



Azure Cloud

NetApp Solutions

NetApp
July 26, 2024

This PDF was generated from https://docs.netapp.com/us-en/netapp-solutions/databases/azure_ora_anf_inc_merge.html on July 26, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Azure Cloud 1
 - TR-4990: Quick Recovery of Oracle VLDB with Incremental Merge on ANF 1
 - TR-4987: Simplified, Automated Oracle Deployment on Azure NetApp Files with NFS 61
 - Oracle Database Deployment and Protection on Azure NetApp Files..... 79

Azure Cloud

TR-4990: Quick Recovery of Oracle VLDB with Incremental Merge on ANF

Allen Cao, Niyaz Mohamed, NetApp

This solution provides overview and details for quick recovery of Oracle VLDB deployed to Azure VM compute instance with NFS mount on Azure NetApp Files capacity pool to stage a standby database copy that is incrementally merged constantly via RMAN.

Purpose

Recovering a Very Large Database (VLDB) in Oracle using the Oracle Recovery Manager (RMAN) backup tool can be a highly challenging task. The database restoration process from backup media in the event of a failure can be time-consuming, delaying the database recovery and potentially impacting your Service Level Agreement (SLA) significantly. However, starting from version 10g, Oracle introduced a RMAN feature that allows users to create staged image copies of the Oracle database data files on additional disk storage located on the DB server host. These image copies can be incrementally updated using RMAN on a daily basis. In the case of a failure, the Database Administrator (DBA) can swiftly switch the Oracle database from the failed media to the image copy, eliminating the need for a complete database media restore. The result is a greatly improved SLA, albeit at the cost of doubling the required database storage.

If you are keen on SLA for your VLDB and contemplating moving the Oracle database to a public cloud such as Azure, you could set up a similar database protection structure using resources such as Microsoft Azure NetApp Files (ANF) for staging your standby database image copy. In this documentation, we demonstrate how to provision and export an NFS file system from ANF capacity pool to be mounted on an Oracle database server for staging a standby database copy for quick recovery in the event of a primary storage failure.

This solution addresses the following use cases:

- An Oracle VLDB image copy incremental merge via RMAN on NFS mount point off Microsoft ANF capacity pool storage.
- Quick recovery of an Oracle VLDB in the event of a failure on the same Azure database server VM.
- Quick recovery of an Oracle VLDB in the event of a failure on a standby Azure database server VM.

Audience

This solution is intended for the following people:

