

Object Storage

Level 100

Rohit Rahi

November 2018

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.

Objectives

After completing this lesson, you should be able to:

- Describe concepts, use cases and tiers of Object Storage service
- Manage buckets and objects

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 and backup
Durability	Non-persistent; survives reboots	Durable (multiple copies in an AD)	Durable (multiple copies in an AD)	Highly durable (multiple copies across ADs)	Highly durable (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	Big Data, OLTP, high performance workloads	Apps that require SAN like features (Oracle DB, VMW, Exchange)	Apps that require shared file system (EBS, HPC)	Unstructured data incl. logs, images, videos	Long term archival and backups (Oracle DB backups)

Object Storage service

- An internet-scale, high-performance storage platform
- Ideal for storing unlimited amount of unstructured data (images, media files, logs, backups)
- Data is managed as objects using an API built on standard HTTP verbs
- Regional service, not tied to any specific compute instance
- Offers two distinct storage classes to address the need for performant, frequently accessed "hot" storage, and less frequently accessed "cold" storage
- Supports private access from Oracle Cloud Infrastructure resources in a VCN through a Service Gateway
- Supports advanced features such as cross-region copy, pre-authenticated requests, lifecycle rules and multipart upload

Object Storage Scenarios

- Content Repository - highly available and durable content repository for data, images, logs, and video etc.
- Archive/Backup - use of object storage for preserving data for longer periods of time
- Log Data - application log data for analysis and debugs/troubleshooting
- Large Data Sets - Large data e.g. pharmaceutical trials data, genome data, and Internet of Things (IoT)
- Big Data/Hadoop Support
 - Use as a primary data repository for big data enables ~50% improvement in performance
 - HDFS connector provides connectivity to various big data analytic engines like Apache Spark and MapReduce

Object Storage Service Features

- Strong consistency
 - Object Storage Service always serves the most recent copy of the data when retrieved
- Durability
 - Data stored redundantly across multiple storage servers across multiple ADs
 - Data integrity is actively monitored and corrupt data detected and auto repaired
- Performance
 - Compute and the Object Storage Services are co-located on the same fast network
- Custom metadata
 - Define your own extensive metadata as key-value pairs
- Encryption
 - Employs 256-bit Advanced Encryption Standard (AES-256) to encrypt object data

Object Storage Resources

- Object
 - All data, regardless of content type, is managed as objects (e.g. logs, videos)
 - Each Object is composed of object itself and metadata of the object
- Bucket
 - A logical container for storing objects; Each object is stored in a bucket
- Namespace
 - A logical entity that serves as a top-level container for all buckets and objects
 - Each tenancy is provided one unique namespace that is global, spanning all compartments and regions
 - Bucket names must be unique within your tenancy, but can be repeated across tenancies
 - Within a namespace, buckets and objects exist in flat hierarchy, but you can simulate a directory structure using prefixes and hierarchies

Object Storage Tiers

- Standard Storage Tier (Hot)
 - Fast, immediate, and frequent access
 - Object Storage Service always serves the most recent copy of the data when retrieved
 - Data retrieval is instantaneous
 - 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
 - Minimum retention requirement for Archive Storage is 90 days
 - Objects need to be restored before download
 - Archive Bucket can't be upgraded to Standard storage tier
 - Time To First Byte (TTFB) after Archive Storage restore request is made: 4 Hours

Create Bucket [help](#) [cancel](#)

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

BUCKET NAME

STORAGE TIER

STANDARD

ARCHIVE

Create Bucket

Object Naming

- Service prepends the Object Storage namespace string and bucket name to object name, `/n/<object_storage_namespace>/b/<bucket>/o/<object_name>`
 - `https://objectstorage.us-phoenix-1.oraclecloud.com/n/gse00014346/b/DatabaseBackup/o/database1.dbf`
- Flat hierarchy
- For large number of objects, use prefixes and hierarchies,
 - `n/ansh8tvru7zp/b/event_photos/o/marathon/finish_line.jpg`
 - `/n/ansh8tvru7zp/b/event_photos/o/marathon/participants/p_21.jpg`
- You can use the CLI to perform bulk downloads and bulk deletes of all objects at a specified level of the hierarchy, without affecting objects in levels above or below
- E.g. above, you can use CLI to download or delete all objects at the `marathon/` level without downloading or deleting objects at the `marathon/participants` sublevel

Cross-region Copy

Cross Region Copy

DESTINATION NAMESPACE ⓘ

intoraclerohit

DESTINATION REGION

Please select destination region ...

