

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, timing, and pricing of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation.



ORACLE

Data Migration

Data Transfer Service and Storage Gateway – L100

Flavio Pereira

Oracle Cloud Infrastructure

October 2019

Objectives

After completing this lesson, you should be able to:

- Plan your data migration to OCI
- Understand the difference between Online and Offline Transport
- Understand how Data Transfer Service works
- Explain Storage Gateway capabilities and use cases

Planning your Data Migration to OCI

The most successful data migrations begin with a detailed data inventory and assessment of on-premises IT resources. The assessment should cover:

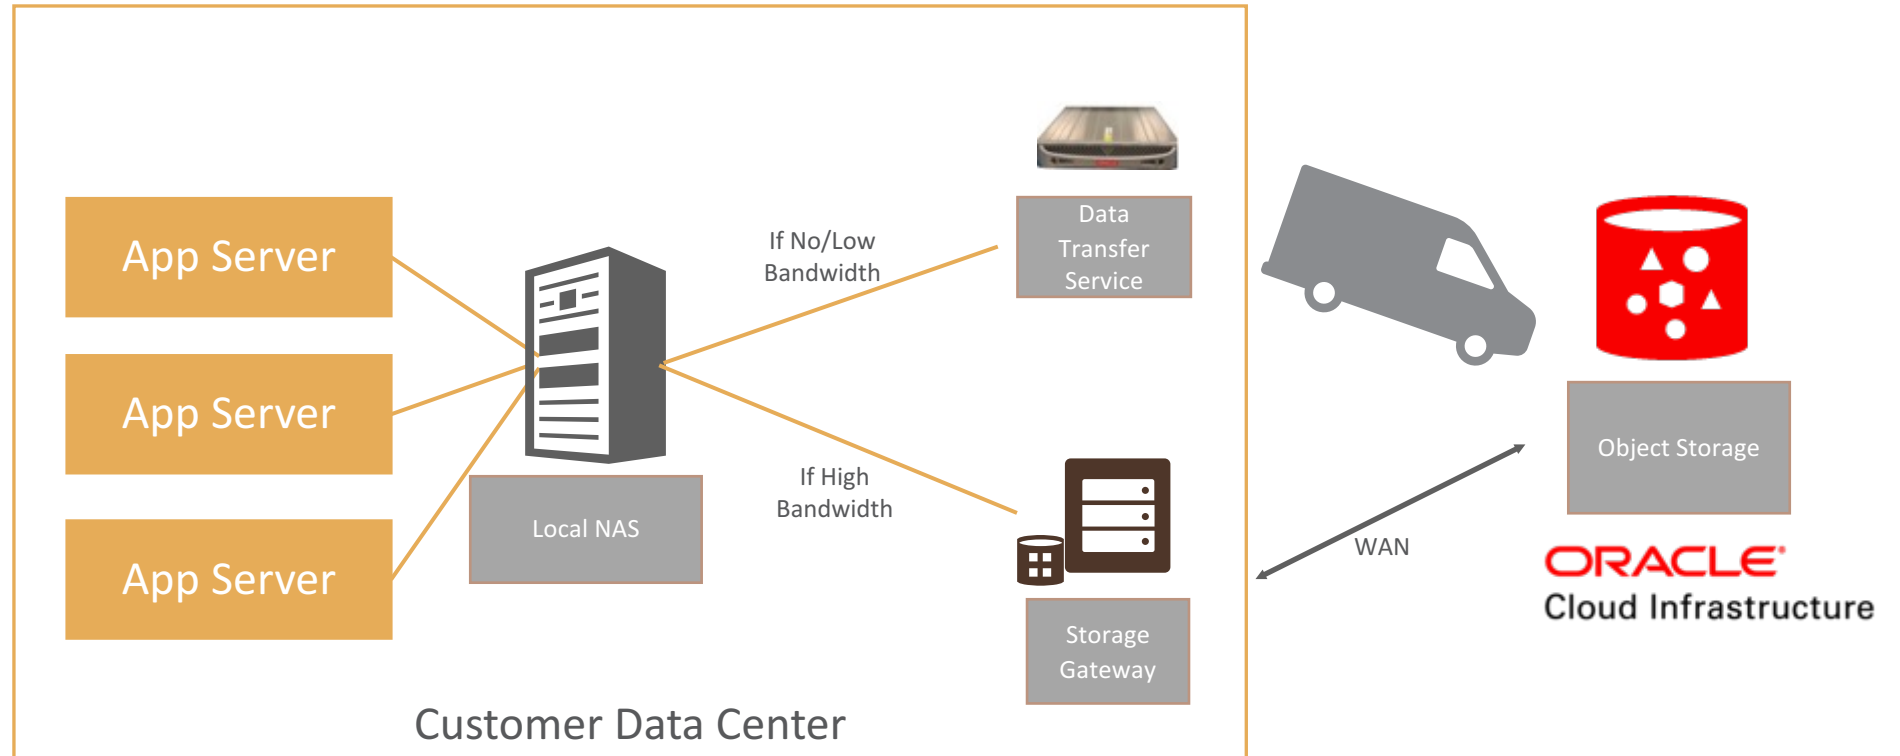

- **Applications:** Identity all software applications, their versions, and their dependencies
- **Database:** Inventory all database versions and make note of the different data types stored within them
- **Regulatory Compliance:** Applications and data types may be subject to regulatory compliance directives like HIPAA and FedRAMP
- **Storage:** Organizations may have several different types of storage in on-premises IT environments. Make note of each storage type and the amount of data housed within each one.
- **Networking:** A detailed assessment of the source environment's network architecture helps migration teams architect an optimal target network in Oracle Cloud Infrastructure.

Planning your Data Migration to OCI

Using the assessment as a guide, the next step is to create a detailed multi-phase data migration plan, with each phase focusing on the migration of specific subsets of related resources. Organizations typically break the migration process into phases based on one or more of following criteria:


- **Business Criticality:** One strategy is to migrate non-critical data first and then move on to more important business-critical data.