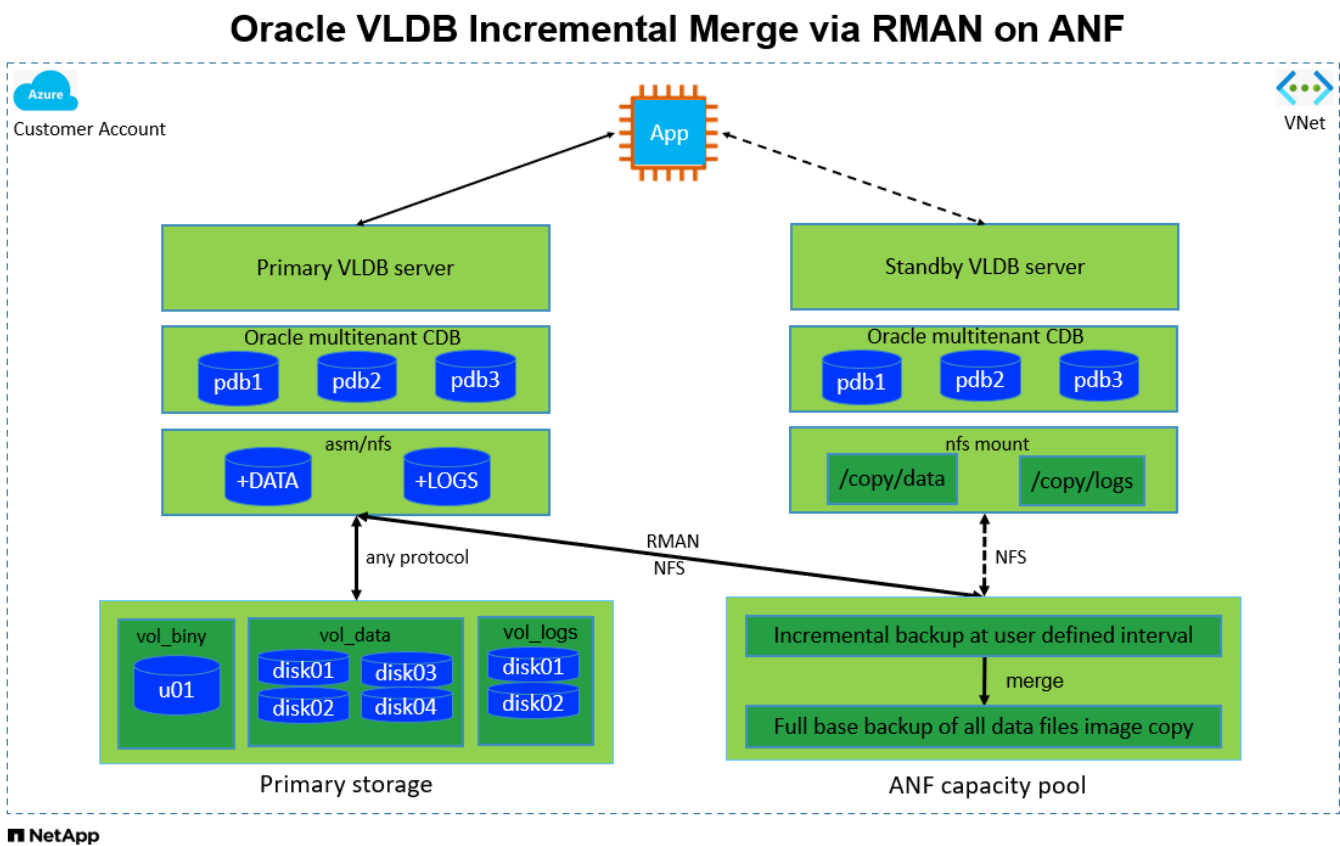
- A DBA who sets up Oracle VLDB image copy incremental merge via RMAN in Azure for faster database recovery.
- A database solution architect who tests Oracle workloads in the Azure public cloud.
- A storage administrator who manages Oracle databases deployed to ANF capacity pool storage.
- An application owner who would like to stand up Oracle databases in Azure cloud environment.

Solution test and validation environment

The testing and validation of this solution was performed in a Microsoft ANF capacity pool storage and Azure

VM compute environment that might not match the final deployment environment. For more information, see the section [Key factors for deployment consideration](#).

Architecture



Hardware and software components

Hardware		
ANF storage	Current version offered by Microsoft	2 TiB ANF capacity pool storage with Premium service level
Azure VM for DB server	Standard_B4ms - 4 vCPUs, 16GiB	2 VMs, one as primary DB server and the other as a standby
Software		
RedHat Linux	RHEL Linux 8.6 (LVM) - x64 Gen2	Deployed RedHat subscription for testing
Oracle Database	Version 19.18	Applied RU patch p34765931_190000_Linux-x86-64.zip
Oracle OPatch	Version 12.2.0.1.36	Latest patch p6880880_190000_Linux-x86-64.zip
NFS	Version 3.0	Oracle dNFS enabled

