



TR-5006: High Throughput Oracle VLDB Implementation on Google Cloud NetApp Volumes with Data Guard

NetApp database solutions

NetApp
August 05, 2025

Table of Contents

TR-5006: High Throughput Oracle VLDB Implementation on Google Cloud NetApp Volumes with Data Guard 1

 Purpose 1

 Audience 1

 Solution test and validation environment 2

 Architecture 2

 Hardware and software components 2

 Oracle VLDB Data Guard configuration with a simulated NY to LA DR setup 3

 Key factors for deployment consideration 3

 Solution deployment 4

 Prerequisites for deployment 4

 Primary Oracle VLDB configuration for Data Guard 9

 Standby Oracle VLDB configuration for Data Guard 21

 Setup Data Guard Broker and FSFO with an Observer 30

 Clone standby database for other use cases via automation 40

 Where to find additional information 40

TR-5006: High Throughput Oracle VLDB Implementation on Google Cloud NetApp Volumes with Data Guard

Allen Cao, Niyaz Mohamed, NetApp

The solution provides an overview and details for configuring high throughput Oracle Very Large Database (VLDB) on Google Cloud NetApp Volumes (GCNV) with Oracle Data Guard in Google cloud.

Purpose

High throughput and mission-critical Oracle VLDB put a heavy demand on backend database storage. To meet service level agreement (SLA), the database storage must deliver the required capacity and high input/output operations per second (IOPS) while maintaining sub milliseconds latency performance. This is particularly challenging when deploying such a database workload in the public cloud with a shared storage resources environment. Not all storage platforms are created equal. GCNV is a premium storage service offered by Google that can support mission-critical Oracle Database deployments in Google cloud requiring sustained IOPS and low-latency performance characteristics. The architecture accommodates OLTP and OLAP workloads, with configurable service tiers supporting various performance profiles. GCNV delivers blazing-fast performance with sub milliseconds latency, achieving up to 4.5 GiBps per volume throughput with mixed read/write workloads.

Leveraging the fast snapshot backup (seconds) and clone (minutes) feature of GCNV, full-size copies of the production database can be cloned from Physical Standby on the fly to serve many other use cases such as DEV, UAT, etc. You can do away with an Active Data Guard license and inefficient and complex Snapshot Standby. The cost savings can be substantial. For a nominal Oracle Data Guard setup with 64 cores CPUs on both primary and standby Oracle servers, just the Active Data Guard licensing cost saving amounts to \$1,472,000 based on the latest Oracle price list.

In this documentation, we demonstrate how to set up an Oracle VLDB with Data Guard configuration on GCNV storage with multiple NFS volumes and Oracle ASM for storage load balancing. The standby database volumes can be quickly backed up via snapshot and cloned for read/write access. NetApp Solutions Engineering team provides an automation toolkit to create and refresh clones with streamlined lifecycle management.

This solution addresses the following use cases:

- Implementation of Oracle VLDB in a Data Guard setting on GCNV storage service across Google cloud regions.
- Snapshot backup and clone the physical standby database to serve use cases such as reporting, dev, test, etc. via automation.

Audience

This solution is intended for the following people:

- A DBA who sets up Oracle VLDB with Data Guard in Google cloud for high availability, data protection, and disaster recovery.

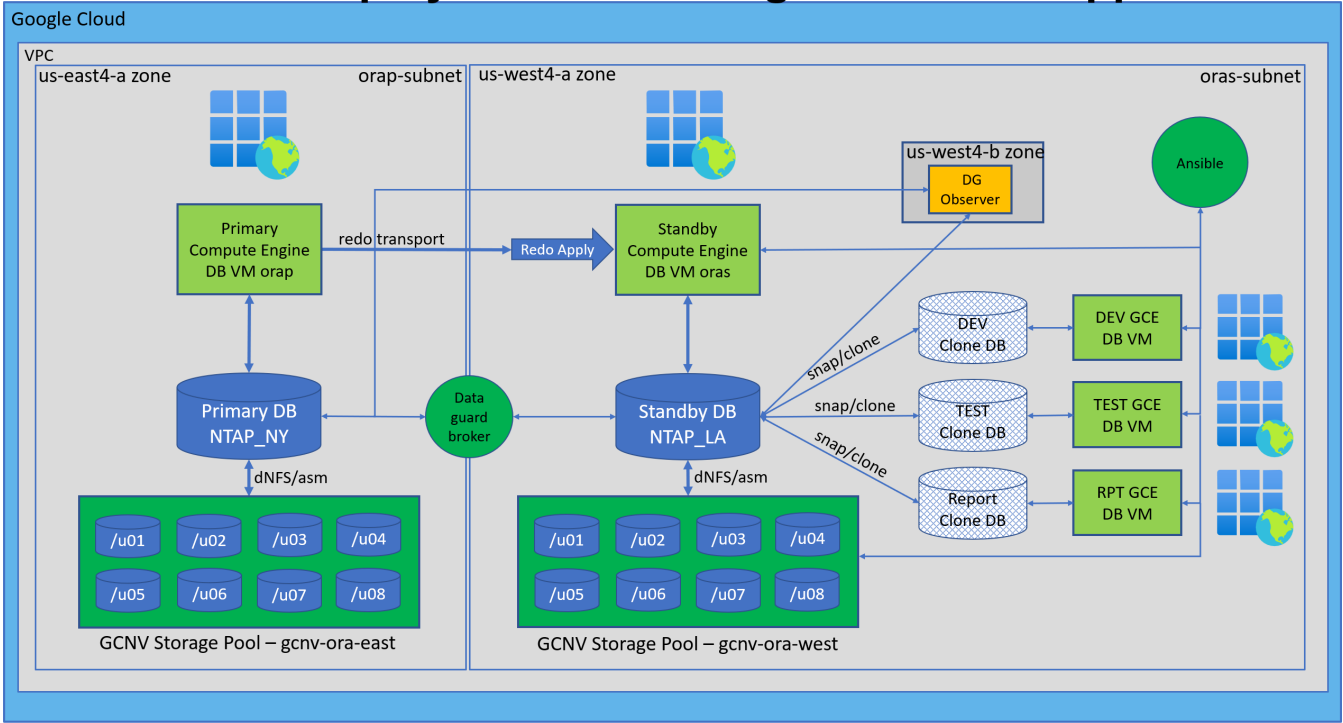
- A database solution architect interested in Oracle VLDB with Data Guard configuration in the Google cloud.
- A storage administrator who manages GCNV storage that supports Oracle database.
- An application owner who likes to stand up Oracle VLDB with Data Guard in a Google cloud environment.

Solution test and validation environment

The testing and validation of this solution was performed in a Google cloud lab setting that might not match the actual user deployment environment. For more information, see the section [Key factors for deployment consideration](#).

Architecture

Oracle VLDB Deployment with Google Cloud NetApp Volumes



Hardware and software components

Hardware		
Google Cloud NetApp Volumes	Current service offered by Google	Two Storage Pools, Premium Service Level, Auto QoS
Google Compute Engine VMs for DB Servers	N1 (4 vcpus, 15 GiB memory)	Four DB VMs, primary DB server, standby DB server, clone DB server, and Data Guard observer.
Software		
RedHat Linux	Red Hat Enterprise Linux 8.10 (Ootpa) - x86/64	RHEL Marketplace image, PAYG

Oracle Grid Infrastructure	Version 19.18	Applied RU patch p34762026_190000_Linux-x86-64.zip
Oracle Database	Version 19.18	Applied RU patch p34765931_190000_Linux-x86-64.zip
dNFS OneOff Patch	p32931941_190000_Linux-x86-64.zip	Applied to both grid and database
Oracle OPatch	Version 12.2.0.1.36	Latest patch p6880880_190000_Linux-x86-64.zip
Ansible	Version core 2.16.2	python version - 3.10.13
NFS	Version 3.0	dNFS enabled for Oracle

Oracle VLDB Data Guard configuration with a simulated NY to LA DR setup

Database	DB_UNIQUE_NAME	Oracle Net Service Name
Primary	NTAP_NY	NTAP_NY.cvs-pm-host-1p.internal
Standby	NTAP_LA	NTAP_LA.cvs-pm-host-1p.internal

