



ORACLE

# OCI Storage Services

Rohit Rahi

Oracle Cloud Infrastructure

Feb 2020

## 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.

# Agenda



Block Volume  
Local NVMe  
File Storage  
Object Storage  
Archive Storage

# Storage Requirements

Persistent v/s non-persistent?

What type of data?

Database, videos, audio, photos, text

Performance?

Max capacity, IOPS, throughput

Durability?

# of copies of data

Connectivity?

Local v/s network, how does app access the data

Protocol

Block v/s File v/s HTTPs

# OCI Storage Services



Block  
Volume



Local  
NVMe



File  
Storage




Object  
Storage



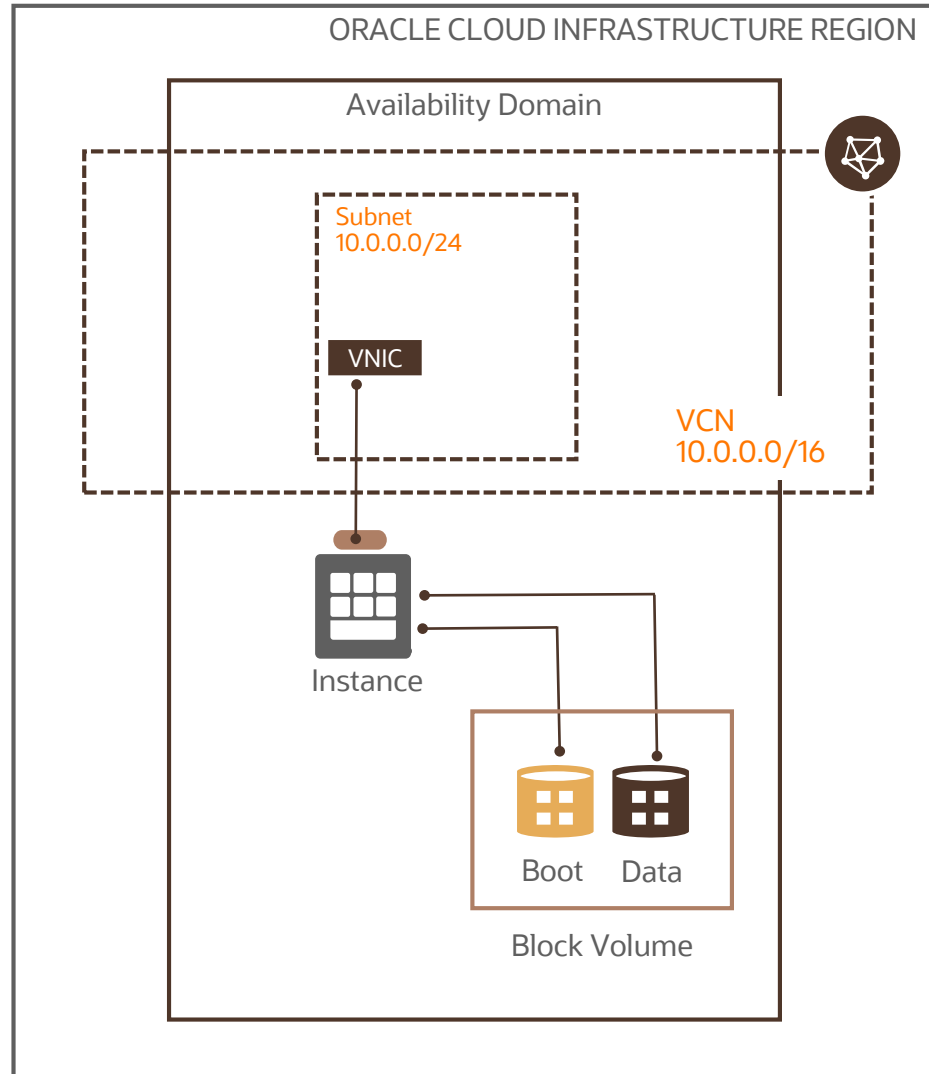
Archive  
Storage

# What is Block Storage?

- Hard drive in a server except the hard drive happens to be installed in a remote chassis
- Data is typically stored on device in fixed sized blocks (e.g. 512 Bytes)
- Accessed by operating system as mounted drive volume
- Applications/file systems decide how blocks are combined and accessed
- Data is stored without any higher-level metadata e.g. for data format, type or ownership
- You can place any kind of file system on block level storage. E.g., Windows uses NTFS; VMware uses VMFS
- Commonly deployed in Storage Area Network (SAN) storage