Key factors for deployment consideration

- **Oracle VLDB storage layout for RMAN incremental merge.** In our tests and validations, the NFS volume for Oracle incremental backup and merge is allocated from a single ANF capacity pool, which has 100 TiB per volume, and 1000 TiB total capacity limit. For deployment over the thresholds, multiple volumes, and ANF capacity pools can be concatenated in parallel with multiple NFS mount points to provide higher capacity.
- **Oracle recoverability using RMAN incremental merge.** The RMAN incremental backup and merge is generally executed at user defined frequency based on your RTO and RPO objectives. If there are total loss of primary data storage and/or archived logs, the data loss can occur. The Oracle database can be recovered up to last incremental backup that is available from ANF database backup image copy. To minimize the data loss, Oracle flash recovery area can be setup on ANF NFS mount point and archived logs are backed up to ANF NFS mount along with database image copy.
- **Running Oracle VLDB off ANF NFS file system.** Unlike other bulk storage for database backup, Microsoft ANF is a cloud enabled production grade storage that delivers high level of performance and storage efficiency. Once Oracle VLDB switches over from primary storage to image copy on ANF NFS file system, database performance can be maintained at high level while the primary storage failure is addressed. You can take comfort to know that user application experience does not suffer as the result of primary storage failure.
- **Azure compute instances.** In these tests and validations, we used Standard_B4ms Azure VMs as the Oracle database servers. There are other Azure VMs that may be optimized and better suited for database workload. You also need to size the Azure VM appropriately for the number of vCPUs and the amount of RAM based on actual workload requirements.
- **ANF capacity pool service level.** ANF capacity pool offers three service level: Standard, Premium, Ultra. By default, an auto QoS applies to a volume created within a capacity pool, which restricts the throughput on the volume. The throughput on a volume can be manually adjusted based on the size of capacity pool and service level.
- **dNFS configuration.** dNFS is built into Oracle kernel and is known to dramatically increase Oracle database performance when Oracle is deployed to NFS storage. dNFS is packaged into Oracle binary but is not turned on by default. It should be turned on for any Oracle database deployment on NFS. For multiple ANF capacity pools deployment for a VLDB, dNFS multi-paths to different ANF capacity pools storage should be properly configured.

Solution deployment

It is assumed that you already have your Oracle VLDB deployed in Azure cloud environment within a VNet. If you need help on Oracle deployment in Azure, please refer to following technical reports for help.

- [Simplified, Automated Oracle Deployment on Azure NetApp Files with NFS](#)
- [Oracle Database Deployment and Protection on Azure NetApp Files](#)

Your Oracle VLDB can be running either on an ANF storage or any other storage of choices within the Azure cloud ecosystem. The following section provides step-by-step deployment procedures for setting up RMAN incremental merge to an image copy of an Oracle VLDB that is staging in an NFS mount off ANF storage.

Prerequisites for deployment

Deployment requires the following prerequisites.

1. An Azure account has been set up, and the necessary Azure VNet and network segments have been created within your Azure account.
2. From the Azure portal console, you must deploy two Azure VM instances, one as the primary Oracle DB server and an optional standby DB server. See the architecture diagram in the previous section for more details about the environment setup. Also review the [Azure Virtual Machine series](#) for more information.
3. From the Azure portal console, deploy ANF storage to host the NFS volumes that stores the Oracle database standby image copy. If you are not familiar with the deployment of ANF, see the documentation [Quickstart: Set up Azure NetApp Files and create an NFS volume](#) for step-by-step instructions.



Ensure that you have allocated at least 128G in Azure VM root volume in order to have sufficient space to stage Oracle installation files.

Provision and export NFS volume to be mounted on primary Oracle VLDB server

In this section, we show provisioning an NFS volume from an ANF capacity pool via Azure portal console. Repeat the procedures on other ANF capacity pools if more than one ANF capacity pools are set up to accommodate the size of the database.


1. First, from Azure portal console, navigating to ANF capacity pool that is used to stage Oracle VLDB image copy.

2. From selected capacity pool - database, click Volumes and then, Add volume to launch add-volume workflow.








3. Fill in Volume name, Quota, Virtual network, and Delegated subnet to move to Protocol page.

Create a volume ...