Key factors for deployment consideration

- **Google Cloud NetApp Volumes Configuration.** GCNV are allocated in the Google cloud as `Storage Pools`. In these tests and validations, we deployed an 2 TiB storage pool to host Oracle primary database at the East4 region and a 4 TiB storage pool to host standby database and DB clone at the West4 region. GCNV storage pool has four service levels: Flex, Standard, Premium, and Extreme. The IO capacity of ANF capacity pool is based on the size of the capacity pool and its service level. At a capacity pool creation, you set storage pool location, service level, availability zone, and capacity of storage pool. For Oracle Data Guard configuration, Zonal availability should be sufficient as the Data Guard provides the database failover protection due to a zone level failure.
- **Sizing the Database Volumes.** For production deployment, NetApp recommends taking a full assessment of your Oracle database throughput requirement from Oracle AWR report. Take into consideration the database size, the throughput requirements, and service level when designing GCNV volumes layout for VLDB database. It is recommended to use only `Premium` or `Extreme` service for Oracle database. The bandwidth is guaranteed at 64 MiB/s per TiB volume capacity up to maximum of 4.5 GiBps for `Premium` service and 128 MiB/s per TiB volume capacity up to 4.5 GiBps for `Extreme` service. Higher throughput will need larger volume sizing to meet the requirement.
- **Multiple Volumes and Load Balancing.** A single large volume can provide similar performance level as multiple volumes with same aggregate volume size as the QoS is strictly enforced based on the volume sizing and storage pool service level. It is recommended to implement multiple volumes (multiple NFS mount points) for Oracle VLDB to better utilize shared backend GCNV storage resource pool and to meet throughput requirement exceeding 4.5 GiBps. Implement Oracle ASM for IO load balancing on multiple NFS volumes.
- **Google Compute Engine VM Instance Consideration.** In these tests and validations, we used Compute Engine VM - N1 with 4 vCPUs and 15 GiB memory. You need to choose the Compute Engine DB VM instance appropriately for Oracle VLDB with high throughput requirement. Besides the number of vCPUs

and the amount of RAM, the VM network bandwidth (ingress and egress or NIC throughput limit) can become a bottleneck before database storage throughput is reached.

- **dNFS Configuration.** By using dNFS, an Oracle database running on a Google Compute Engine VM with GCNV storage can drive significantly more I/O than the native NFS client. Ensure that Oracle dNFS patch p32931941 is applied to address potential bugs.

Solution deployment

The following section demonstrates the configuration for Oracle VLDB on GCNV in an Oracle Data Guard setting between a primary Oracle DB in Google cloud of East region with GCNV storage to a physical standby Oracle DB in Google cloud of West region with GCNV storage.

Prerequisites for deployment

Deployment requires the following prerequisites.

1. A Google cloud account has been set up and a project has been created within your Google account to deploy resources for setting up Oracle Data Guard.
2. Create a VPC and subnets that span the regions that are desired for Data Guard. For a resilient DR setup, consider to place the primary and standby DBs in different geographic locations that can tolerate major disaster in a local region.
3. From the Google cloud portal console, deploy four Google compute engine Linux VM instances, one as the primary Oracle DB server, one as the standby Oracle DB server, a clone target DB server, and an Oracle Data Guard observer. See the architecture diagram in the previous section for more details about the environment setup. Follow Google documentation [Create a Linux VM instance in Compute Engine](#) for detailed instructions.



Ensure that you have allocated at least 50G in the Azure VMs root volume in order to have sufficient space to stage Oracle installation files. Google compute engine VMs are locked down at instance level by default. To enable communication between VMs, a specific firewall rules should be created to open the TCP port traffic flow such as typical Oracle port 1521.

4. From the Google cloud portal console, deploy two GCNV storage pools to host Oracle database volumes. Referred to documentation [Create a storage pool quickstart](#) for step-by-step instructions. Following are some screen shots for quick reference.

Status	Name	Location	Service level	Capacity	Allocated to Volumes	Volume count	Created	Labels	Show more
Ready	gcnv-storage	us-east4	Premium	2048 GB	2000 GB	8	Apr 7, 2025, 11:50:52 AM	diskless-oracle	
Ready	artifici-pool	us-east5-a	Flex	1024 GB	100 GB	1	Mar 27, 2025, 2:24:04 PM	user-class:herkanto-test owner:parul	
Ready	test-region-01	us-west2	Flex	1024 GB	0 GB	0	Mar 24, 2025, 3:05:14 AM	creator:parul	
Ready	ai-data-pool	us-central1	Extreme	2048 GB	2048 GB	1	Mar 22, 2025, 10:21:12 AM	creator:parul	
Ready	ai-artifact-pool	us-central1-c	Flex	1024 GB	100 GB	1	Mar 21, 2025, 1:34:44 PM	creator:parul	
Ready	test-pool	us-central1	Premium	2048 GB	0 GB	0	Mar 18, 2025, 7:02:03 PM	owner:chiv	
Ready	ai-storage-pool	us-central1	Extreme	2048 GB	0 GB	0	Mar 18, 2025, 1:57:48 AM	owner:parul	
Ready	mike-pool-01	eu-west4-a	Flex	1024 GB	0 GB	0	Mar 12, 2025, 11:10:05 AM		
Ready	mike-storage	us-central1-f	Flex	1111 GB	0 GB	0	Mar 12, 2025, 5:11:29 AM		
Ready	myke-test	us-west4-a	Flex	1024 GB	0 GB	0	Feb 27, 2025, 12:24:41 AM	purpose:ai-artifact-test owner:parul	
Ready	gcnv-flex-pool-mountai	asia-south1-a	Flex	1024 GB	700 GB	2	Jan 15, 2025, 12:25:57 AM		
Ready	gcnv-flex-pool-bake	asia-northeast1-a	Flex	1024 GB	100 GB	1	Jan 15, 2025, 12:12:37 AM		
Ready	damian-pool	mex-central2-b	Flex	1024 GB	0 GB	0	Nov 8, 2024, 12:05:04 PM		
Ready	toronto-premium	northamerica-northeast2	Premium	2048 GB	300 GB	3	Nov 7, 2024, 8:19:50 AM	owner:team	
Ready	gcnv-data-pool	asia-south1	Flex	2048 GB	1423 GB	7	Aug 20, 2024, 4:04:52 AM	owner:parul	
Ready	prince1-govc	asia-southeast1	Premium	2048 GB	1100 GB	4	Aug 10, 2023, 6:43:49 PM	owner:kati purpose:govc	
Ready	montreal-premium	northamerica-northeast1	Premium	2048 GB	1300 GB	5	Aug 4, 2023, 8:13:52 AM	purpose:ai-artifact-general owner:team	

Google Cloud | NetApp Volumes / Storage pools / Create storage pool

Storage pools

A storage pool specifies the following for all volumes within the pool: Capacity, service level, Virtual Private Network (VPN), Active Directory policy, LDAP and encryption. [Learn more](#)

Storage pool details

Storage pool name *
gcmv-ora-west

Description (optional)

Location

Region *
us-west4 (Las Vegas)

[COMPARE THE LOCATION AND SERVICE LEVEL AVAILABILITY](#)

Service Level

Service level of the pool defines the service level of all volumes within the pool. Each volume will get an individual or shared throughput limit based on its allocated capacity and storage pool service level. [Learn more](#)

☐ Flex
Highly available general purpose storage with advanced data management. 2048–10,485,760 GB capacity. Up to 5 GB/s per storage pool with custom performance.

☐ Standard
Highly available general purpose storage with advanced data management. 2,048–10,485,760 GB capacity. Up to 15 GB/s per GB.

☒ Premium
Highly available high-performance storage with advanced data management. 2,048–10,485,760 GB capacity. Up to 64 KB/s per GB.

☐ Extreme
Highly available high-throughput storage with advanced data management. 2,048–10,485,760 GB capacity. Up to 128 KB/s per GB.

Storage pools and all volumes within the pool are highly available within the location selected.

[COMPARE THE DIFFERENT SERVICE LEVELS](#)

Capacity

Capacity *
2096 GB

Capacity must be between 2,048 GB and 10,485,760 GB in increments of 1 GB.

[COMPARE THE DIFFERENT SERVICE LEVELS](#)

Google Cloud | NetApp Volumes / Storage pools / Create storage pool

Storage pools

Set up connections

Select the network and address range that applications will use to access your volume. [Learn more](#)

Network
shared-vpc-prod

Private service access connection for network shared-vpc-prod has been successfully created. You will now be able to use the same network across all your project's managed services. If you would like to change this connection, please visit the [following page](#).

Active Directory policy

Assign an Active Directory policy to provision a storage pool for volumes using LDAP, Kerberos, or any SMB protocol variations. [Learn more](#)

☐ Assign an Active Directory policy to the storage pool. Choice is permanent once assigned.

LDAP configuration for NFS protocol(s)

☐ Enable LDAP (Available only at storage pool creation)
Enables user look up from Active Directory LDAP server for your NFS volumes. Choice is permanent.

Encryption

The following encryption policy applies for all volumes in your region. [Learn more](#)

Encryption

☒ Google-managed encryption key
No configuration required.

☐ Customer-managed encryption key (CMEK)
Manage via CMEK policies for NetApp Volumes.

Auto-tiering

Optimize storage costs by automatically moving cold data on volumes with enabled auto-tiering to the most cost-effective access tier depending on access pattern changes. [Learn more](#)

☐ Allow auto-tiering for volumes. Choice is permanent.

Labels

Label your storage pools for reports, queries.