- **Deployment Environment Type:** For example, low-risk environments like development and testing typically go first, followed by user acceptance training, integration, and finally, high-risk production environments.
- **Disaster Recovery:** One low-risk migration strategy involves creating a complete disaster recovery environment in Oracle Cloud Infrastructure. Organizations then switch to using the Oracle Cloud Infrastructure backup deployment as the production environment, and the on-premises environment for disaster recovery.

Offline and Online Migration



Offline Transport

Data Transfer Appliance and Disk



Offline Transport

Petabyte scale offline Data Transport Service



Data Transfer Disk



SATA



USB

Customer purchases disks, loads data and sends to Data Transfer site (up to 100TB and 10 disks per Transfer Packet)

Data encrypted (AES-256) before transit

Data uploaded to bucket, verified, then erased (NIST 800-88) from disk(s) and returned to customer



Data Transfer Appliance



Customer rents one or more Data Transfer Appliance (DTA)s from Oracle (150TBs), loads data and sends to Data Transfer Site

Data encrypted (AES-256) before transit

Data uploaded, verified, erased from DTA (NIST 800-88)

Data Transfer Appliance Specifications



Appliance Capacity	150 TB
Management Interface	NFS v3.0, v4.0, v4.1
Appliance Weight	38 Pounds (17.24kg) (64 pounds (29.03kg) with shipping case)
Form Factor	2u Device. Can be standalone on a desk or on a standard rack shelf
Data Security	AES 256 Encryption
Appliance Security	<ul style="list-style-type: none">✓ Tamper resistant and Tamper evident enclosure with physical and digital controls✓ Only network power and serial ports are exposed✓ Appliance is secure wiped after each use
Network Connectivity	10 GbE RJ45 10GbE SFP+
Power	554W
Shipping Case Dimensions	11 x 25 x 28 inches (27.94 x 63.5 x 71.12 centimeters)

How is the data secured in transit?