Basics Protocol Tags Review + create

This page will help you create an Azure NetApp Files volume in your subscription and enable you to access the volume from within your virtual network. [Learn more about Azure NetApp Files](#) 

Volume details

Volume name *	<input type="text" value="ora-01-u02-copy"/> 
Available quota (GiB) ⓘ	<input type="text" value="748"/> 748 GiB
Quota (GiB) * ⓘ	<input type="text" value="500"/>  500 GiB
Available throughput (MiB/s) ⓘ	<input type="text" value="46.75"/>
Max. Throughput (MiB/s) ⓘ	<input type="text" value="31.25"/>
Enable Cool Access ⓘ	<input type="checkbox"/>
Coolness Period ⓘ	<input type="text" value="31"/>
Cool Access Retrieval Policy ⓘ	<input type="text" value="Default"/> 
Virtual network * ⓘ	<input type="text" value="ANFAVSVal (172.30.136.64/26,172.30.137.128/25,172.30.152.0/27)"/>  Create new virtual network
Delegated subnet * ⓘ	<input type="text" value="ANF_Sub (172.30.136.64/26)"/>  Create new subnet
Network features ⓘ	<input type="radio"/> Basic <input checked="" type="radio"/> Standard
Availability Zone ⓘ	<input type="text" value="None"/> 
Encryption key source ⓘ	<input type="text"/> 
Show advanced section	<input type="checkbox"/>

Review + create

[< Previous](#)

[Next : Protocol >](#)

- Take a note of the file path, enter allowed clients CIDR range, and enable `Root Access` for the volume.

Create a volume ...

Basics Protocol Tags Review + create

Configure access to your volume.

Access

Protocol type ☒ NFS ☐ SMB ☐ Dual-protocol

Configuration

File path * ⓘ

Versions * ▼

Kerberos ☐ Enabled ☒ Disabled

LDAP ☐ Enabled ☒ Disabled

Unix Permissions ⓘ

Azure VMware Solution DataStore ⓘ ☐

Export policy

Configure the volume's export policy. This can be edited later. [Learn more](#) ⓘ

↑ Move up ↓ Move down ↕ Move to top ⬇ Move to bottom 🗑 Delete

<input type="checkbox"/> Index	Allowed clients	Access	Root Access	Chown Mode
<input type="checkbox"/> 1	<input type="text" value="172.30.137.128/25,1"/> ✓	<input type="text" value="Read & Write"/> ▼	<input type="text" value="On"/> ▼	<input type="text" value="Restricted"/>
	<input type="text"/>	<input type="text"/> ▼	<input type="text"/> ▼	<input type="text"/>



[Review + create](#)

[< Previous](#)

[Next : Tags >](#)


5. Add a volume tag if desired.

Create a volume ...

Basics Protocol Tags Review + create

Tags are name/value pairs that enable you to categorize resources and view consolidated billing by applying the same tag to multiple resources and resource groups. [Learn more about tags](#) 

Note that if you create tags and then change resource settings on other tabs, your tags will be automatically updated.

Name ⓘ		Value ⓘ	
<input type="text" value="database"/>	:	<input type="text" value="oracle"/>	
<input type="text"/>	:	<input type="text"/>	

[Review + create](#)

[< Previous](#)

[Next : Review + create >](#)

6. Review and create the volume.

Create a volume ...

✓ Validation passed

Basics Protocol Tags Review + create

Basics

Subscription	Hybrid Cloud TME Onprem
Resource group	ANFAVSRG
Region	South Central US
Volume name	ora-01-u02-copy
Capacity pool	database
Service level	Premium
Quota	500 GiB
Encryption key source	None
Availability Zone	None

Networking

Virtual network	ANFAVSval (172.30.136.64/26,172.30.137.128/25,172.30.152.0/27)
Delegated subnet	ANF_Sub (172.30.136.64/26)
Network features	Standard

Protocol

Protocol	NFSv3
File path	ora-01-u02-copy
Unix Permissions	0770

Tags

database	oracle
----------	--------

Create

< Previous

Next >

[Download a template for automation](#)

7. Login to primary Oracle VLDB server as a user with sudo privilege and mount the NFS volume exported from ANF storage. Change to your ANF NFS server IP address and file path as necessary. The ANF NFS server IP address can be retrieved from ANF volume console page.

```
sudo mkdir /nfsanf
```

```
sudo mount 172.30.136.68:/ora-01-u02-copy /nfsanf -o  
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsiz=262144,wsiz=262144,noin  
tr
```