Key 1 *
database

Value 1
oracle

[+ ADD LABEL](#)

[CREATE](#) [CANCEL](#)

Google Cloud | NetApp Volumes / Storage pools

Storage pools

A storage pool specifies the following for all volumes within the pool: Capacity, service level, Virtual Private Cloud (VPC), Active Directory policy, LDAP and encryption. [Learn more](#)

Filter Name: gcmv-ora* Search for storage pools by name, location, etc.

Status	Name	Location	Service level	Capacity	Allocated to Volumes	Volume count	Created	Labels	Show more
Ready	gcmv-ora-west	us-west4	Premium	2048 GiB	2000 GiB	8	Apr 9, 2025, 10:53:36 AM	database: oracle	
Ready	gcmv-ora-east	us-east4	Premium	2048 GiB	2000 GiB	8	Apr 7, 2025, 11:50:52 AM	database: oracle	

5. Create database volumes in storage pools. Referred to documentation [Create a volume quickstart](#) for step-by-step instructions. Following are some screen shots for quick reference.

Google Cloud

cvs-pm-host-1p

Search (/) for resources, docs, products, and more

Q Search

NetApp Volumes / Volumes / Create volume

Storage

Storage pools

Volumes

Data protection

Backups

Backup vaults

Policies

Active Directory policies

CMEK policies

Backup policies

Create volume

A volume provides NFS or SMB file services for your application with integrated data protection services. A volume is allocated from a storage pool and gets an individual or shared throughput limit based on its allocated capacity and storage pool service level. [Learn more](#)

Volume name *

orap-u01

Choice is permanent. Must be unique to the region. Use lowercase letters, numbers and hyphens. Start with a letter.

Description

Storage pool details

Select a storage pool in which to create the volume

Selected Storage Pool

gcnv-ora-east

Location

us-east4

Storage pool available capacity

2048 GiB

Number of volumes in the pool

0

Service level

Premium

VPC

shared-vpc-prod

Active Directory policy

No value

LDAP enabled

No

Encryption

Google-managed

[SELECT STORAGE POOL](#)

[CREATE NEW STORAGE POOL](#)

CREATE

CANCEL

Google Cloud

cvs-pm-host-1p

Search (/) for resources, docs, products, and more

Q Search

NetApp Volumes / Volumes / Create volume

Storage

Storage pools

Volumes

Data protection

Backups

Backup vaults

Policies

Active Directory policies

CMEK policies

Backup policies

Create volume

[SELECT STORAGE POOL](#)

[CREATE NEW STORAGE POOL](#)

Volume details

Share name *

orap-u01

Must be unique to a location

Capacity *

100

GiB

Capacity must be between 100 GiB and 102,400 GiB. Increments of 1 GiB

Auto-tiering

Optimize storage costs by automatically moving cold data on volumes with enabled auto-tiering to the most cost-effective access tier depending on access pattern changes. [Learn more](#)

Enable auto-tiering

All data identified as unaccessed according to the cooling threshold will be moved to lower cost cold tier storage. Once enabled, it can be paused but not disabled.

Protocol(s) *

NFSv3

Configuration for selected protocol(s)

Block volume from deletion when clients are connected

Required for volumes used as GCVE datastores. Choice is permanent.

Export rules

Rules are evaluated in order. First matching rule applies.

Rules

CREATE

CANCEL

Google Cloud

cvs-pm-host-1p

Search (/) for resources, docs, products, and more

Q Search

NetApp Volumes / Volumes / Create volume

Storage

Storage pools

Volumes

Data protection

Backups

Backup vaults

Policies

Active Directory policies

CMEK policies

Backup policies

Create volume

Block volume from deletion when clients are connected

Required for volumes used as GCVE datastores. Choice is permanent.

Export rules

Rules are evaluated in order. First matching rule applies.

Rules

Edit Rule

Allowed clients *

0.0.0.0/0

Comma-separated list of IPv4 addresses or CIDRs (up to 4096 characters).

Access *

Read & Write

Read Only

Root Access (no_root_squash)

On

Off

DONE

ADD RULE

Snapshot configuration

Make snapshot directory visible

Make the snapshot (NFS) or --snapshot (SMB) directory visible to clients. For SMB volumes, it also enforces "ShareName\version" format. For NFSv4.1 volumes, the directory itself will not.

CREATE

CANCEL

7

Google Cloud

cvs-pm-host-1p

Search (/) for resources, docs, products, and more

Q Search

NetApp Volumes / Volumes / Create volume

Storage

Storage pools

Volumes

Data protection

Backups

Backup vaults

Policies

Active Directory policies

CMEK policies

Backup policies

Create a volume

Snapshot configuration

Make snapshot directory visible

Makes snapshot (NFS) or --snapshot (SMB) directory visible to clients. For SMB volumes, it also enables "Previous versions" support. For NFSv4.1 volumes, the directory itself will not be listed but can be accessed to list contents, etc.

Allow scheduled snapshots

When enabled, snapshots are created according to the schedule configured below.

HOURLY

DAILY

WEEKLY

MONTHLY

Snapshots To Keep

0

Hour (UTC)

Every hour

Minute (UTC)

0

The snapshot schedule will not be applied to the volume. To apply the snapshot schedule, adjust the retention controls.

Backup configuration

Use backup policy to automate volume backups.

Allow scheduled backups

When enabled, backups are automatically created according to the following specified policy.

Labels

Label your volumes for reports, queries.

+ ADD LABEL

CREATE

CANCEL

Google Cloud

cvs-pm-host-1p

Search (/) for resources, docs, products, and more

Q Search

NetApp Volumes / Volumes

Storage

Storage pools

Volumes

Data protection

Backups

Backup vaults

Policies

Active Directory policies

CMEK policies

Backup policies

Volumes

CREATE

REFRESH

A volume provides NFS or SMB file services for your application, with integrated data protection services. A volume is allocated from a pool and will get an individual throughput limit based on its allocated size and the pool service level. [Learn more](#)

Filter

Name: orap*

Name:

Status	Name	Location	Service level	Share name	Capacity	Used	Protocol(s)	Storage pool	Labels	Show more
Ready	orap-u08	us-east4	Premium	orap-u08	400	80.25% (321 GiB)	NFSV3	gcnv-ora-east	database: oracle	⋮
Ready	orap-u07	us-east4	Premium	orap-u07	250	80% (200 GiB)	NFSV3	gcnv-ora-east	database: oracle	⋮
Ready	orap-u06	us-east4	Premium	orap-u06	250	80% (200 GiB)	NFSV3	gcnv-ora-east	database: oracle	⋮
Ready	orap-u05	us-east4	Premium	orap-u05	250	80% (200 GiB)	NFSV3	gcnv-ora-east	database: oracle	⋮
Ready	orap-u04	us-east4	Premium	orap-u04	250	80% (200 GiB)	NFSV3	gcnv-ora-east	database: oracle	⋮
Ready	orap-u03	us-east4	Premium	orap-u03	250	80% (200 GiB)	NFSV3	gcnv-ora-east	database: oracle	⋮
Ready	orap-u02	us-east4	Premium	orap-u02	250	80% (200 GiB)	NFSV3	gcnv-ora-east	database: oracle	⋮
Ready	orap-u01	us-east4	Premium	orap-u01	100	21% (21 GiB)	NFSV3	gcnv-ora-east	database: oracle	⋮

Google Cloud

cvs-pm-host-1p

Search (/) for resources, docs, products, and more

Q Search

NetApp Volumes / Volumes

Storage

Storage pools

Volumes

Data protection

Backups

Backup vaults

Policies

Active Directory policies

CMEK policies

Backup policies

Volumes

CREATE

REFRESH

A volume provides NFS or SMB file services for your application, with integrated data protection services. A volume is allocated from a pool and will get an individual throughput limit based on its allocated size and the pool service level. [Learn more](#)

Filter

Name: oras*

Name:

Status	Name	Location	Service level	Share name	Capacity	Used	Protocol(s)	Storage pool	Labels	Show more
Ready	oras-u08	us-west4	Premium	oras-u08	400	79% (316 GiB)	NFSV3	gcnv-ora-west	database: oracle	⋮
Ready	oras-u07	us-west4	Premium	oras-u07	250	80% (200 GiB)	NFSV3	gcnv-ora-west	database: oracle	⋮
Ready	oras-u06	us-west4	Premium	oras-u06	250	80% (200 GiB)	NFSV3	gcnv-ora-west	database: oracle	⋮
Ready	oras-u05	us-west4	Premium	oras-u05	250	80% (200 GiB)	NFSV3	gcnv-ora-west	database: oracle	⋮
Ready	oras-u04	us-west4	Premium	oras-u04	250	80% (200 GiB)	NFSV3	gcnv-ora-west	database: oracle	⋮
Ready	oras-u03	us-west4	Premium	oras-u03	250	80% (200 GiB)	NFSV3	gcnv-ora-west	database: oracle	⋮
Ready	oras-u02	us-west4	Premium	oras-u02	250	80% (200 GiB)	NFSV3	gcnv-ora-west	database: oracle	⋮
Ready	oras-u01	us-west4	Premium	oras-u01	100	19% (19 GiB)	NFSV3	gcnv-ora-west	database: oracle	⋮

- The primary Oracle database should have been installed and configured in the primary Oracle DB server. On the other hand, in the standby Oracle DB server or the clone Oracle DB server, only Oracle software is installed and no Oracle databases are created. Ideally, the Oracle files directories layout should be exactly matching on all Oracle DB servers. Refer to TR-4974 for help on Oracle grid infrastructure and database installation and configuration with NFS/ASM. Although the solution is

8

validated on AWS FSx/EC2 environment, it can be equally applied to Google GCNV/Compute Engine environment.