# Block Volume Service

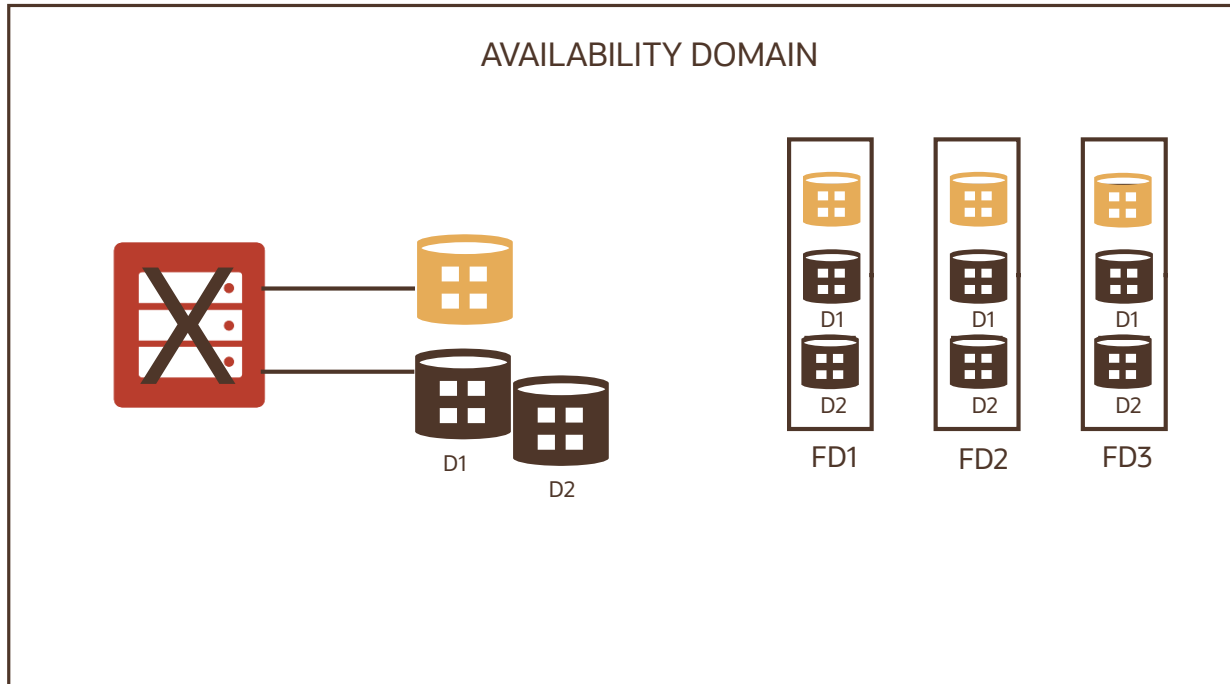


- Storage for compute instances
- 2 types: Boot Volume (OS disk), Block Volume (data disks)
- Service lets you store data independently and beyond the lifespan of compute instances

## Use cases:

- Databases
- Exchange (supports block level storage only)
- VMware (common to deploy VMware servers that use shared VMFS volumes on block level storage)
- Server boot (in public clouds, instances are configured to boot from block level storage)