DESTINATION BUCKET

OVERWRITE RULE

ETag matching rules allow you to control the copying or overwriting of objects based on their ETag values

[Learn more about ETag matching in cross region copy.](#)

OVERWRITE DESTINATION OBJECT

DO NOT OVERWRITE ANY DESTINATION OBJECT

OVERWRITE DESTINATION OBJECT ONLY IF IT MATCHES THE SPECIFIED ETAG

COPY OBJECT ONLY IF THE SOURCE MATCHES THE SPECIFIED ETAG

ETAG MATCH

- Copy objects to other buckets in the same region and to buckets in other regions
- Must authorize the service to manage objects on your behalf (separate policy for each region). E.g.
 - [allow service objectstorage-us-ashburn-1 to manage object-family in tenancy](#)
- Must specify an existing target bucket
- Bulk copying is not supported
- Objects cannot be copied from Archive storage

Object Lifecycle Management

- Define lifecycle rules to automatically archive or delete objects after a specified number of days
- Must authorize the service to manage objects on your behalf (separate policy/region). E.g.
 - [allow service objectstorage-us-ashburn-1 to manage object-family in tenancy](#)
- Applied at the bucket or object name prefix level. If no prefix is specified, the rule will apply to all objects in the bucket
- A rule that deletes an object always takes priority over a rule that would archive that same object
- Enable or disable a rule to make it active or inactive

For objects, [/n/ansh8tvru7zp/b/apparel/o/gloves_27_dark_green.jpg](#), [/n/ansh8tvru7zp/b/apparel/o/gloves_27_light_blue.jpg](#), gloves_27 is the prefix

Create Lifecycle Rule

Create a rule to automatically archive or delete objects after a specified number of days. prefix is specified, the rule will apply to all objects in the bucket.

NAME

The name of the lifecycle rule.

LIFECYCLE ACTION

Objects deleted by a lifecycle policy cannot be restored.

NUMBER OF DAYS

Delete objects that are older than 30 days and that match any of the specified object name prefixes.

Object Name Prefix

PREFIX

STATE

ENABLED DISABLED

Your rule will be active upon creation.

Managing Access and Authentications

Pre-Authenticated Requests

- Provides a way to let users access a bucket or an object without having their own credentials
- Can access via a unique URL. E.g. <https://objectstorage.us-ashburn-1.oraclecloud.com/p/p09Nx-f4UaLCN-MM0xGQIpobmMchgHQRsQv4Lr-aSzs/n/intoraclerohit/b/Image/o/kvm>
- Can revoke the links any time (much easier than S3)

Public Buckets

- At creation, a bucket is considered private and access to the bucket requires authentication and authorization
- Service supports anonymous, unauthenticated access to a bucket by making a bucket public (read access to the bucket)
- Changing the type of access doesn't affect existing pre-authenticated requests. Existing pre-authenticated requests still work

