reduction in cost along with a slight reduction in overall run times as compared to an x86 alternative

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

VM Standard Instances (Intel X7 and AMD)

Shape	Instance type	OCPU	RAM (GB)	Local Disk (TB)	Network Bandwidth	Max vNIC
VM.Standard2.1	Standard	1	15	Block Storage only	1 Gbps	2
VM.Standard2.2	Standard	2	30	Block Storage only	2 Gbps	2
VM.Standard2.4	Standard	4	60	Block Storage only	4.1 Gbps	4
VM.Standard2.8	Standard	8	120	Block Storage only	8.2 Gbps	8
VM.Standard2.16	Standard	16	240	Block Storage only	16.4 Gbps	16
VM.Standard2.24	Standard	24	320	Block Storage only	24.6 Gbps	24
VM.Standard.E2.1	Standard	1	8	Block Storage only	700 Mbps	2
VM.Standard.E2.2	Standard	2	16	Block Storage only	1.4 Gbps	2
VM.Standard.E2.4	Standard	4	32	Block Storage only	2.8 Gbps	4
VM.Standard.E2.8	Standard	8	64	Block Storage only	5.6 Gbps	4

- Compute Standard E2 is based of AMD EPYC™ processor
- Network bandwidth is based on expected bandwidth for traffic within a VCN

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

VM Dense I/O Instances

Shape	Instance type	OCPU	RAM (GB)	Local Disk (TB)	Network Bandwidth	Max vNIC
VM.DenseIO2.8	Dense I/O	8	60	6.4 TB NVMe SSD	8.2 Gbps	8
VM.DenseIO2.16	Dense I/O	16	240	12.8 TB NVMe SSD	16.4 Gbps	16
VM.DenseIO2.24	Dense I/O	24	320	25.6 NVMe SSD	24.6 Gbps	24

- Network bandwidth is based on expected bandwidth for traffic within a VCN

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

GPU Instances

Shape	Instance type	GPU	OCPU	RAM (GB)	Local Disk (TB)	Network Bandwidth
VM.GPU3.1	Volta VM GPU	1 x V100	6	90	Block Storage only	400 Mbps
VM.GPU3.2	Volta VM GPU	2 x V100	12	180	Block Storage only	800 Mbps
BM.GPU3.4	Volta BM GPU	4 x V100	24	360	Block Storage only	1 x 25 Gbps
BM.GPU3.8	Volta BM GPU	8 x V100	52	768	Block Storage only	2 x 25 Gbps
VM.GPU2.1	Pascal VM GPU	1 x P100	12	104	Block Storage only	1 x 25 Gbps
VM.GPU2.2	Pascal BM GPU	2 x P100	28	192	Block Storage only	2 x 25 Gbps

- Network bandwidth is based on expected bandwidth for traffic within a VCN

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Oracle provided Images

A template of a virtual hard drive that determines the operating system and other software for an instance. Images can be Oracle-provided, Custom, or BYOI

Oracle provides several pre-built images for Oracle Linux, Microsoft Windows, Ubuntu and CentOS

Image	Name
Oracle Linux	Oracle-Linux-7.x-<date>-<number>, Oracle-Linux-6.x-<date>-<number>
CentOS 7	CentOS-7-x-<date>-<number>, CentOS-6.x-<date>-<number>
Ubuntu 16.04 LTS	Canonical-Ubuntu-16.x-<date>-<number>, Canonical-Ubuntu-14.x-<date>-<number>
Windows Server 2012 R2	Windows-Server-2012-R2-<edition>-<gen>.<date>-<number>
Windows Server 2008 R2 - VM	Windows-Server-2008-R2-Standard-Edition-VM-<date>-<number>
Windows Server 2016	Windows-Server-2016-Datacenter-Edition-Gen2.<date>-<number>

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Oracle Provided Images

Linux Images

- User name opc created automatically for instances created from Oracle Linux/CentOS
- User name ubuntu created automatically for instances created from Ubuntu image
- These users have sudo privileges and are configured for remote access over the SSH v2
- Default set of firewall rules allow only SSH access (port 22)
- Provide a startup script using cloud-init

Windows Images

- User name opc created automatically with an OTP (one time password)
- Include the Windows Update utility to get the latest Windows updates from Microsoft

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Custom Images

- Create a custom image of an instance's boot disk and use it to launch other instances
- Instances you launch from your custom image include customizations, configuration, and software installed when you created the image
- During the process, instance shuts down and remains unavailable for several minutes. The instance restarts when the process completes
- Custom images do not include the data from any attached block volumes
- A custom image **cannot exceed 300 GB**
- **Windows custom images cannot be exported or downloaded out of the tenancy**

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Image Import/Export

- Compute service enables you to share custom images across tenancies and regions using image import/export
- Image import/export uses OCI Object Storage service
- You can import Linux and Windows Operating System
- Supports:
  - **Emulation Mode**: Fully emulated NIC, block boot, legacy BIOS.
  - **Native Mode:** Offer maximum performance with modern OS's

- You can also find more information about custom images here: https://cloud.oracle.com/iaas/whitepapers/deploying_custom_os_images.pdf