Data fully encrypted at each stage of the process

AES 256 bit encryption

Data encrypted at rest in OCI Object Storage bucket (AES-256)

Network communication is encrypted using Transit Layer Security (TLS)



Manifest contains an index of all copied files with data integrity hashes

Appliance is tamper resistant and tamper evident

Chain of custody is maintained through the transfer process

Data Transfer Utility

The Data Transfer Utility is the command-line software for you to prepare transfer disks for your data and for shipment to Oracle. To install and use the Data Transfer Utility for either disk-based or appliance-based transfers, obtain the following:

- Required Oracle Cloud Infrastructure users and groups with the required IAM policies.
- A host machine with the following installed:
 - Oracle Linux 6 or greater, Ubuntu 14.04 or greater, or SUSE 11 or greater
 - Java 1.8 or Java 1.11
 - hdparm 9.0 or later
- Firewall access:
 - You may need to open your firewall configuration to the following IP address ranges: 140.91.0.0/16.
 - You also need to open access to the object storage IP address ranges: 134.70.0.0/17

Command Line Interface for Appliance Transfers

Use the OCI command-line software "DTS" functionality for to prepare the appliance for your data and for shipment to Oracle. To install and use the OCI Command Line Interface (CLI) for appliance-based transfers obtain the following:

- Required Oracle Cloud Infrastructure users and groups with the required IAM policies.
- A keypair used for signing API requests, with the public key uploaded to Oracle.
 - To use the CLI without a keypair, you can use token-based authentication (see online documentation)
- Host Linux machine running Python version 2.7.5 or 3.5 or later*
 - If you use the CLI Installer and do not have Python on your machine, the Installer offers to automatically install Python for you.
- For FIPS compliance see the documentation on FIPS-validated Libraries.
- Refer to <https://docs.cloud.oracle.com/iaas/Content/API/SDKDocs/cliinstall.htm> if needed.

* Note: Though OCI-CLI can run on Mac, Windows, Linux, DTS operations through OCI-CLI is ONLY supported on Linux. That is because CLI performs complex local operations and those are validated ONLY on Linux.



How Data Transfer Works



Step 1 |

- ✓ Create a Transfer Job (Console or CLI)



Step 2 |

- ✓ Request one or more Appliances



Step 4 |

- ✓ Retrieve encryption key using the data transfer utility/CLI
- ✓ Prepare the appliance
- ✓ Create Datasets (mount point)



Step 4 |

- ✓ 'Seal' dataset to generate manifest file
- ✓ 'Finalize' the Transfer Appliance



Step 3 |

- ✓ Connect appliance to the network
- ✓ Connect to the appliance via serial console and assign IP



Step 6 |

- ✓ Ship the Transfer Appliance back to Oracle

Transporting VMs, Data, and Files to Oracle Cloud

The table below provides information about how long it takes to migrate datasets online and offline. The table accounts for dataset sizes and available bandwidth.

⌵


Approximate Data Upload Time					
Dataset Size	10 Mbps	100 Mbps	1 Gbps (FastConnect)	10 Gbps (FastConnect)	Using Data Transfer Service
10 TB	92 Days	9 Days	22 Hrs	2 Hrs	1 week
100 TB	1,018 Days	101 Days	10 Days	24 Hrs	1 week
500 TB	5,092 Days	509 Days	50 Days	5 Days	1 week
1 PB	10,185 Days	1,018 Days	101 Days	10 Days	2 weeks



Sweet Spot for Offline Transport (Data Transfer Appliance/Disk)