- [TR-4974: Oracle 19c in Standalone Restart on AWS FSx/EC2 with NFS/ASM](#)

Primary Oracle VLDB configuration for Data Guard

In this demonstration, we have setup a primary Oracle database called NTAP on the primary DB server with eight NFS mount points: /u01 for the Oracle binary, /u02, /u03, /u04, /u05, /u06, /u07 for the Oracle data files, and load balanced with Oracle ASM disk group +DATA; /u08 for the Oracle active logs, archived log files, and load balanced with Oracle ASM disk group +LOGS. Oracle control files are placed on both +DATA and +LOGS disk groups for redundancy. This setup serves as a reference configuration. Your actual deployment should take into consideration of your specific needs and requirements in terms of the storage pool sizing, the service level, the number of database volumes and the sizing of each volume.

For detailed step by step procedures for setting up Oracle Data Guard on NFS with ASM, please referred to TR-5002 - [Oracle Active Data Guard Cost Reduction with Azure NetApp Files](#). Although the procedures in TR-5002 were validated on Azure ANF environment, they are equally applicable to Google GCNV environment.

Following illustrates the details of a primary Oracle VLDB in a Data Guard configuration in Google GCNV environment.

1. The primary database NTAP in the primary compute engine DB server is deployed as a single instance database in a standalone Restart configuration on the GCNV storage with NFS protocol and ASM as database storage volume manager.

```
orap.us-east4-a.c.cvs-pm-host-1p.internal:
Zone: us-east-4a
size: n1-standard-4 (4 vCPUs, 15 GB Memory)
OS: Linux (redhat 8.10)
pub_ip: 35.212.124.14
pri_ip: 10.70.11.5

[oracle@orap ~]$ df -h
Filesystem                                Size  Used Avail Use% Mounted on
devtmpfs                                  7.2G   0    7.2G   0% /dev
tmpfs                                      7.3G   0    7.3G   0% /dev/shm
tmpfs                                      7.3G  8.5M    7.2G   1% /run
tmpfs                                      7.3G   0    7.3G   0% /sys/fs/cgroup
/dev/sda2                                 50G   40G   11G   80% /
/dev/sda1                                 200M   5.9M  194M   3% /boot/efi
10.165.128.180:/orap-u05                  250G  201G   50G   81% /u05
10.165.128.180:/orap-u08                  400G  322G   79G   81% /u08
10.165.128.180:/orap-u04                  250G  201G   50G   81% /u04
10.165.128.180:/orap-u07                  250G  201G   50G   81% /u07
10.165.128.180:/orap-u02                  250G  201G   50G   81% /u02
10.165.128.180:/orap-u06                  250G  201G   50G   81% /u06
10.165.128.180:/orap-u01                  100G   21G   80G   21% /u01
10.165.128.180:/orap-u03                  250G  201G   50G   81% /u03

[oracle@orap ~]$ cat /etc/oratab
#
```

```
# This file is used by ORACLE utilities.  It is created by root.sh
# and updated by either Database Configuration Assistant while
# creating
# a database or ASM Configuration Assistant while creating ASM
# instance.

# A colon, ':', is used as the field terminator.  A new line
# terminates
# the entry.  Lines beginning with a pound sign, '#', are comments.
#
# Entries are of the form:
#   $ORACLE_SID:$ORACLE_HOME:<N|Y>:
#
# The first and second fields are the system identifier and home
# directory of the database respectively.  The third field indicates
# to the dbstart utility that the database should , "Y", or should
# not,
# "N", be brought up at system boot time.
#
# Multiple entries with the same $ORACLE_SID are not allowed.
#
#
+ASM:/u01/app/oracle/product/19.0.0/grid:N
NTAP:/u01/app/oracle/product/19.0.0/NTAP:N
```

2. Login to primary DB server as the oracle user. Validate grid configuration.

```
$GRID_HOME/bin/crsctl stat res -t
```

```

[oracle@orap ~]$ $GRID_HOME/bin/crsctl stat res -t
-----
-----
Name          Target    State          Server          State
details
-----
-----
Local Resources
-----
-----
ora.DATA.dg
          ONLINE    ONLINE        orap            STABLE
ora.LISTENER.lsnr
          ONLINE    ONLINE        orap            STABLE
ora.LOGS.dg
          ONLINE    ONLINE        orap            STABLE
ora.asm
          ONLINE    ONLINE        orap
Started, STABLE
ora.ons
          OFFLINE   OFFLINE        orap            STABLE
-----
-----
Cluster Resources
-----
-----
ora.cssd
      1      ONLINE    ONLINE        orap            STABLE
ora.diskmon
      1      OFFLINE   OFFLINE        STABLE
ora.evmd
      1      ONLINE    ONLINE        orap            STABLE
ora.ntap.db
      1      ONLINE    ONLINE        orap
Open, HOME=/u01/app/o
racle/product/19.0.0
/NTAP, STABLE
-----
-----
[oracle@orap ~]$

```

3. ASM disk group configuration.

asmcmd

```

[oracle@orap ~]$ asmcmd
ASMCMDB> lsdg
State      Type      Rebal  Sector  Logical_Sector  Block      AU
Total_MB   Free_MB   Req_mir_free_MB  Usable_file_MB  Offline_disks
Voting_files  Name
MOUNTED    EXTERN    N      512      512      4096      4194304
1228800    1219888      0      1219888      0
N  DATA/
MOUNTED    EXTERN    N      512      512      4096      4194304
327680     326556      0      326556      0
N  LOGS/
ASMCMDB> lsdk
Path
/u02/oradata/asm/orap_data_disk_01
/u02/oradata/asm/orap_data_disk_02
/u02/oradata/asm/orap_data_disk_03
/u02/oradata/asm/orap_data_disk_04
/u03/oradata/asm/orap_data_disk_05
/u03/oradata/asm/orap_data_disk_06
/u03/oradata/asm/orap_data_disk_07
/u03/oradata/asm/orap_data_disk_08
/u04/oradata/asm/orap_data_disk_09
/u04/oradata/asm/orap_data_disk_10
/u04/oradata/asm/orap_data_disk_11
/u04/oradata/asm/orap_data_disk_12
/u05/oradata/asm/orap_data_disk_13
/u05/oradata/asm/orap_data_disk_14
/u05/oradata/asm/orap_data_disk_15
/u05/oradata/asm/orap_data_disk_16
/u06/oradata/asm/orap_data_disk_17
/u06/oradata/asm/orap_data_disk_18
/u06/oradata/asm/orap_data_disk_19
/u06/oradata/asm/orap_data_disk_20
/u07/oradata/asm/orap_data_disk_21
/u07/oradata/asm/orap_data_disk_22
/u07/oradata/asm/orap_data_disk_23
/u07/oradata/asm/orap_data_disk_24
/u08/oralogs/asm/orap_logs_disk_01
/u08/oralogs/asm/orap_logs_disk_02
/u08/oralogs/asm/orap_logs_disk_03
/u08/oralogs/asm/orap_logs_disk_04
ASMCMDB>

```

4. Parameters setting for Data Guard on primary DB.


```
SQL> show parameter name
```

NAME	TYPE	VALUE
-----	-----	

cdb_cluster_name	string	
cell_offloadgroup_name	string	
db_file_name_convert	string	
db_name	string	ntap
db_unique_name	string	ntap_ny
global_names	boolean	FALSE
instance_name	string	NTAP
lock_name_space	string	
log_file_name_convert	string	
pdb_file_name_convert	string	
processor_group_name	string	
NAME	TYPE	VALUE
-----	-----	

service_names	string	ntap_ny.cvs-pm-
host-lp.interna		

```
SQL> sho parameter log_archive_dest
```

NAME	TYPE	VALUE
-----	-----	

log_archive_dest	string	
log_archive_dest_1	string	
LOCATION=USE_DB_RECOVERY_FILE_		
		DEST
VALID_FOR=(ALL_LOGFILES,A		
		LL_ROLES)
DB_UNIQUE_NAME=NTAP_		
		NY
log_archive_dest_10	string	
log_archive_dest_11	string	
log_archive_dest_12	string	
log_archive_dest_13	string	
log_archive_dest_14	string	
log_archive_dest_15	string	