# Block Volume – highly durable

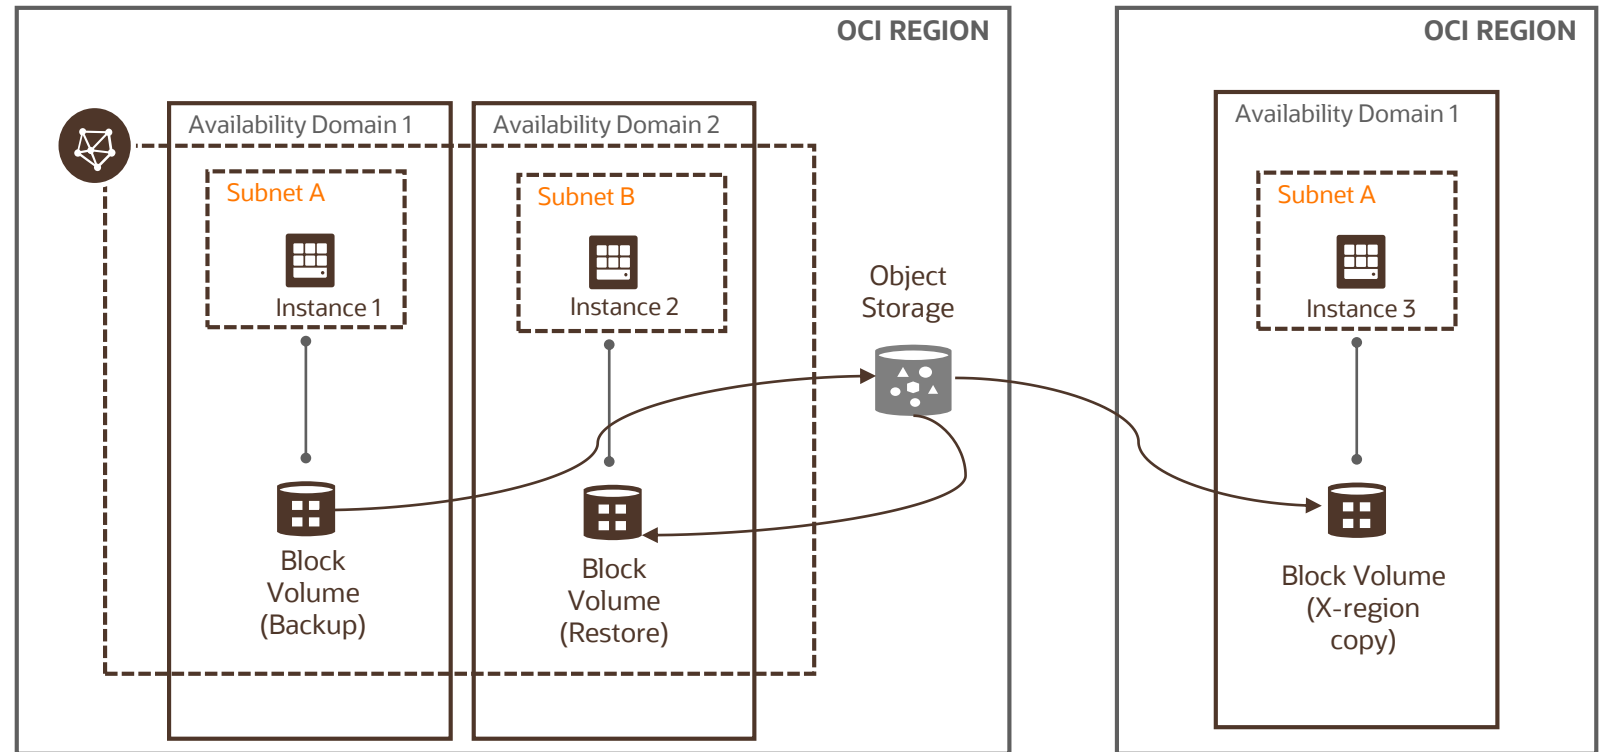


- Storage is highly durable and persistent
- Block Volume stores replica of data in 3 separate Fault Domains
- You don't need to configure any software based protection (RAID-10 etc.)
- To minimize loss of data due to deletes or corruption, we recommend to take periodic backups of block volumes. OCI allows automated scheduled backups