Online Transport

Storage Gateway



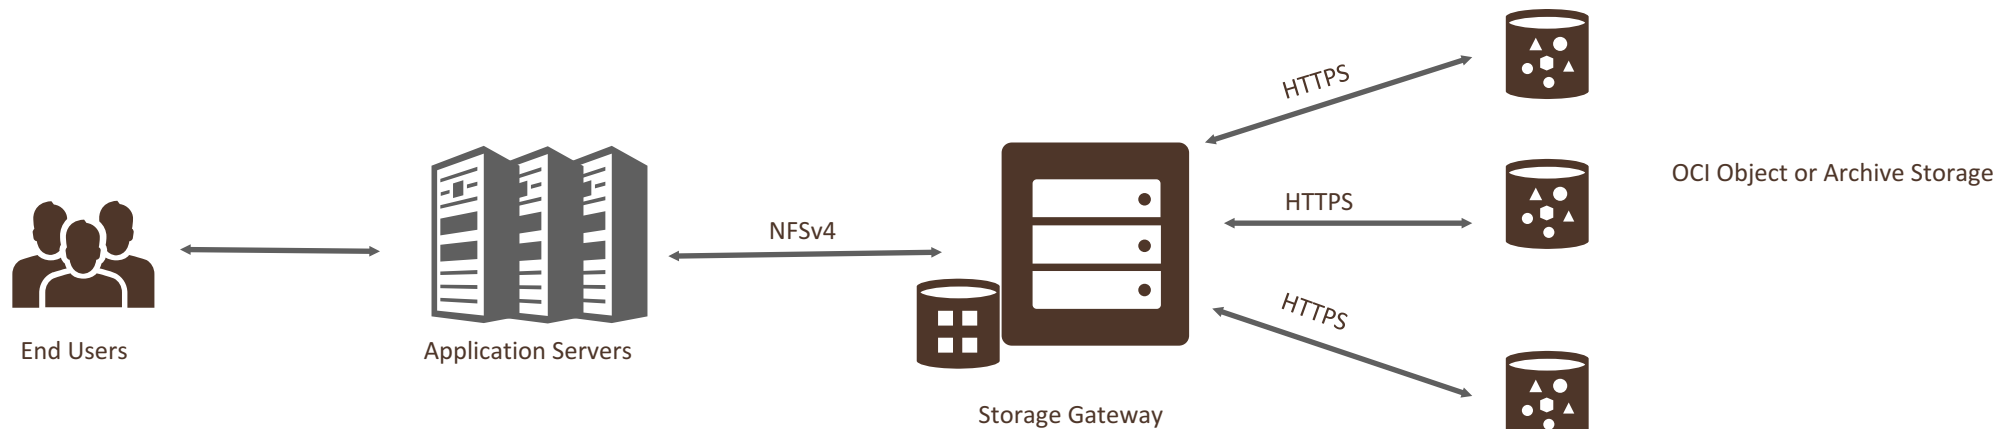
Online Transport

It's important to consider bandwidth and security when transporting data over the wire. Data should always be encrypted at rest and in transit:

- **VPN over Internet** : Relatively small datasets—up to approximately 2 terabytes (TBs)—can typically be transported over the public internet without problems
- **FastConnect** : It's the right choice for organizations that need to transport large datasets
- **Storage Gateway**: Once a secure connection has been established, organizations can use the Oracle Cloud Infrastructure Storage Gateway to securely create copies of on-premises files and place them into Oracle object storage without the need to modify applications

Storage Gateway Service

- Storage Gateway is installed as a Linux Docker instance on one or more hosts in your on-premises data center.
- Storage Gateway exposes an NFS mount point that can be mounted to any host that supports an NFSv4 client. The Storage Gateway mount point maps to an Object Storage bucket.





Storage Gateway – Two Main Use Case Categories:

Hybrid cloud: on-premises applications actively use cloud storage content

- Use cloud storage and archive as a low cost, high durability data tier
- Create a permanent data archive in cloud
- Extend on-premises data center to the cloud with limitless back-end storage
- Enhance disaster recovery and business continuity using remote cloud resources