Create Pre-Authenticated Request

NAME

PRE-AUTHENTICATED REQUEST TARGET

BUCKET

OBJECT

ACCESS TYPE

PERMIT READS ON THE OBJECT

PERMIT WRITES TO THE OBJECT

PERMIT READS ON AND WRITES TO THE OBJECT

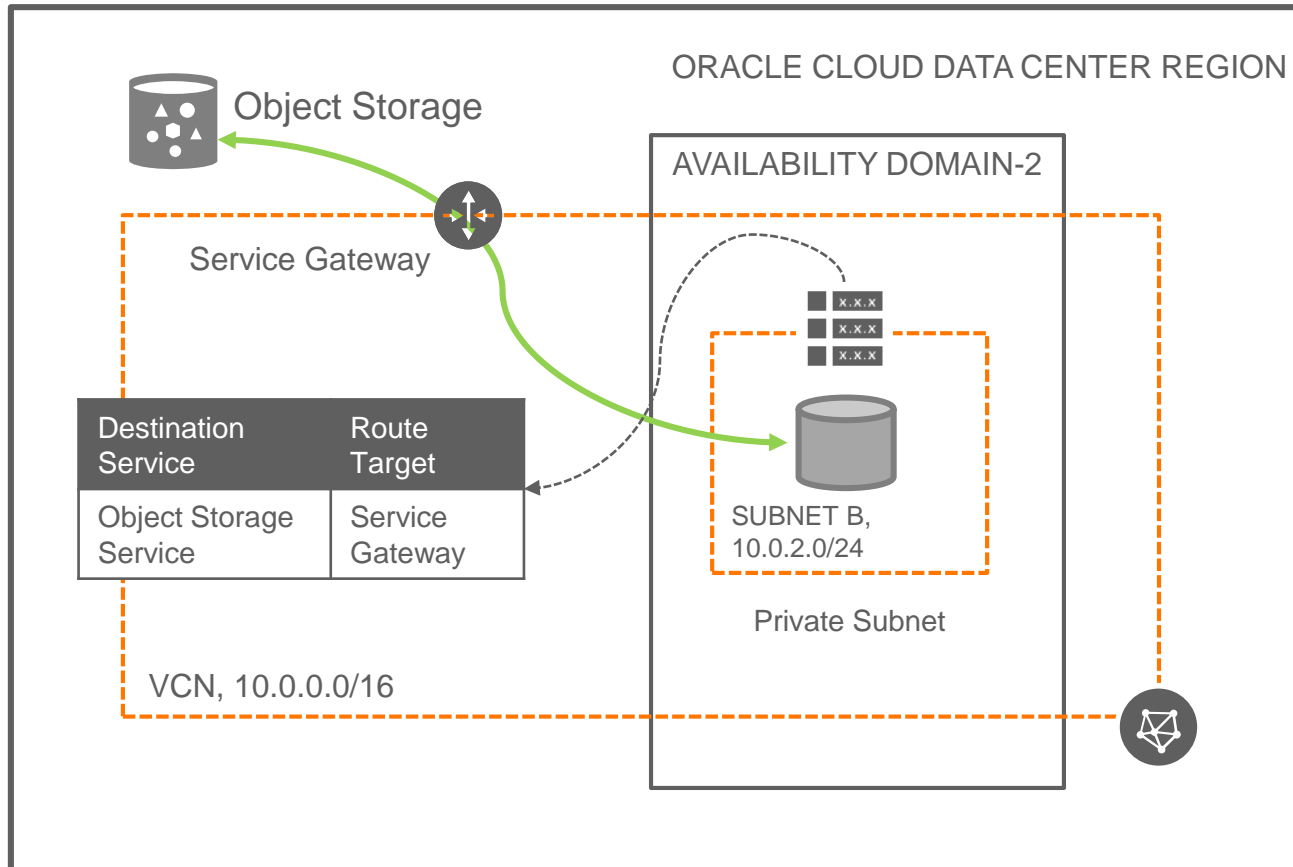
EXPIRATION DATE/TIME

[Create Pre-Authenticated Request](#)

Managing Multipart Uploads

- With multipart uploads, individual parts of an object can be uploaded in parallel to reduce the amount of time you spend uploading. Steps involved -
- Create object parts
 - Perform a multipart upload to upload objects larger than 100 MiB. Individual parts can be as large as 50 GiB or as small as 10 MB
 - Assign part numbers from 1 to 10,000
- Initiate an upload
 - Initiate a multipart upload by making a CreateMultipartUpload REST API call
- Upload object parts
 - Make an UploadPart request for each object part upload
 - If you have network issues, you can restart a failed upload for an individual part. You do not need to restart the entire upload
- Commit the upload
 - When you have uploaded all object parts, complete the multipart upload by committing it; add a bullet on checksum etc.

Service Gateway



Service gateway lets resources in VCN access public OCI services such as Object Storage, but without using an internet or NAT gateway

Any traffic from VCN that is destined for one of the supported OCI public services uses the instance's private IP address for routing, travels over OCI network fabric, and never traverses the internet

Use case: back up DB Systems in VCN to Object Storage)

Object Storage Demo

Pricing

Oracle Object Storage is charged on the following usage elements

- Storage used per month, measured in Storage-Byte Hours, aggregated per month
- Total number of requests received per month. Delete requests are excluded

	Metric	Pay as You Go	Monthly Flex
OCI Object Storage Requests	10,000 requests/month	\$0.034	\$0.034
OCI Object Storage	GB/month	\$0.0255	\$0.0255
OCI Archive Storage	GB/month	\$0.0026	\$0.0026

- Outbound Internet Transfer

	Metric	Pay as You Go	Monthly Flex
Outbound Data Transfer - First 10 TB / Month	GB/month	Free	Free
Outbound Data Transfer - Over 10 TB / Month	GB/Month	\$0.0085	\$0.0085
Inbound Data Transfer	GB/Month	Free	Free

Summary

- An internet-scale, high-performance storage platform
- Regional service, not tied to any specific compute instance
- Offers two distinct storage classes to address the need for performant, frequently accessed "hot" storage, and less frequently accessed "cold" storage
- Supports private access from Oracle Cloud Infrastructure resources in a VCN through a Service Gateway
- Supports advanced features such as cross-region copy, life cycle management, pre-authenticated requests and multipart uploads

ORACLE[®]
Cloud Infrastructure

cloud.oracle.com/iaas

cloud.oracle.com/tryit