8. Change mount point ownership to oracle:oinstall, change to your oracle user name and primary group as necessary.

```
sudo chown oracle:oinstall /nfsanf
```

Setup Oracle RMAN incremental merge to image copy on ANF

RMAN incremental merge update the staging database data files image copy continuously at every incremental backup/merge interval. The image copy of database backup will be as up to date as the frequency you execute the incremental backup/merge. So, take into consideration of database performance, your RTO and RPO objectives when deciding the frequency of RMAN incremental backup and merge.

1. Login to primary Oracle VLDB server as oracle user.
2. Create an oracopy directory under mount point /nfsanf to store oracle data files image copies and archlog directory for Oracle flash recovery area.

```
mkdir /nfsanf/oracopy
```

```
mkdir /nfsanf/archlog
```

3. Login to Oracle database via sqlplus, enable block change tracking for faster incremental backup and change Oracle flash recovery area to ANF NFS mount if it is currently on primary storage. This allows the RMAN default control file/spfile autobackup and archived logs to be backed up to ANF NFS mount for recovery.

```
sqlplus / as sysdba
```

From sqlplus prompt, execute following command.

```
alter database enable block change tracking using file  
'/nfsanf/oracopy/bct_ntap1.ctf'
```

```
alter system set db_recovery_file_dest='/nfsanf/archlog/'  
scope=both;
```

Expected output:

```
[oracle@ora-01 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Wed Mar 20 16:44:21
2024
Version 19.18.0.0.0

Copyright (c) 1982, 2022, Oracle. All rights reserved.

Connected to:
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 -
Production
Version 19.18.0.0.0

SQL> alter database enable block change tracking using file
'/nfsanf/oracopy/bct_ntap1.ctf';

Database altered.

SQL> alter system set db_recovery_file_dest='/nfsanf/archlog/'
scope=both;

System altered.

SQL>
```

4. Create a RMAN backup and incremental merge script. The script allocates multiple channels for parallel RMAN backup and merge. First execution would generate the initial full baseline image copy. In a complete run, it first purges obsolete backups that are outside of retention window to keep staging area clean. It then switches current log file before merge and backup. The incremental backup follows the merge so that the database image copy is trailing current database state by one backup/merge cycle. The merge and backup order can be reversed for quicker recovery at user's preference. The RMAN script can be integrated into a simple shell script to be executed from crontab on the primary DB server. Ensure control file autobackup is on in RMAN setting.

```
vi /home/oracle/rman_bkup_merge.cmd
```

Add following lines:

RUN

```
{
    allocate channel c1 device type disk format '/nfsanf/oracopy/%U';
    allocate channel c2 device type disk format '/nfsanf/oracopy/%U';
    allocate channel c3 device type disk format '/nfsanf/oracopy/%U';
    allocate channel c4 device type disk format '/nfsanf/oracopy/%U';
    delete obsolete;
    sql 'alter system archive log current';
    recover copy of database with tag 'OraCopyBKUPonANF_level_0';
    backup incremental level 1 copies=1 for recover of copy with tag
'OraCopyBKUPonANF_level_0' database;
}
```

5. At the primary Oracle VLDB server, login to RMAN locally as oracle user with or without RMAN catalog. In this demonstration, we are not connecting to a RMAN catalog.

```
rman target / nocatalog;
```

output:

```
[oracle@ora-01 ~]$ rman target / nocatalog
```

```
Recovery Manager: Release 19.0.0.0.0 - Production on Wed Mar 20
16:54:24 2024
Version 19.18.0.0.0
```

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

```
connected to target database: NTAP1 (DBID=2441823937)
using target database control file instead of recovery catalog
```