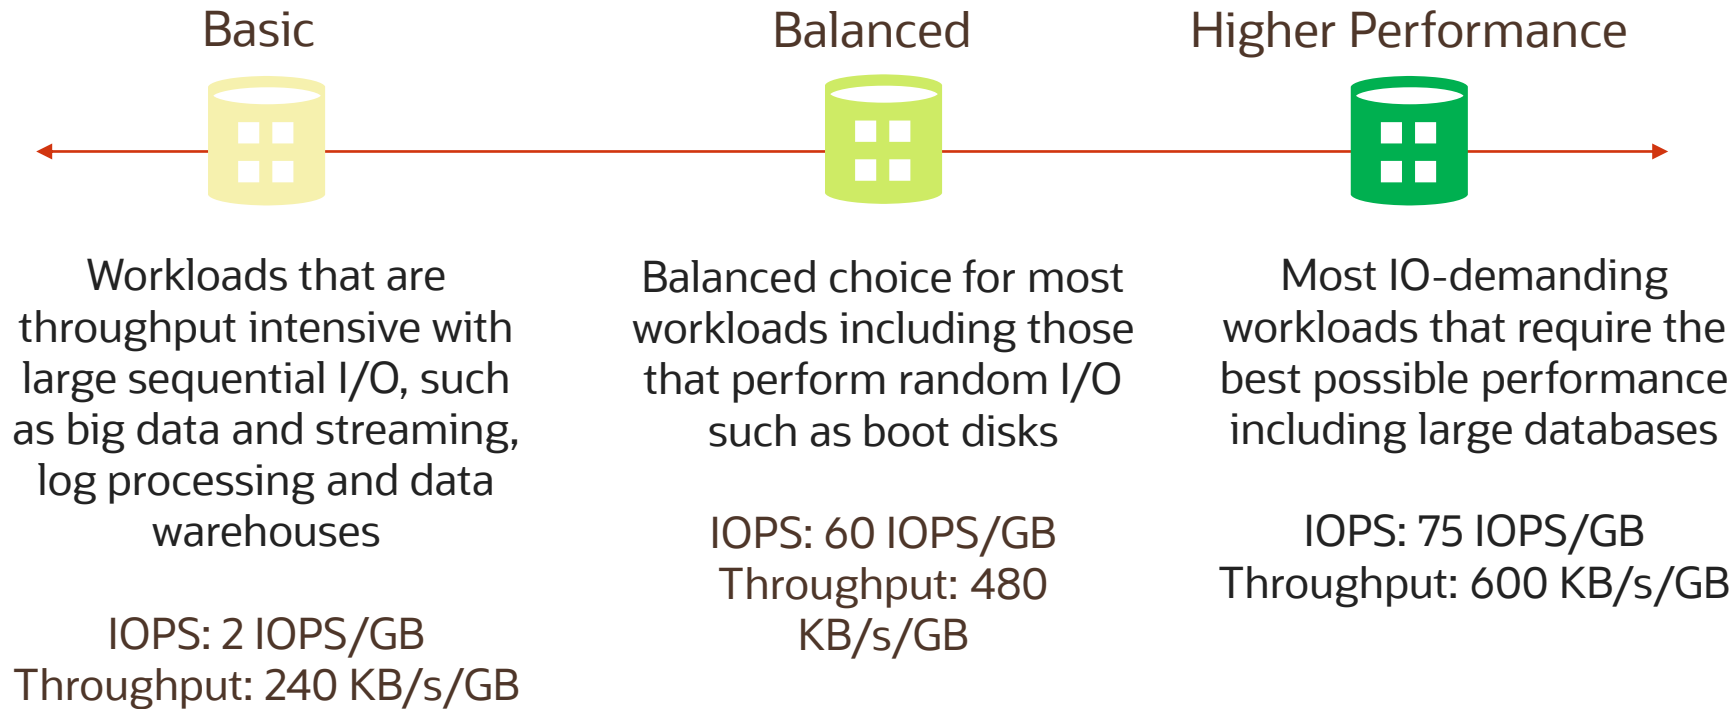


# Block Volume Backup

- Complete point-in-time snapshot copy of your block volumes
- Encrypted and stored in the Object Storage, and can be restored as new volumes to any AD within the same region (for multi-AD regions)
- Can copy block volume backups from one-region to another
- Backups can be scheduled