NAME	TYPE	VALUE
-----	-----	

```

log_archive_dest_16          string
log_archive_dest_17          string
log_archive_dest_18          string
log_archive_dest_19          string
log_archive_dest_2           string      SERVICE=NTAP_LA
ASYNC VALID_FO

R=(ONLINE_LOGFILES,PRIMARY_ROL

                                E)

DB_UNIQUE_NAME=NTAP_LA
log_archive_dest_20          string
log_archive_dest_21          string
log_archive_dest_22          string

```

5. Primary DB configuration.

```
SQL> select name, open_mode, log_mode from v$database;
```

NAME	OPEN_MODE	LOG_MODE
NTAP	READ WRITE	ARCHIVELOG

```
SQL> show pdbs
```

CON_ID	CON_NAME	OPEN MODE	RESTRICTED
2	PDB\$SEED	READ ONLY	NO
3	NTAP_PDB1	READ WRITE	NO
4	NTAP_PDB2	READ WRITE	NO
5	NTAP_PDB3	READ WRITE	NO

```
SQL> select name from v$datafile;
```

NAME
+DATA/NTAP/DATAFILE/system.257.1198026005
+DATA/NTAP/DATAFILE/sysaux.258.1198026051
+DATA/NTAP/DATAFILE/undotbs1.259.1198026075
+DATA/NTAP/86B637B62FE07A65E053F706E80A27CA/DATAFILE/system.266.1198027075
+DATA/NTAP/86B637B62FE07A65E053F706E80A27CA/DATAFILE/sysaux.267.1198027075
+DATA/NTAP/DATAFILE/users.260.1198026077

```

+DATA/NTAP/86B637B62FE07A65E053F706E80A27CA/DATAFILE/undotbs1.268.11
98027075
+DATA/NTAP/32639B76C9BC91A8E063050B460A2116/DATAFILE/system.272.1198
028157
+DATA/NTAP/32639B76C9BC91A8E063050B460A2116/DATAFILE/sysaux.273.1198
028157
+DATA/NTAP/32639B76C9BC91A8E063050B460A2116/DATAFILE/undotbs1.271.11
98028157
+DATA/NTAP/32639B76C9BC91A8E063050B460A2116/DATAFILE/users.275.11980
28185

```

NAME

```

-----
-----
+DATA/NTAP/32639D40D02D925FE063050B460A60E3/DATAFILE/system.277.1198
028187
+DATA/NTAP/32639D40D02D925FE063050B460A60E3/DATAFILE/sysaux.278.1198
028187
+DATA/NTAP/32639D40D02D925FE063050B460A60E3/DATAFILE/undotbs1.276.11
98028187
+DATA/NTAP/32639D40D02D925FE063050B460A60E3/DATAFILE/users.280.11980
28209
+DATA/NTAP/32639E973AF79299E063050B460AFBAD/DATAFILE/system.282.1198
028209
+DATA/NTAP/32639E973AF79299E063050B460AFBAD/DATAFILE/sysaux.283.1198
028209
+DATA/NTAP/32639E973AF79299E063050B460AFBAD/DATAFILE/undotbs1.281.11
98028209
+DATA/NTAP/32639E973AF79299E063050B460AFBAD/DATAFILE/users.285.11980
28229

```

19 rows selected.

```
SQL> select member from v$logfile;
```

MEMBER

```

-----
-----
+DATA/NTAP/ONLINELOG/group_3.264.1198026139
+LOGS/NTAP/ONLINELOG/group_3.259.1198026147
+DATA/NTAP/ONLINELOG/group_2.263.1198026137
+LOGS/NTAP/ONLINELOG/group_2.258.1198026145
+DATA/NTAP/ONLINELOG/group_1.262.1198026137
+LOGS/NTAP/ONLINELOG/group_1.257.1198026145
+DATA/NTAP/ONLINELOG/group_4.286.1198511423

```

```
+LOGS/NTAP/ONLINELOG/group_4.265.1198511425
+DATA/NTAP/ONLINELOG/group_5.287.1198511445
+LOGS/NTAP/ONLINELOG/group_5.266.1198511447
+DATA/NTAP/ONLINELOG/group_6.288.1198511459
```

MEMBER

```
-----
-----
```

```
+LOGS/NTAP/ONLINELOG/group_6.267.1198511461
+DATA/NTAP/ONLINELOG/group_7.289.1198511477
+LOGS/NTAP/ONLINELOG/group_7.268.1198511479
```

14 rows selected.

```
SQL> select name from v$controlfile;
```

NAME

```
-----
-----
```

```
+DATA/NTAP/CONTROLFILE/current.261.1198026135
+LOGS/NTAP/CONTROLFILE/current.256.1198026135
```

6. Oracle listener configuration.

```
lsnrctl status listener
```

```
[oracle@orap admin]$ lsnrctl status
```

```
LSNRCTL for Linux: Version 19.0.0.0.0 - Production on 15-APR-2025
16:14:00
```

Copyright (c) 1991, 2022, Oracle. All rights reserved.

```
Connecting to (ADDRESS=(PROTOCOL=tcp) (HOST=) (PORT=1521))
STATUS of the LISTENER
```

```
-----
```

Alias	LISTENER
Version	TNSLSNR for Linux: Version 19.0.0.0.0 -
Production	
Start Date	14-APR-2025 19:44:21
Uptime	0 days 20 hr. 29 min. 38 sec
Trace Level	off
Security	ON: Local OS Authentication

```

SNMP                                OFF
Listener Parameter File
/u01/app/oracle/product/19.0.0/grid/network/admin/listener.ora
Listener Log File
/u01/app/oracle/diag/tnslsnr/orap/listener/alert/log.xml
Listening Endpoints Summary...
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp) (HOST=orap.us-east4-a.c.cvs-
pm-host-1p.internal) (PORT=1521)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=ipc) (KEY=EXTPROC1521)))
Services Summary...
Service "+ASM" has 1 instance(s).
  Instance "+ASM", status READY, has 1 handler(s) for this
service...
Service "+ASM_DATA" has 1 instance(s).
  Instance "+ASM", status READY, has 1 handler(s) for this
service...
Service "+ASM_LOGS" has 1 instance(s).
  Instance "+ASM", status READY, has 1 handler(s) for this
service...
Service "32639b76c9bc91a8e063050b460a2116.cvs-pm-host-1p.internal"
has 1 instance(s).
  Instance "NTAP", status READY, has 1 handler(s) for this
service...
Service "32639d40d02d925fe063050b460a60e3.cvs-pm-host-1p.internal"
has 1 instance(s).
  Instance "NTAP", status READY, has 1 handler(s) for this
service...
Service "32639e973af79299e063050b460afbad.cvs-pm-host-1p.internal"
has 1 instance(s).
  Instance "NTAP", status READY, has 1 handler(s) for this
service...
Service "86b637b62fdf7a65e053f706e80a27ca.cvs-pm-host-1p.internal"
has 1 instance(s).
  Instance "NTAP", status READY, has 1 handler(s) for this
service...
Service "NTAPXDB.cvs-pm-host-1p.internal" has 1 instance(s).
  Instance "NTAP", status READY, has 1 handler(s) for this
service...
Service "NTAP_NY_DGMGRL.cvs-pm-host-1p.internal" has 1 instance(s).
  Instance "NTAP", status UNKNOWN, has 1 handler(s) for this
service...
Service "ntap.cvs-pm-host-1p.internal" has 1 instance(s).
  Instance "NTAP", status READY, has 1 handler(s) for this
service...
Service "ntap_pdb1.cvs-pm-host-1p.internal" has 1 instance(s).
  Instance "NTAP", status READY, has 1 handler(s) for this

```

```

service...
Service "ntap_pdb2.cvs-pm-host-1p.internal" has 1 instance(s).
  Instance "NTAP", status READY, has 1 handler(s) for this
service...
Service "ntap_pdb3.cvs-pm-host-1p.internal" has 1 instance(s).
  Instance "NTAP", status READY, has 1 handler(s) for this
service...
The command completed successfully

```

7. Flashback is enabled at primary database.

```
SQL> select name, database_role, flashback_on from v$database;
```

NAME	DATABASE_ROLE	FLASHBACK_ON
NTAP	PRIMARY	YES

8. dNFS configuration on primary DB.

```
SQL> select svrname, dirname from v$dnfs_servers;
```

SVRNAME	DIRNAME
10.165.128.180	/orap-u04
10.165.128.180	/orap-u05
10.165.128.180	/orap-u07
10.165.128.180	/orap-u03

```
10.165.128.180
/orap-u06

10.165.128.180
/orap-u02

SVRNAME
-----
-----
DIRNAME
-----
-----
10.165.128.180
/orap-u08

10.165.128.180
/orap-u01

8 rows selected.
```

This completes the demonstration of a Data Guard setup for VLDB NTAP at the primary site on GCNV with NFS/ASM.

Standby Oracle VLDB configuration for Data Guard

Oracle Data Guard requires OS kernel configuration and Oracle software stacks including patch sets on standby DB server to match with primary DB server. For easy management and simplicity, the database storage configuration of the standby DB server ideally should match with the primary DB server as well, such as the database directory layout and sizes of NFS mount points.