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Bring Your Own Image (BYOI)

- You can bring your own Operating System images or older Operating Systems
- Build new RHEL 7.4 images for Bare Metal and VM instances using a Terraform template
- Bring your own Hypervisor (KVM) – install and configure KVM on Bare Metal instances
- Bring Oracle VM – bring Oracle VM workloads to OCI
- Bring Image for Emulation Mode VM – import new and legacy OS using the QCOW2 or VMDK formats, and run them on Compute VMs using emulated hardware

Image Name	Supported versions
RHEL	4.5, 5.9, 5.11, 6.9, 7.4
CentOS	4.0, 4.8, 5.11, 6.9, 7.x
Oracle Linux	4.5, 4.8, 5.11, 6.2, 6.5, 6.9, 7.4
Ubuntu	12.04, 14.04, 16.04
Windows Server	2008 R2 Standard, Enterprise, Datacenter 2012 Standard, Datacenter 2012 R2 Standard, Datacenter 2016 Standard, Datacenter

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Boot Volumes

- A compute instance is launched using OS image stored on a remote boot volume
- Boot volume is created automated and associated with an instance until you terminate the instance
- Boot volumes are encrypted, have faster performance, lower launch times, and higher durability for BM and VM instances
- Compute instance can be scaled to a larger shape by using boot volumes
- You can preserve the boot volume when you terminate a compute instance
- Boot volumes are only terminated when you manually delete them
- Boot volumes cannot be detached from a running instance
- Possible to take a manual backup, assign backup policy or create clone of boot volumes

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Custom Boot Volumes

- You have the option of specifying a custom boot volume size
- In order to take advantage of the larger size, you must first extend the root (Linux-based images) or system (Windows-based images) partition

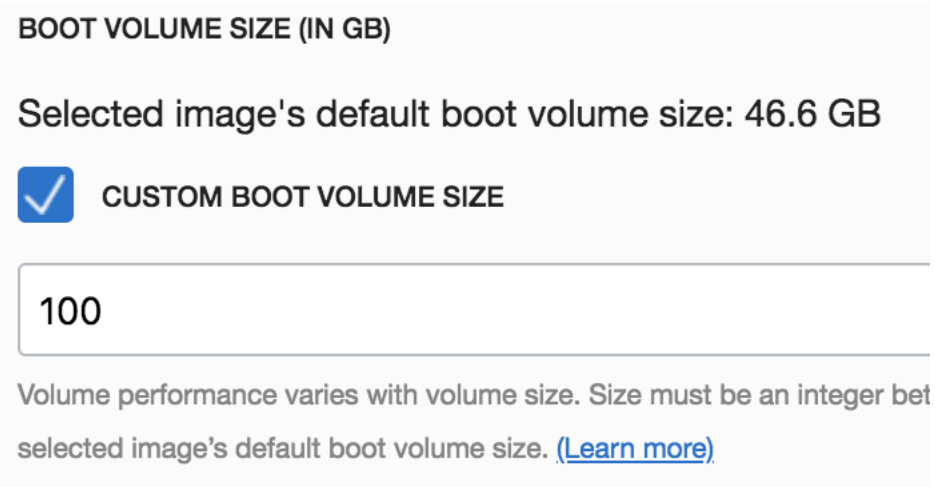

Linux default size is 46.6 GB

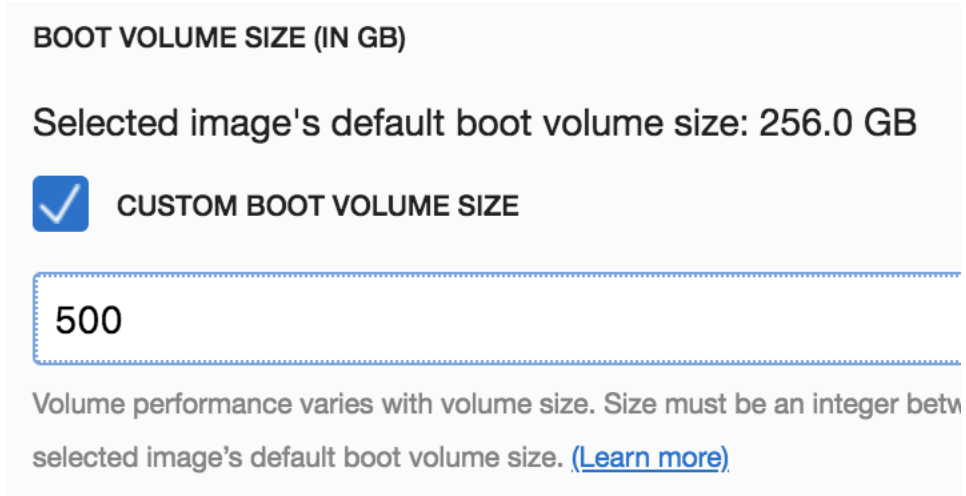

Windows default size is 256GB

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Custom Image v/s Boot Volume Backup