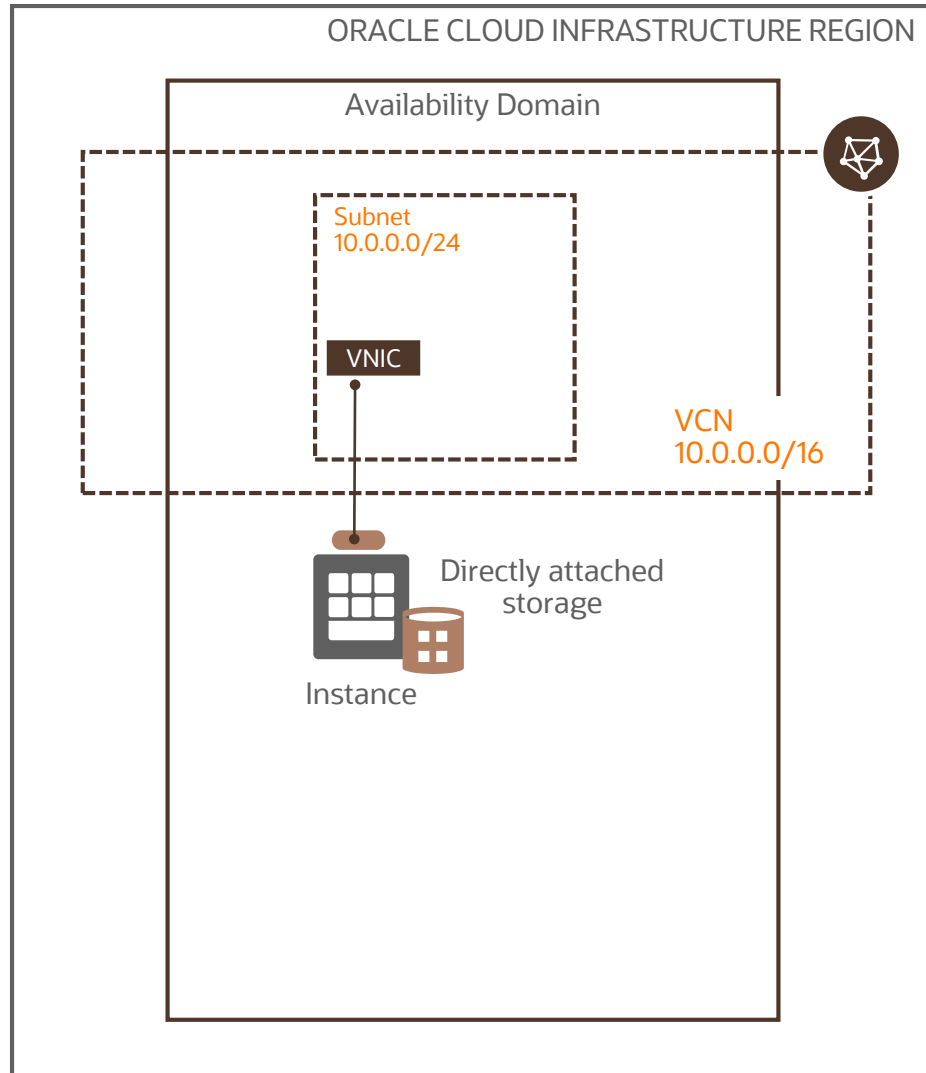
# Block Volume Tiers



Volumes can be 50 GB - 32 TB in size; you can attach up to 32 volumes/instance.

Data encrypted at rest and in-transit (oracle managed) or customer managed keys

# Local NVMe



- Temporary NVMe based storage locally attached to the compute instances
- Designed for applications that require high-performance local storage
- Use cases:
  - NoSQL databases (e.g. Cassandra, MongoDB, Redis),
  - in-memory databases,
  - Scale-out transactional databases,
  - Data warehousing
- Storage is non-persistent (survives reboot)



# Local NVMe



OCI uses NVMe (Non-Volatile Memory Express) interface for very high performance

OCI provides no RAID, snapshots, backups capabilities for these devices and customers are responsible for data durability

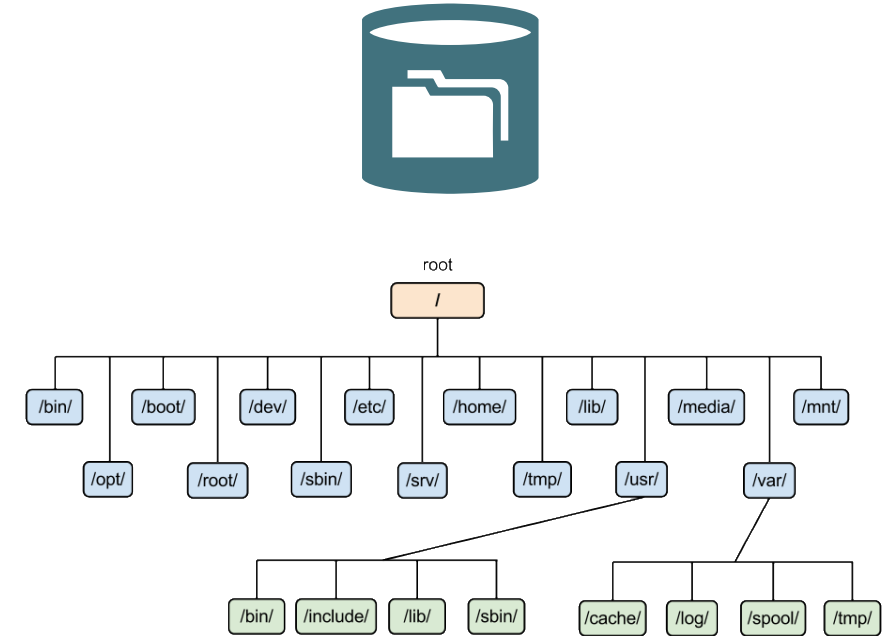
Block based Protocol (like Block Volume)