6. From RMAN prompt, execute the script. First execution creates a baseline database image copy and subsequent executions merge and update the baseline image copy incrementally. The following is how to execute the script and the typical output. Set the number of channels to match the CPU cores on the host.

```
RMAN> @/home/oracle/rman_bkup_merge.cmd
```

```
RMAN> RUN
```

```

2> {
3>   allocate channel c1 device type disk format
'/nfsanf/oracopy/%U';
4>   allocate channel c2 device type disk format
'/nfsanf/oracopy/%U';
5>   allocate channel c3 device type disk format
'/nfsanf/oracopy/%U';
6>   allocate channel c4 device type disk format
'/nfsanf/oracopy/%U';
7>   delete obsolete;
8>   sql 'alter system archive log current';
9>   recover copy of database with tag 'OraCopyBKUPonANF_level_0';
10>  backup incremental level 1 copies=1 for recover of copy with
tag 'OraCopyBKUPonANF_level_0' database;
11> }
allocated channel: c1
channel c1: SID=142 device type=DISK

allocated channel: c2
channel c2: SID=277 device type=DISK

allocated channel: c3
channel c3: SID=414 device type=DISK

allocated channel: c4
channel c4: SID=28 device type=DISK

RMAN retention policy will be applied to the command
RMAN retention policy is set to redundancy 1
Deleting the following obsolete backups and copies:

```

Type	Key	Completion Time	Filename/Handle
Backup Set	1	18-MAR-24	
Backup Piece	1	18-MAR-24	/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163958359__04h19dgr_.bkp
Backup Set	2	18-MAR-24	
Backup Piece	2	18-MAR-24	/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163961675__0711m21g_.bkp
Backup Set	3	18-MAR-24	
Backup Piece	3	18-MAR-24	/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163962888__08p6y71x_.bkp
Backup Set	4	18-MAR-24	
Backup Piece	4	18-MAR-24	


```

/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163963796__09k8g1m
4_.bkp
Backup Set          5          18-MAR-24
  Backup Piece      5          18-MAR-24
/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163964697__0bd3tqg
3_.bkp
Backup Set          6          18-MAR-24
  Backup Piece      6          18-MAR-24
/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163965895__0chx6mz
t_.bkp
Backup Set          7          18-MAR-24
  Backup Piece      7          18-MAR-24
/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163966806__0dbyx34
4_.bkp
Backup Set          8          18-MAR-24
  Backup Piece      8          18-MAR-24
/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163968012__0fgvg80
5_.bkp
Backup Set          9          18-MAR-24
  Backup Piece      9          18-MAR-24
/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163968919__0g9x5t1
v_.bkp
Backup Set         10          18-MAR-24
  Backup Piece     10          18-MAR-24
/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163969821__0h4rfdz
j_.bkp
Backup Set         11          18-MAR-24
  Backup Piece     11          18-MAR-24
/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163971026__0j8o4wk
8_.bkp
Backup Set         12          18-MAR-24
  Backup Piece     12          18-MAR-24
/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163971931__0k3pnn2
o_.bkp
Backup Set         13          18-MAR-24
  Backup Piece     13          18-MAR-24
/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163972835__0kyg92t
1_.bkp
deleted backup piece
backup piece
handle=/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163963796__
09k8g1m4_.bkp RECID=4 STAMP=1163963804
deleted backup piece
backup piece
handle=/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163962888__
08p6y7lx_.bkp RECID=3 STAMP=1163962897