Again, for detailed step by step procedures for setting up Oracle Data Guard standby on NFS with ASM, please refer to [TR-5002 - Oracle Active Data Guard Cost Reduction with Azure NetApp Files](#) and [TR-4974 - Oracle 19c in Standalone Restart on AWS FSx/EC2 with NFS/ASM](#) relevant sections. Following illustrates the detail of standby Oracle VLDB configuration on standby DB server in a Data Guard setting in Google GCNV environment.

1. The standby Oracle DB server configuration at standby site in the demo lab.

```
oras.us-west4-a.c.cvs-pm-host-1p.internal:
Zone: us-west4-a
size: n1-standard-4 (4 vCPUs, 15 GB Memory)
OS: Linux (redhat 8.10)
pub_ip: 35.219.129.195
pri_ip: 10.70.14.16

[oracle@oras ~]$ df -h
Filesystem                Size      Used Avail Use% Mounted on
devtmpfs                  7.2G         0  7.2G   0% /dev
tmpfs                     7.3G      1.1G   6.2G  16% /dev/shm
tmpfs                     7.3G      8.5M   7.2G   1% /run
tmpfs                     7.3G         0  7.3G   0% /sys/fs/cgroup
/dev/sda2                  50G       40G   11G  80% /
/dev/sda1                  200M     5.9M  194M   3% /boot/efi
10.165.128.197:/oras-u07   250G     201G   50G  81% /u07
10.165.128.197:/oras-u06   250G     201G   50G  81% /u06
10.165.128.197:/oras-u02   250G     201G   50G  81% /u02
10.165.128.196:/oras-u03   250G     201G   50G  81% /u03
10.165.128.196:/oras-u01  100G       20G   81G  20% /u01
10.165.128.197:/oras-u05   250G     201G   50G  81% /u05
10.165.128.197:/oras-u04   250G     201G   50G  81% /u04
10.165.128.197:/oras-u08  400G     317G   84G  80% /u08

[oracle@oras ~]$ cat /etc/oratab
#Backup file is
/u01/app/oracle/crsdata/oras/output/oratab.bak.oras.oracle line
added by Agent
#

# This file is used by ORACLE utilities.  It is created by root.sh
# and updated by either Database Configuration Assistant while
```



```

creating
# a database or ASM Configuration Assistant while creating ASM
instance.

# A colon, ':', is used as the field terminator. A new line
terminates
# the entry. Lines beginning with a pound sign, '#', are comments.
#
# Entries are of the form:
#   $ORACLE_SID:$ORACLE_HOME:<N|Y>:
#
# The first and second fields are the system identifier and home
# directory of the database respectively. The third field indicates
# to the dbstart utility that the database should , "Y", or should
not,
# "N", be brought up at system boot time.
#
# Multiple entries with the same $ORACLE_SID are not allowed.
#
#
+ASM:/u01/app/oracle/product/19.0.0/grid:N
NTAP:/u01/app/oracle/product/19.0.0/NTAP:N           # line added
by Agent

```

2. Grid infrastructure configuration on standby DB server.

```
[oracle@oras ~]$ $GRID_HOME/bin/crsctl stat res -t
```

```
-----
Name          Target    State          Server          State
details
-----
Local Resources
-----
ora.DATA.dg
          ONLINE    ONLINE        oras            STABLE
ora.LISTENER.lsnr
          ONLINE    ONLINE        oras            STABLE
ora.LOGS.dg
          ONLINE    ONLINE        oras            STABLE
ora.asm
          ONLINE    ONLINE        oras
Started, STABLE
ora.ons
          OFFLINE   OFFLINE        oras            STABLE
-----
Cluster Resources
-----
ora.cssd
    1      ONLINE    ONLINE        oras            STABLE
ora.diskmon
    1      OFFLINE   OFFLINE
ora.evmd
    1      ONLINE    ONLINE        oras            STABLE
ora.ntap_la.db
    1      ONLINE    INTERMEDIATE oras
Dismounted, Mount Ini
tiated, HOME=/u01/app
/oracle/product/19.0
.0/NTAP, STABLE
-----
-----
```

3. ASM disk groups configuration on standby DB server.

```

[oracle@oras ~]$ asmcmd
ASMCMD> lsdg
State      Type      Rebal  Sector  Logical_Sector  Block      AU
Total_MB   Free_MB   Req_mir_free_MB  Usable_file_MB  Offline_disks
Voting_files  Name
MOUNTED    EXTERN    N       512      512      4096      4194304
1228800    1228420              0        1228420              0
N  DATA/
MOUNTED    EXTERN    N       512      512      4096      4194304
322336     322204              0        322204              0
N  LOGS/
ASMCMD> lsdk
Path
/u02/oradata/asm/oras_data_disk_01
/u02/oradata/asm/oras_data_disk_02
/u02/oradata/asm/oras_data_disk_03
/u02/oradata/asm/oras_data_disk_04
/u03/oradata/asm/oras_data_disk_05
/u03/oradata/asm/oras_data_disk_06
/u03/oradata/asm/oras_data_disk_07
/u03/oradata/asm/oras_data_disk_08
/u04/oradata/asm/oras_data_disk_09
/u04/oradata/asm/oras_data_disk_10
/u04/oradata/asm/oras_data_disk_11
/u04/oradata/asm/oras_data_disk_12
/u05/oradata/asm/oras_data_disk_13
/u05/oradata/asm/oras_data_disk_14
/u05/oradata/asm/oras_data_disk_15
/u05/oradata/asm/oras_data_disk_16
/u06/oradata/asm/oras_data_disk_17
/u06/oradata/asm/oras_data_disk_18
/u06/oradata/asm/oras_data_disk_19
/u06/oradata/asm/oras_data_disk_20
/u07/oradata/asm/oras_data_disk_21
/u07/oradata/asm/oras_data_disk_22
/u07/oradata/asm/oras_data_disk_23
/u07/oradata/asm/oras_data_disk_24
/u08/oralogs/asm/oras_logs_disk_01
/u08/oralogs/asm/oras_logs_disk_02
/u08/oralogs/asm/oras_logs_disk_03
/u08/oralogs/asm/oras_logs_disk_04
ASMCMD>

```

4. Parameters setting for Data Guard on standby DB.

```
SQL> show parameter name
```

NAME	TYPE	VALUE

cdb_cluster_name	string	
cell_offloadgroup_name	string	
db_file_name_convert	string	
db_name	string	NTAP
db_unique_name	string	NTAP_LA
global_names	boolean	FALSE
instance_name	string	NTAP
lock_name_space	string	
log_file_name_convert	string	
pdb_file_name_convert	string	
processor_group_name	string	
NAME	TYPE	VALUE

service_names	string	NTAP_LA.cvs-pm-
host-lp.interna		1

```
SQL> show parameter log_archive_config
```

NAME	TYPE	VALUE

log_archive_config	string	
DG_CONFIG=(NTAP_NY,NTAP_LA)		

```
SQL> show parameter fal_server
```

NAME	TYPE	VALUE

fal_server	string	NTAP_NY

5. Standby DB configuration.

```
SQL> select name, open_mode, log_mode from v$database;
```

NAME	OPEN_MODE	LOG_MODE

NTAP	MOUNTED	ARCHIVELOG

```
SQL> show pdbs
```

CON_ID	CON_NAME	OPEN MODE	RESTRICTED
2	PDB\$SEED	MOUNTED	
3	NTAP_PDB1	MOUNTED	
4	NTAP_PDB2	MOUNTED	
5	NTAP_PDB3	MOUNTED	

```
SQL> select name from v$datafile;
```

NAME

```
+DATA/NTAP_LA/DATAFILE/system.261.1198520347
+DATA/NTAP_LA/DATAFILE/sysaux.262.1198520373
+DATA/NTAP_LA/DATAFILE/undotbs1.263.1198520399
+DATA/NTAP_LA/32635CC1DCF58A60E063050B460AB746/DATAFILE/system.264.1
198520417
+DATA/NTAP_LA/32635CC1DCF58A60E063050B460AB746/DATAFILE/sysaux.265.1
198520435
+DATA/NTAP_LA/DATAFILE/users.266.1198520451
+DATA/NTAP_LA/32635CC1DCF58A60E063050B460AB746/DATAFILE/undotbs1.267
.1198520455
+DATA/NTAP_LA/32639B76C9BC91A8E063050B460A2116/DATAFILE/system.268.1
198520471
+DATA/NTAP_LA/32639B76C9BC91A8E063050B460A2116/DATAFILE/sysaux.269.1
198520489
+DATA/NTAP_LA/32639B76C9BC91A8E063050B460A2116/DATAFILE/undotbs1.270
.1198520505
+DATA/NTAP_LA/32639B76C9BC91A8E063050B460A2116/DATAFILE/users.271.11
98520513
```