SLA around Performance

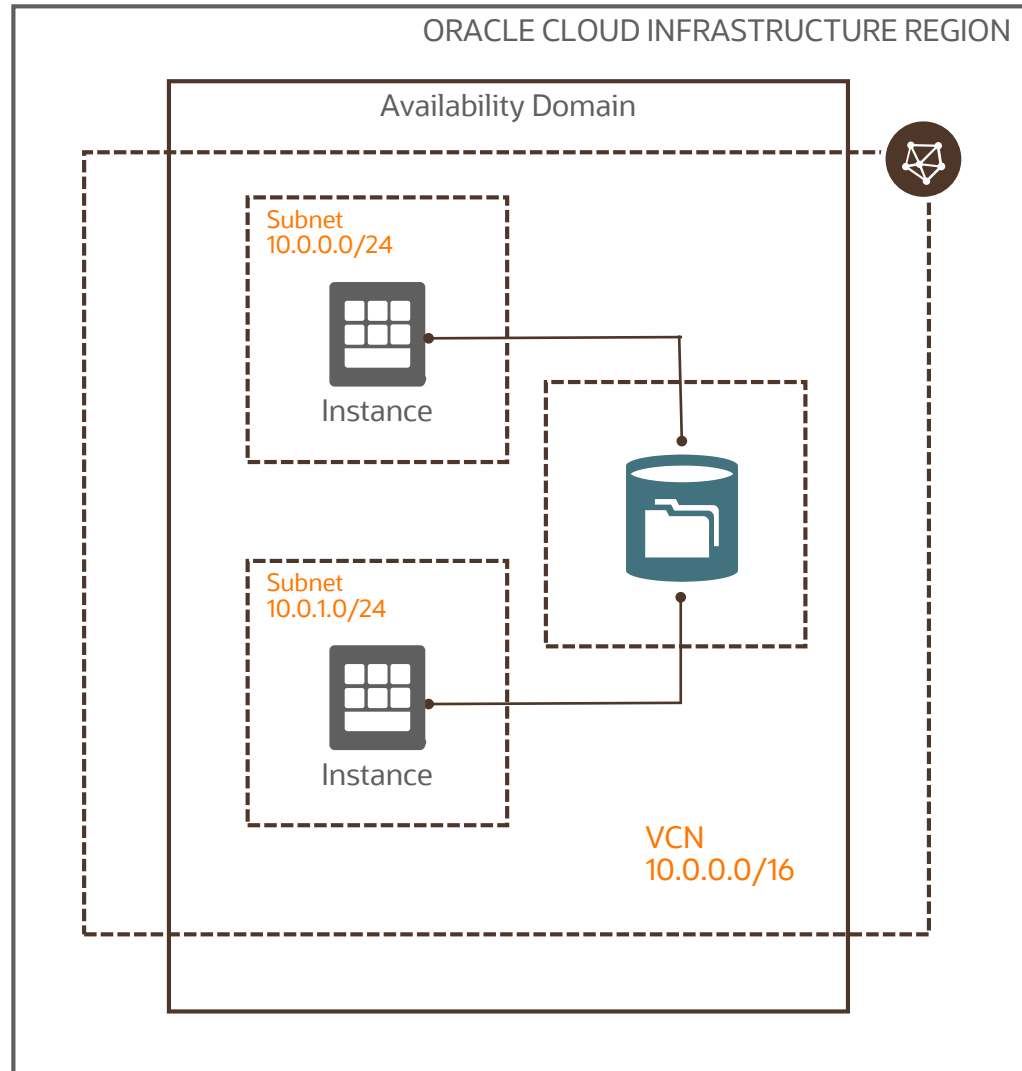
Instance type	NVMe SSD Devices	IOPS
BM.DenseIO2.52	8 drives = 51.2 TB raw	250k
VM.DenseIO2.8	2 drive = 6.4 TB raw	400k
VM.DenseIO2.16	4 drives = 12.8 TB raw	800k
VM.DenseIO2.24	8 drives = 25.6 TB raw	3.0MM

# What is File Storage?

- Hierarchical collection of documents organized into named directories which are themselves structured files
- Distributed file systems make distributed look exactly like local file systems
- Distributed file standards – NFS and SMB
  - Supported by Unix and Windows
  - Allow creation, deletion, reading, writing, sharing and locking
  - Supported by all major OSes and hypervisors
  - (typically) no extra client software needed
  - Provide access over networks