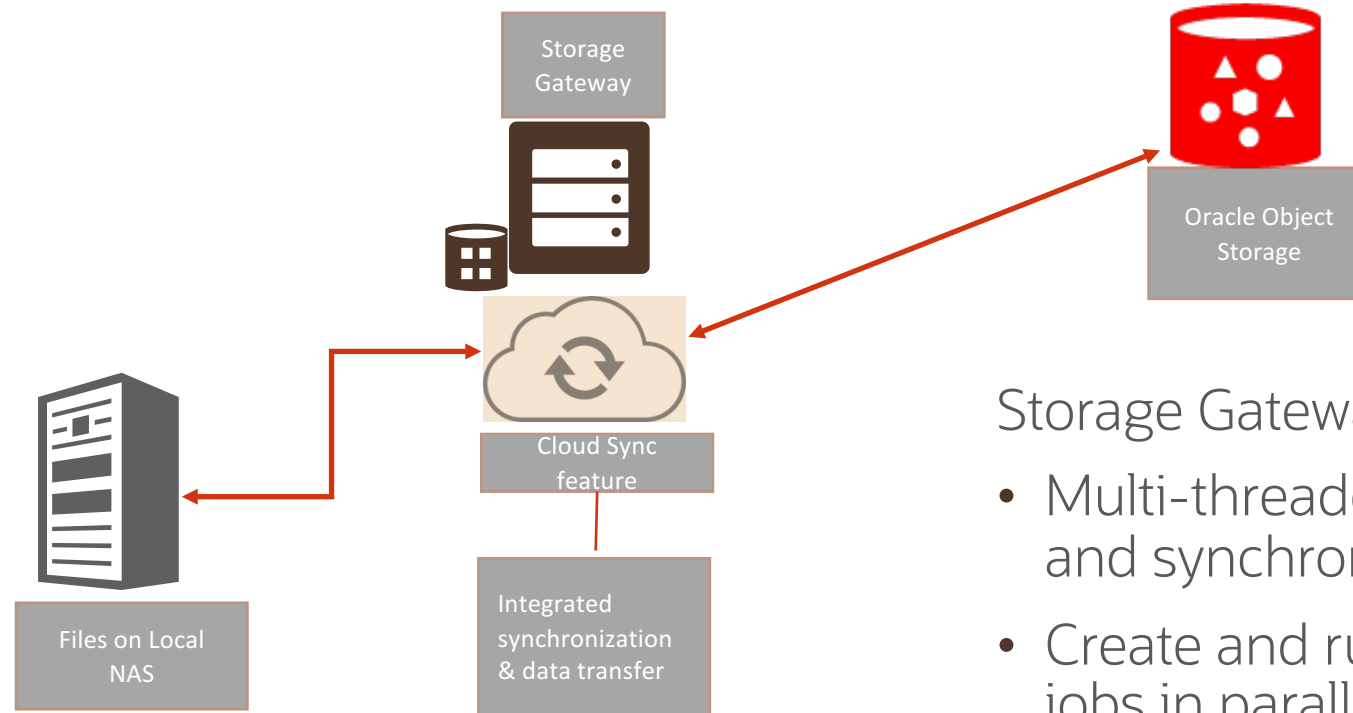
One-time data migration or periodic transfers

- Move data to cloud for app migration or adjacent analysis
- Copy data to cloud as it's written
- Move existing bulk data one time



Hybrid Cloud – Tiering, using Cloud Sync Feature

Move on-premises data sets from local NAS, via the gateway, into the cloud



Storage Gateway Cloud Sync:

- Multi-threaded data transfer and synchronization
- Create and run multiple sync jobs in parallel
- Reports upload status