NAME

```
+DATA/NTAP_LA/32639D40D02D925FE063050B460A60E3/DATAFILE/system.272.1
198520517
+DATA/NTAP_LA/32639D40D02D925FE063050B460A60E3/DATAFILE/sysaux.273.1
198520533
+DATA/NTAP_LA/32639D40D02D925FE063050B460A60E3/DATAFILE/undotbs1.274
.1198520551
+DATA/NTAP_LA/32639D40D02D925FE063050B460A60E3/DATAFILE/users.275.11
98520559
+DATA/NTAP_LA/32639E973AF79299E063050B460AFBAD/DATAFILE/system.276.1
```

```

198520563
+DATA/NTAP_LA/32639E973AF79299E063050B460AFBAD/DATAFILE/sysaux.277.1
198520579
+DATA/NTAP_LA/32639E973AF79299E063050B460AFBAD/DATAFILE/undotbs1.278
.1198520595
+DATA/NTAP_LA/32639E973AF79299E063050B460AFBAD/DATAFILE/users.279.11
98520605

```

19 rows selected.

```
SQL> select name from v$controlfile;
```

NAME

```

-----
-----
+DATA/NTAP_LA/CONTROLFILE/current.260.1198520303
+LOGS/NTAP_LA/CONTROLFILE/current.257.1198520305

```

```
SQL> select group#, type, member from v$logfile order by 2, 1;
```

GROUP#	TYPE	MEMBER
1	ONLINE	+DATA/NTAP_LA/ONLINELOG/group_1.280.1198520649
1	ONLINE	+LOGS/NTAP_LA/ONLINELOG/group_1.259.1198520651
2	ONLINE	+DATA/NTAP_LA/ONLINELOG/group_2.281.1198520659
2	ONLINE	+LOGS/NTAP_LA/ONLINELOG/group_2.258.1198520661
3	ONLINE	+DATA/NTAP_LA/ONLINELOG/group_3.282.1198520669
3	ONLINE	+LOGS/NTAP_LA/ONLINELOG/group_3.260.1198520671
4	STANDBY	+DATA/NTAP_LA/ONLINELOG/group_4.283.1198520677
4	STANDBY	+LOGS/NTAP_LA/ONLINELOG/group_4.261.1198520679
5	STANDBY	+DATA/NTAP_LA/ONLINELOG/group_5.284.1198520687
5	STANDBY	+LOGS/NTAP_LA/ONLINELOG/group_5.262.1198520689
6	STANDBY	+DATA/NTAP_LA/ONLINELOG/group_6.285.1198520697

GROUP#	TYPE	MEMBER
6	STANDBY	+LOGS/NTAP_LA/ONLINELOG/group_6.263.1198520699
7	STANDBY	+DATA/NTAP_LA/ONLINELOG/group_7.286.1198520707
7	STANDBY	+LOGS/NTAP_LA/ONLINELOG/group_7.264.1198520709

14 rows selected.

6. Validate the standby database recovery status. Notice the recovery logmerger in APPLYING_LOG action.

```
SQL> SELECT ROLE, THREAD#, SEQUENCE#, ACTION FROM
V$DATAGUARD_PROCESS;
```

ROLE	THREAD#	SEQUENCE#	ACTION
post role transition	0	0	IDLE
recovery apply slave	0	0	IDLE
recovery apply slave	0	0	IDLE
recovery apply slave	0	0	IDLE
recovery apply slave	0	0	IDLE
recovery logmerger	1	24	APPLYING_LOG
managed recovery	0	0	IDLE
RFS ping	1	24	IDLE
archive redo	0	0	IDLE
archive redo	0	0	IDLE
gap manager	0	0	IDLE

ROLE	THREAD#	SEQUENCE#	ACTION
archive local	0	0	IDLE
redo transport timer	0	0	IDLE
archive redo	0	0	IDLE
RFS async	1	24	IDLE
redo transport monitor	0	0	IDLE
log writer	0	0	IDLE

17 rows selected.

7. Flashback is enabled at standby database.

```
SQL> select name, database_role, flashback_on from v$database;
```

NAME	DATABASE_ROLE	FLASHBACK_ON
NTAP	PHYSICAL STANDBY	YES

8. dNFS configuration on standby DB.

```
SQL> select svrname, dirname from v$dnfs_servers;
```

```
SVRNAME
```

```
-----
```

```
-----
```

```
DIRNAME
```

```
-----
```

```
-----
```

```
10.165.128.197
```

```
/oras-u04
```

```
10.165.128.197
```

```
/oras-u05
```

```
10.165.128.197
```

```
/oras-u06
```

```
10.165.128.197
```

```
/oras-u07
```

```
10.165.128.197
```

```
/oras-u02
```

```
10.165.128.197
```

```
/oras-u08
```

```
10.165.128.196
```

```
/oras-u03
```

```
10.165.128.196
```

```
/oras-u01
```

```
8 rows selected.
```

This completes the demonstration of a Data Guard setup for VLDB NTAP with managed standby recovery enabled at standby site.

Setup Data Guard Broker and FSFO with an Observer

Setup Data Guard Broker

Oracle Data Guard broker is a distributed management framework that automates and centralizes the creation, maintenance, and monitoring of Oracle Data Guard configurations. Following section demonstrate how to setup Data Guard Broker to manage Data Guard environment.

1. Start data guard broker on both the primary and the standby databases with following command via sqlplus.

```
alter system set dg_broker_start=true scope=both;
```

2. From primary database, connect to Data Guard Borker as SYSDBA.

```
[oracle@orap ~]$ dgmgrl sys@NTAP_NY
DGMGRL for Linux: Release 19.0.0.0.0 - Production on Wed Dec 11
20:53:20 2024
Version 19.18.0.0.0

Copyright (c) 1982, 2019, Oracle and/or its affiliates. All rights
reserved.

Welcome to DGMGRL, type "help" for information.
Password:
Connected to "NTAP_NY"
Connected as SYSDBA.
DGMGRL>
```

3. Create and enable Data Guard Broker configuration.

```
DGMGRL> create configuration dg_config as primary database is
NTAP_NY connect identifier is NTAP_NY;
Configuration "dg_config" created with primary database "ntap_ny"
DGMGRL> add database NTAP_LA as connect identifier is NTAP_LA;
Database "ntap_la" added
DGMGRL> enable configuration;
Enabled.
DGMGRL> show configuration;

Configuration - dg_config

Protection Mode: MaxPerformance
Members:
  ntap_ny - Primary database
  ntap_la - Physical standby database

Fast-Start Failover: Disabled

Configuration Status:
SUCCESS (status updated 3 seconds ago)
```

4. Validate the database status within the Data Guard Broker management framework.

```
DGMGRL> show database ntap_ny;
```

```
Database - ntap_ny
```

```
Role:                PRIMARY
Intended State:       TRANSPORT-ON
Instance(s):         NTAP
```

```
Database Status:
SUCCESS
```

```
DGMGRL> show database ntap_la;
```

```
Database - ntap_la
```

```
Role:                PHYSICAL STANDBY
Intended State:       APPLY-ON
Transport Lag:        0 seconds (computed 0 seconds ago)
Apply Lag:            0 seconds (computed 0 seconds ago)
Average Apply Rate:   3.00 KByte/s
Real Time Query:      OFF
Instance(s):         NTAP
```

```
Database Status:
SUCCESS
```

```
DGMGRL>
```

In the event of a failure, Data Guard Broker can be used to failover the primary database to the standby instantaneously. If `Fast-Start Failover` is enabled, Data Guard Broker can failover the primary database to the standby when a failure is detected without an user intervention.

Configure FSFO with an Observer

Optionally, Fast Start Fail Over (FSFO) can be enabled for Data Guard Broker to failover the primary database to standby database in the event of a failure automatically. Following are the procedures to setup FSFO with an observer instance.