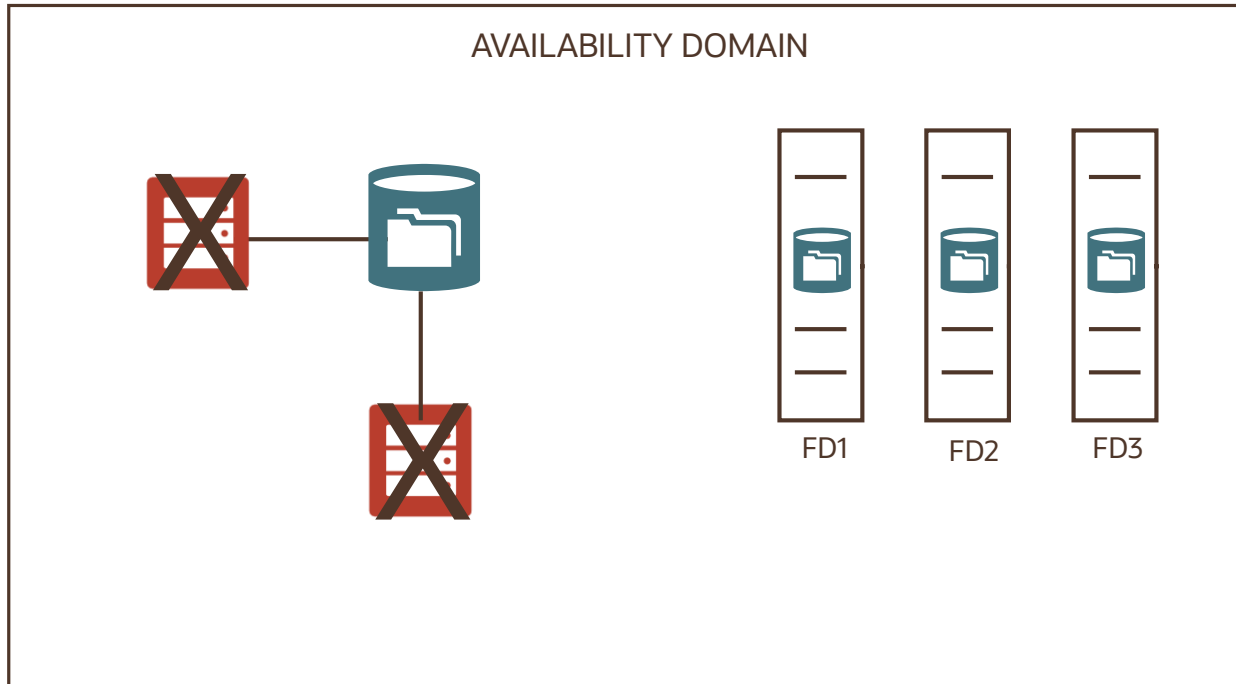


# File Storage Service (FSS)



- Shared file system storage for compute instances
- Supports NFS v.3 distributed file system
- Data protection: Snapshots (10,000 snapshots per file system)
- Security: data-at-rest and in-transit encryption for all file systems & metadata
- Use cases:
  - Oracle Applications (e.g. EBS)
  - HPC
  - Big Data and Analytics
  - General purpose File systems

# File Storage – highly durable



Storage is highly durable and persistent

File Storage stores replica of data in 3 separate Fault Domains

You can take snapshots of file systems that provide a read-only, space efficient, point-in-time backup of a file system

You can restore a file within the snapshot, or an entire snapshot using the `cp` or `rsync` command

# What is Object Storage?

- All data, regardless of content type, is managed as objects
- Each object is stored in a bucket. A bucket is a logical container for storing objects
- Objects are stored in a single, flat structure without a folder hierarchy. This means that accessing individual objects is fast and easy
- Each Object is composed of object itself and metadata of the object. This makes it easier to index and access data
- Object storage is quite common in cloud-based storage scenarios with very high scalability and reliability
- While files and blocks are generally available to an operating system (by mount operation), object storage relies on standard HTTP verbs





# Object Storage Service



An internet-scale, high-performance storage platform

Ideal for storing unlimited amount of unstructured data (images, media files, logs, backups)

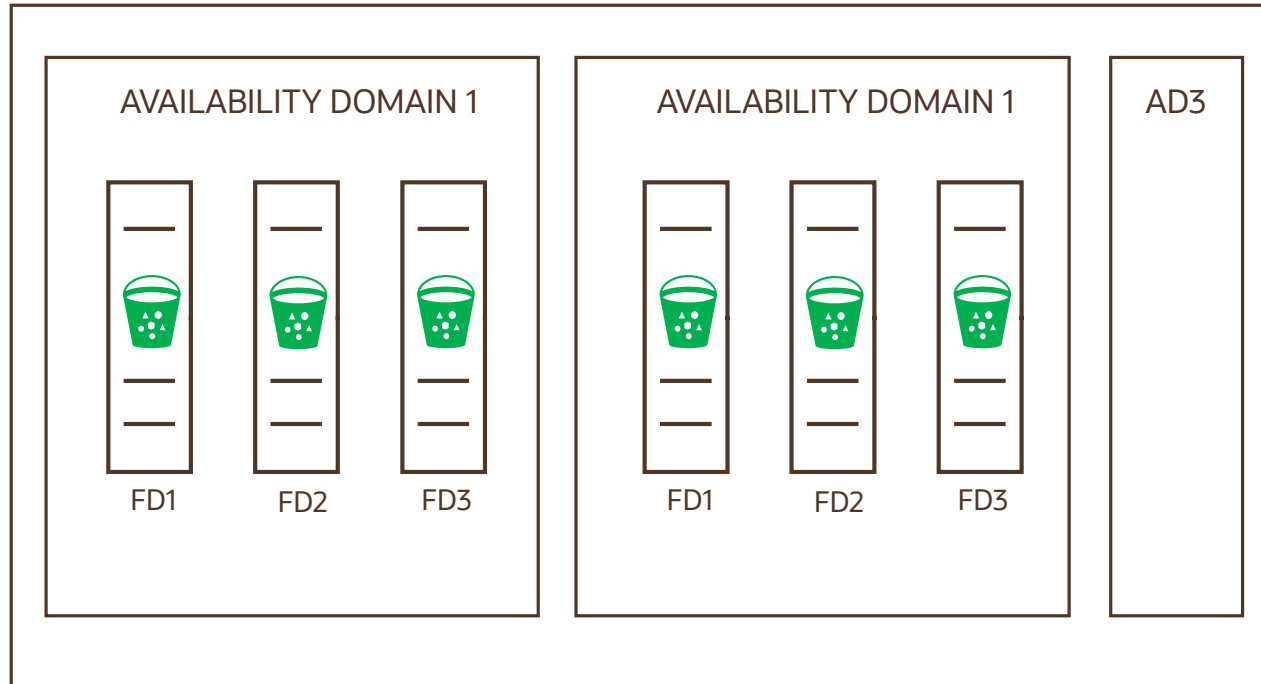
Regional service, not tied to any specific compute instance