Custom Images

Pros	Cons
You can export a custom image across regions and tenancies	Instance shuts down and remains unavailable for several minutes until the process finished
No cost associated to store your custom images	Limit of 25 custom images per compartment

Boot volume Backup

Pros	Cons
It doesn’t require a downtime	Cost associated with the amount of Object Storage used to store your backup
Preserve the entire state of your running operating system as a backup	Creating a boot volume backup while instance is running creates a crash-consistent backup

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Instance Configuration and Pool

Instance Configurations

- Define the configuration information to launch a compute instance
- Include parameters (OS image, metadata, shape) and related resources as a single configuration entity, so you don’t have to specify them every time you launch a new instance
- Configure attached storage volumes; VNIC, Subnets and AD placements all with a single request

Instance Pools

- Provision and create multiple compute instances based off of the same instance configuration, within the same region
- Scale up/down
- 1 pool:1 configuration, but 1 configuration: n pools

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Instance Metadata

- Instance Metadata includes its OCID, name, compartment, shape, region, AD, creation date, state, image, and any custom metadata such as an SSH public key
- Service runs on every instance and is an HTTP endpoint listening on 169.254.169.254
- Get instance metadata by logging in to the instance and using the metadata service
- Oracle provided Linux instances
  - curl http://169.254.169.254/opc/v1/instance/
  - curl http://169.254.169.254/opc/v1/instance/metadata/
  - curl http://169.254.169.254/opc/v1/instance/metadata/<key-name>/
- Add and update custom metadata for an instance using CLI or SDK

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Instance Life Cycle

- Start – Restarts a stopped instance. After the instance is restarted, the Stop action is enabled
- Stop – Shuts down the instance. After the instance is powered off, the Start action is enabled
- Reboot – Shuts down the instance, and then restarts it
- Terminate – Permanently delete instances that you no longer need (vNICs and volumes are auto detached; boot volume is not deleted, but you can delete it)
- Resource Billing
  - Standard VM and BM instances, billing pauses in a STOP state
  - High I/O BM instances and dense I/O BM and VM instances , billing continues even in STOP state

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Compute Demo

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Compute Pricing (Bare Metal)

Shape	Instance type	OCPU	RAM (GB)	PAYG (OCPU/Hr.)
Bare Metal servers	X7 Standard compute	52	768	$0.0638
Bare Metal servers	X5 Standard compute	36	256	$0.0638
AMD Bare Metal servers	AMD Standard compute	64	512	$0.03
Bare Metal Dense IO	X7 Dense I/O compute	52	768	$0.1275
Bare Metal Dense IO	X5 Dense I/O compute	36	512	$0.1275
Bare Metal HPC	X7 High Frequency	36	384	$0.075

https://cloud.oracle.com/compute/pricing

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Compute Pricing (VM)

Shape	Instance type	OCPU	RAM (GB)	PAYG (OCPU/Hr.)
VM.Standard2.1 – VM.Standard2.24*	Gen 2 Standard VMs	1-24	15-320	$0.0638
VM.Standard1.1 – VM.Standard1.16	Gen 1 Standard VMs	1-16	7-112	$0.0638
VM.Standard.E2.1 – VM.Standard.E2.8	AMD Standard VMs	1-8	8-64	$0.03
VM.DenseIO2.8 – VM.DenseIO2.24	Gen 2 Dense IO VMs	8-24	60-320	$0.1275
VM.DenseIO1.4 VM.DenseIO1.16	Gen 1 Dense IO VMs	4-16	60-240	$0.1275
Compute Windows OS	1 OCPU Windows Server Standard			$0.0204

An OCPU is defined as the CPU capacity equivalent of one physical core of an Intel Xeon processor with hyper threading enabled. For Intel Xeon processor each OCPU corresponds to two hardware execution threads, known as vCPUs

https://cloud.oracle.com/compute/pricing

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Compute Pricing (GPU)

Shape	Instance type	GPU	OCPU	RAM (GB)	PAYG (GPU/hr.)
VM.GPU3.1	Volta VM GPU	1 x V100	6	90	$1.275
VM.GPU3.2	Volta VM GPU	2 x V100	12	180	$1.275
BM.GPU3.4	Volta BM GPU	4 x V100	24	360	$2.25
BM.GPU3.8	Volta BM GPU	8 x V100	52	768	$2.25
VM.GPU2.1	Pascal VM GPU	1 x P100	12	104	$2.25
VM.GPU2.2	Pascal BM GPU	2 x P100	28	192	$2.25

https://cloud.oracle.com/compute/pricing

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

Summary

- OCI Compute Service offers both Bare Metal and Virtual Machine instances
- Bare Metal instances provide direct hardware access and highest level of performance and isolation
- VMs and BMs run on the same cloud-optimized hardware, firmware, software stack, and networking infrastructure
- Offer a wide variety of shapes with industry leading price/performance
- Supports both x7 and AMD EPYC based instances with industry leading price/performance
- Image options - Oracle-provided images, BYOI, custom images, image import/export
- Advanced features include instance configuration and instance pools

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

ORACLE®
Cloud Infrastructure

cloud.oracle.com/iaas

cloud.oracle.com/tryit