Hybrid Cloud – Nearline Content Repositories

Offsite protection and data distribution

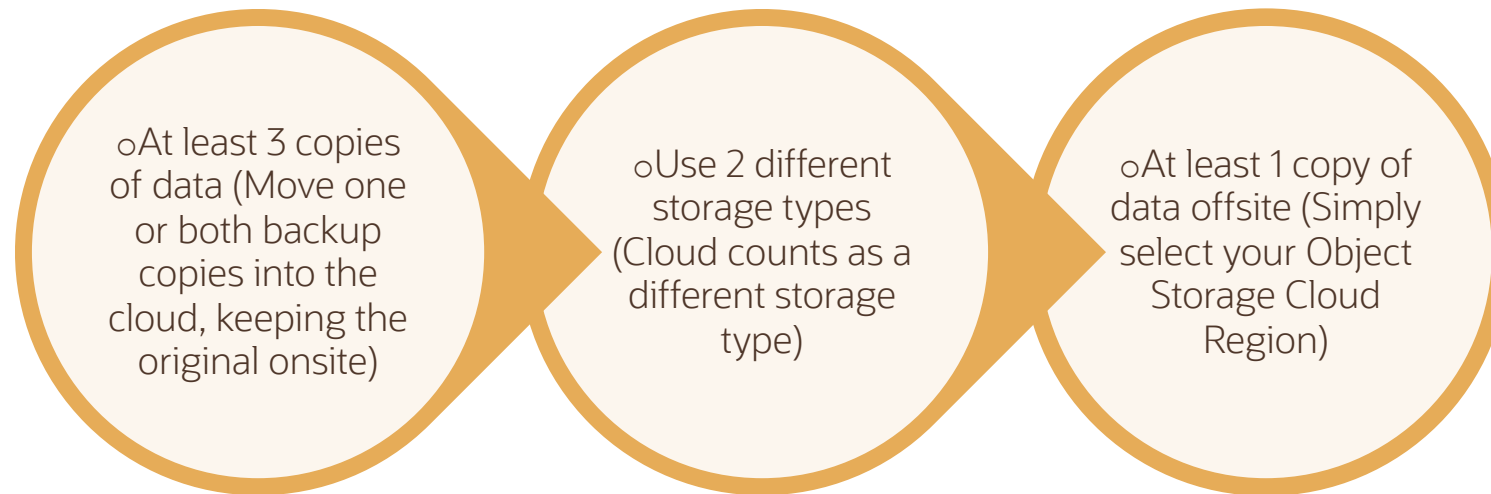


- Data is more readily shareable for re-use
- Provide access to a bucket or object with pre-authenticated request (unique URLs)
- Use object lifecycle policies between standard and archive tiers for aged out content

Hybrid Cloud – Backup Copies





The 3-2-1 best practice rule for backup and recovery



Hybrid Cloud – Cold Disaster Recovery

Resume access to data in the cloud



- Offsite data protection with geographic separation
- In DR scenario 2nd Storage Gateway claims the bucket (file system) and accesses the data

Options for Tuning & Pinning Frequently Used Data

Storage Gateway **caches frequently retrieved data** on the local host, minimizing the number of REST API calls to Oracle Cloud Infrastructure Object Storage and enabling faster data retrieval. You can configure the cache for a file system when you create the file system.

- File data is first written to local storage accessed by Storage Gateway
- Then asynchronously uploaded to the cloud bucket (data tiering)
- Frequently used data stays local for low latency read operations, files age out over time
- “pin” known hot files to local storage to maintain quick access
- Files are always uploaded to cloud storage, even when pinned to local storage

Storage Gateway Performance

A single Storage Gateway instance is able to utilize most of the available network available:

- Sustained upload speeds: 450–500 MB/s
 - using 4–8 files ranging in size from 10 GB to 50 GB over 10Gb/s FastConnect
- Sustained download speeds : 700–750 MB/s
 - using 3 files ranging in size from 50 GB to 200 GB



Limits on Storage Gateway Resources

- Ensure that the number of file systems per Storage Gateway doesn't exceed 10.
- Ensure that the number of objects stored in a Storage Gateway file system doesn't exceed 100 million.
- Ensure that you configure adequate local storage for file system cache. Storage Gateway warns you if you have configured less than the recommended 500 GB.
- The number of files in the cache is limited to 20,000 regardless of the specified cache size in bytes.
- To improve the efficiency of file ingestion, cloud upload operations, and to reduce the number of objects in the namespace, bin-pack or zip small files before writing them to Storage Gateway.

Storage Gateway FAQ

1. Can I use Storage Gateway as a general purpose Network Attached Storage server?

No, Storage Gateway must not be used as a general-purpose network-attached storage. Its primary function is data movement to the cloud. Storage Gateway is not built to handle the transactional workload that a full-featured shared file system can handle.

2. How much does Storage Gateway cost?

Storage Gateway is a free service for all Oracle Cloud Infrastructure customers.


3. Do I need to use Object Storage REST APIs to use Storage Gateway?

No, you do not need to use any Object Storage APIs to use Storage Gateway. Storage Gateway performs the NFS to REST API translation for you.

[Click here for full Storage Gateway FAQ](#)

ORACLE

Storage Gateway Demo



Summary

- Data Transfer is the best option for large data sets over 100 TB
- You can use a combination of Online and Offline migration based your workload requirements
- You can encrypted using your own keys for Data Transfer
- Storage Gateway exposes an NFS mount point that can be mounted to any host that supports an NFSv4 client.
- Storage Gateway transfers data to Oracle Cloud Infrastructure using HTTPS, which encrypts data packets in flight between Storage Gateway and the cloud.



Oracle Cloud always free tier:

oracle.com/cloud/free/

OCI training and certification:

oracle.com/cloud/iaas/training

oracle.com/cloud/iaas/training/certification

education.oracle.com/oracle-certification-path/pFamily_647

OCI hands-on labs:

ocitraining.qcloudable.com/provider/oracle

Oracle learning library videos on YouTube:

youtube.com/user/OracleLearning