Offers two distinct storage classes "hot" storage (standard), "cold" storage (Archive)

Use cases

- Content repository for data, images, logs, and video etc.
- Archive/Backup for longer periods of time
- Storing log data for analysis and debugs/troubleshooting
- Storing large data sets (genome data, IoT)
- Big Data/Hadoop storage

# Object Storage – highly durable



Storage is highly durable and persistent

Object storage stores replica of data in 3 separate Fault Domains in an AD

In a multi-AD region, it stores replica of data in more than one AD

Data integrity is actively monitored and corrupt data detected and auto repaired

You can leverage cross-region copy for disaster recovery scenarios

# Object Storage Tiers

## Standard Storage Tier (Hot)

- Fast, immediate, and frequent access
- Data retrieval is instantaneous
- Always serves the most recent copy of the data when retrieved
- Standard buckets can't be downgraded to archive storage

## Archive Storage Tier (Cold)

- Seldom or rarely accessed data but must be retained and preserved for long periods of time
- 10X cheaper than Standard Tier ( \$0.0026 v/s \$0.0255 Gb/month)
- 90 days minimum retention requirement
- Objects need to be restored before download; Time To First Byte (TTFB) after restore request is made: 4 Hours
- Archive Bucket can't be upgraded to Standard storage tier

Create Bucket

[help](#) [cancel](#)

Specify the storage tier for this bucket. Storage tier for a bucket can only be specified during creation.

BUCKET NAME

ObjectStorageBucketName

STORAGE TIER

☒ STANDARD

☐ ARCHIVE

Create Bucket

# OCI Storage Services

	Local NVMe	Block Volume	File Storage	Object Storage	Archive Storage
Type	NVMe SSD based temporary storage	NVMe SSD based block storage	NFSv3 compatible file system	Highly durable Object storage	Long-term archival & backup
Access	Block	Block	File	Object	Object
Structure	Block level structured	Block level structured	Hierarchical	Unstructured	Unstructured
Durability	Non-persistent; survives reboots	Durable (multiple copies in an AD)	Durable (multiple copies in an AD)	Multiple copies across ADs*	Multiple copies across ADs*
Capacity	Terabytes+	Petabytes+	Exabytes+	Petabytes+	Petabytes+
Unit Size	51.2 TB for BM, 6.4-25.6 TB for VM	50 GB to 32 TB/vol 32 vols/instance	Up to 8 Exabyte	10 TB/object	10 TB/object
Use cases	OLTP, NoSQL, Data warehousing	Database, VMFS, NTFS, boot and data disks for instances	Oracle apps (EBS), HPC, general purpose file systems	Unstructured data incl. logs, images, videos	Backups and long term archival (DB backups)

\* in multi-AD regions

# Summary



Block Volume  
Local NVMe  
File Storage  
Object Storage  
Archive Storage



### **Oracle Cloud always free tier:**

[oracle.com/cloud/free/](https://oracle.com/cloud/free/)

### **OCI training and certification:**

[cloud.oracle.com/en\\_US/iaas/training](https://cloud.oracle.com/en_US/iaas/training)

[cloud.oracle.com/en\\_US/iaas/training/certification](https://cloud.oracle.com/en_US/iaas/training/certification)

[education.oracle.com/oracle-certification-path/pFamily\\_647](https://education.oracle.com/oracle-certification-path/pFamily_647)

### **OCI hands-on labs:**

[ocitraining.clouddable.com/provider/oracle](https://ocitraining.clouddable.com/provider/oracle)

### **Oracle learning library videos on YouTube:**

[youtube.com/user/OracleLearning](https://youtube.com/user/OracleLearning)



Thank you