1. Create a lightweight Google compute engine instance to run Observer in a different zone than primary or standby DB server. In the test case, we used a N1 instance with 2 vCPU with 7.5G memory. Have same version of Oracle installed on the host.
2. Login in as oracle user and set oracle environment in the oracle user .bash_profile.

```
vi ~/.bash_profile
```

```
# .bash_profile

# Get the aliases and functions
if [ -f ~/.bashrc ]; then
    . ~/.bashrc
fi

# User specific environment and startup programs

export ORACLE_HOME=/u01/app/oracle/product/19.0.0/NTAP
export PATH=$ORACLE_HOME/bin:$PATH
```

3. Add primary and standby DB TNS name entries to tnsname.ora file.

```
vi $ORACLE_HOME/network/admin/tnsnames.ora
```

```

NTAP_NY =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = orap.us-east4-a.c.cvs-pm-
host-1p.internal) (PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = NTAP_NY.cvs-pm-host-1p.internal)
      (UR=A)
    )
  )

NTAP_LA =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = oras.us-west4-a.c.cvs-pm-
host-1p.internal) (PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = NTAP_LA.cvs-pm-host-1p.internal)
      (UR=A)
    )
  )

```

4. Create and initialize wallet with a password.

```
mkdir -p /u01/app/oracle/admin/NTAP/wallet
```

```
mkstore -wrl /u01/app/oracle/admin/NTAP/wallet -create
```

```

[oracle@orao NTAP]$ mkdir -p /u01/app/oracle/admin/NTAP/wallet
[oracle@orao NTAP]$ mkstore -wrl /u01/app/oracle/admin/NTAP/wallet
-create
Oracle Secret Store Tool Release 19.0.0.0.0 - Production
Version 19.4.0.0.0
Copyright (c) 2004, 2022, Oracle and/or its affiliates. All rights
reserved.

Enter password:
Enter password again:
[oracle@orao NTAP]$

```

5. Enable passwordless authentication for user sys of both primary and standby database. Enter sys password first, then wallet password from previous step.

```
mkstore -wrl /u01/app/oracle/admin/NTAP/wallet -createCredential NTAP_NY sys
```

```
mkstore -wrl /u01/app/oracle/admin/NTAP/wallet -createCredential NTAP_LA sys
```

```
[oracle@orao NTAP]$ mkstore -wrl /u01/app/oracle/admin/NTAP/wallet
-createCredential NTAP_NY sys
Oracle Secret Store Tool Release 19.0.0.0.0 - Production
Version 19.4.0.0.0
Copyright (c) 2004, 2022, Oracle and/or its affiliates. All rights
reserved.
```

Your secret/Password is missing in the command line

Enter your secret/Password:

Re-enter your secret/Password:

Enter wallet password:

```
[oracle@orao NTAP]$ mkstore -wrl /u01/app/oracle/admin/NTAP/wallet
-createCredential NTAP_LA sys
Oracle Secret Store Tool Release 19.0.0.0.0 - Production
Version 19.4.0.0.0
Copyright (c) 2004, 2022, Oracle and/or its affiliates. All rights
reserved.
```

Your secret/Password is missing in the command line

Enter your secret/Password:

Re-enter your secret/Password:

Enter wallet password:

```
[oracle@orao NTAP]$
```

6. Update sqlnet.ora with wallet location.

```
vi $ORACLE_HOME/network/admin/sqlnet.ora
```

```
WALLET_LOCATION =
  (SOURCE =
    (METHOD = FILE)
    (METHOD_DATA = (DIRECTORY =
/u01/app/oracle/admin/NTAP/wallet))
  )
SQLNET.WALLET_OVERRIDE = TRUE
```

7. Validate the credentials.

```
mkstore -wrl /u01/app/oracle/admin/NTAP/wallet -listCredential
```

```
sqlplus /@NTAP_LA as sysdba
```

```
sqlplus /@NTAP_NY as sysdba
```

```
[oracle@orao NTAP]$ mkstore -wrl /u01/app/oracle/admin/NTAP/wallet  
-listCredential  
Oracle Secret Store Tool Release 19.0.0.0.0 - Production  
Version 19.4.0.0.0  
Copyright (c) 2004, 2022, Oracle and/or its affiliates. All rights  
reserved.
```

```
Enter wallet password:  
List credential (index: connect_string username)  
2: NTAP_LA sys  
1: NTAP_NY sys
```

8. Configure and enable Fast-Start Failover.

```
mkdir /u01/app/oracle/admin/NTAP/fsfo
```

```
dgmgrl
```

Welcome to DGMGRL, type "help" for information.

```
DGMGRL> connect /@NTAP_NY
```

Connected to "ntap_ny"

Connected as SYSDBA.

```
DGMGRL> show configuration;
```

Configuration - dg_config

Protection Mode: MaxAvailability

Members:

ntap_ny - Primary database

ntap_la - Physical standby database

Fast-Start Failover: Disabled

Configuration Status:

SUCCESS (status updated 58 seconds ago)

```
DGMGRL> enable fast_start failover;
```

Enabled in Zero Data Loss Mode.

```
DGMGRL> show configuration;
```

Configuration - dg_config

Protection Mode: MaxAvailability

Members:

ntap_ny - Primary database

Warning: ORA-16819: fast-start failover observer not started

ntap_la - (*) Physical standby database

Fast-Start Failover: Enabled in Zero Data Loss Mode

Configuration Status:

WARNING (status updated 43 seconds ago)

9. Start and validate observer.

```
nohup dgmgrl /@NTAP_NY "start observer  
file='/u01/app/oracle/admin/NTAP/fsfo/fsfo.dat'" >>  
/u01/app/oracle/admin/NTAP/fsfo/dgmgrl.log &
```

```
[oracle@orao NTAP]$ nohup dgmgrl /@NTAP_NY "start observer  
file='/u01/app/oracle/admin/NTAP/fsfo/fsfo.dat'" >>
```



```
/u01/app/oracle/admin/NTAP/fsfo/dgmgrr.log &
```

```
[1] 94957
```

```
[oracle@orao fsfo]$ dgmgrr
```

```
DGMGRL for Linux: Release 19.0.0.0.0 - Production on Wed Apr 16
```

```
21:12:09 2025
```

```
Version 19.18.0.0.0
```

```
Copyright (c) 1982, 2019, Oracle and/or its affiliates. All rights reserved.
```

```
Welcome to DGMGRL, type "help" for information.
```

```
DGMGRL> connect /@NTAP_NY
```

```
Connected to "ntap_ny"
```

```
Connected as SYSDBA.
```

```
DGMGRL> show configuration verbose;
```

```
Configuration - dg_config
```

```
Protection Mode: MaxAvailability
```

```
Members:
```

```
ntap_ny - Primary database
```

```
ntap_la - (*) Physical standby database
```

```
(*) Fast-Start Failover target
```

```
Properties:
```

FastStartFailoverThreshold	= '30'
OperationTimeout	= '30'
TraceLevel	= 'USER'
FastStartFailoverLagLimit	= '30'
CommunicationTimeout	= '180'
ObserverReconnect	= '0'
FastStartFailoverAutoReinststate	= 'TRUE'
FastStartFailoverPmyShutdown	= 'TRUE'
BystandersFollowRoleChange	= 'ALL'
ObserverOverride	= 'FALSE'
ExternalDestination1	= ''
ExternalDestination2	= ''
PrimaryLostWriteAction	= 'CONTINUE'
ConfigurationWideServiceName	= 'ntap_CFG'

```
Fast-Start Failover: Enabled in Zero Data Loss Mode
```

```
Lag Limit: 30 seconds (not in use)
```

```
Threshold: 30 seconds
```

```
Active Target: ntap_la
```

```
Potential Targets: "ntap_la"
  ntap_la      valid
Observer:      orao
Shutdown Primary: TRUE
Auto-reinstate: TRUE
Observer Reconnect: (none)
Observer Override: FALSE
```

```
Configuration Status:
SUCCESS
```

```
DGMGRL>
```



To achieve zero data loss, Oracle Data Guard protection mode needs to be set to `MaxAvailability` or `MaxProtection` mode. The default protection mode of `MaxPerformance` can be changed from the Data Guard Broker interface by editing the Data Guard configuration and changing `LogXptMode` from `ASync` to `Sync`. The Oracle archive log destination log mode needs to be changed accordingly. When real-time log application is enabled for Data Guard as required for `MaxAvailability`, avoid rebooting the database automatically because automatic database reboot may inadvertently open the standby database in `READ ONLY WITH APPLY` mode, which requires an Active Data Guard license. Instead, boot the database manually to ensure it remains in a `MOUNT` state with managed recovery in real time.

Clone standby database for other use cases via automation

Following automation toolkit is specifically designed to create or refresh clones of an Oracle Data Guard standby DB deployed to GCNV with NFS/ASM configuration for a complete clone lifecycle management.

```
https://bitbucket.ngage.netapp.com/projects/NS-
BB/repos/na_oracle_clone_gcnv/browse
```



The toolkit can only be accessed by NetApp internal user with bitbucket access at this moment. For interested external users, please request access from your account team or reach out to NetApp Solutions Engineering team. Refer to [Automated Oracle Clone Lifecycle on GCNV with ASM](#) for usage instructions.

Where to find additional information

To learn more about the information described in this document, review the following documents and/or websites:

- TR-5002: Oracle Active Data Guard Cost Reduction with Azure NetApp Files

[TR-5002: Oracle Active Data Guard Cost Reduction with Azure NetApp Files](#)

- [TR-4974: Oracle 19c in Standalone Restart on AWS FSx/EC2 with NFS/ASM](#)

[TR-4974: Oracle 19c in Standalone Restart on AWS FSx/EC2 with NFS/ASM](#)

- NetApp's best-in-class file storage service, in Google Cloud

<https://cloud.google.com/netapp-volumes?hl=en>

- Oracle Data Guard Concepts and Administration

<https://docs.oracle.com/en/database/oracle/oracle-database/19/sbydb/index.html#Oracle%C2%AE-Data-Guard>

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.