```

```
deleted backup piece
backup piece
handle=/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163961675__
0711m2lg_.bkp RECID=2 STAMP=1163961683
deleted backup piece
backup piece
handle=/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163958359__
04h19dgr_.bkp RECID=1 STAMP=1163958361
deleted backup piece
backup piece
handle=/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163964697__
0bd3tqg3_.bkp RECID=5 STAMP=1163964705
deleted backup piece
backup piece
handle=/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163965895__
0chx6mzt_.bkp RECID=6 STAMP=1163965906
deleted backup piece
backup piece
handle=/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163966806__
0dbyx344_.bkp RECID=7 STAMP=1163966814
deleted backup piece
backup piece
handle=/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163968012__
0fgvg805_.bkp RECID=8 STAMP=1163968018
deleted backup piece
backup piece
handle=/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163968919__
0g9x5t1v_.bkp RECID=9 STAMP=1163968926
deleted backup piece
backup piece
handle=/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163969821__
0h4rfdzj_.bkp RECID=10 STAMP=1163969827
Deleted 3 objects

deleted backup piece
backup piece
handle=/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163971026__
0j8o4wk8_.bkp RECID=11 STAMP=1163971032
Deleted 3 objects

deleted backup piece
backup piece
handle=/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163971931__
0k3pnn2o_.bkp RECID=12 STAMP=1163971938
Deleted 3 objects
```

```
deleted backup piece
backup piece
handle=/u03/orareco/NTAP1/autobackup/2024_03_18/o1_mf_s_1163972835_
0kyg92t1_.bkp RECID=13 STAMP=1163972837
Deleted 4 objects
```

```
sql statement: alter system archive log current
```

```
Starting recover at 20-MAR-24
no copy of datafile 1 found to recover
no copy of datafile 3 found to recover
no copy of datafile 4 found to recover
.
.
no copy of datafile 31 found to recover
no copy of datafile 32 found to recover
Finished recover at 20-MAR-24
```

```
Starting backup at 20-MAR-24
no parent backup or copy of datafile 1 found
no parent backup or copy of datafile 3 found
no parent backup or copy of datafile 4 found
.
.
no parent backup or copy of datafile 19 found
no parent backup or copy of datafile 20 found
channel c1: starting datafile copy
input datafile file number=00021
name=/u02/oradata/NTAP1/NTAP1_pdb1/soe_01.dbf
channel c2: starting datafile copy
input datafile file number=00022
name=/u02/oradata/NTAP1/NTAP1_pdb1/soe_02.dbf
channel c3: starting datafile copy
input datafile file number=00023
name=/u02/oradata/NTAP1/NTAP1_pdb1/soe_03.dbf
channel c4: starting datafile copy
input datafile file number=00024
name=/u02/oradata/NTAP1/NTAP1_pdb1/soe_04.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SOE_FNO-22_0g2m6brl tag=ORACOPYBKUPONANF_LEVEL_0 RECID=4
STAMP=1164132108
channel c2: datafile copy complete, elapsed time: 01:06:39
channel c2: starting datafile copy
input datafile file number=00025
name=/u02/oradata/NTAP1/NTAP1_pdb1/soe_05.dbf
```

```
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-  
SOE_FNO-24_0i2m6brl tag=ORACOPYBKUPONANF_LEVEL_0 RECID=5  
STAMP=1164132121  
channel c4: datafile copy complete, elapsed time: 01:06:45  
channel c4: starting datafile copy  
input datafile file number=00026  
name=/u02/oradata/NTAP1/NTAP1_pdb1/soe_06.dbf  
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-  
SOE_FNO-23_0h2m6brl tag=ORACOPYBKUPONANF_LEVEL_0 RECID=6  
STAMP=1164132198  
channel c3: datafile copy complete, elapsed time: 01:08:05  
channel c3: starting datafile copy  
input datafile file number=00027  
name=/u02/oradata/NTAP1/NTAP1_pdb1/soe_07.dbf  
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-  
SOE_FNO-21_0f2m6brl tag=ORACOPYBKUPONANF_LEVEL_0 RECID=7  
STAMP=1164132248  
channel c1: datafile copy complete, elapsed time: 01:08:57  
channel c1: starting datafile copy  
input datafile file number=00028  
name=/u02/oradata/NTAP1/NTAP1_pdb1/soe_08.dbf  
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-  
SOE_FNO-25_0j2m6fol tag=ORACOPYBKUPONANF_LEVEL_0 RECID=9  
STAMP=1164136123  
channel c2: datafile copy complete, elapsed time: 01:06:46  
channel c2: starting datafile copy  
input datafile file number=00029  
name=/u02/oradata/NTAP1/NTAP1_pdb1/soe_09.dbf  
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-  
SOE_FNO-26_0k2m6fot tag=ORACOPYBKUPONANF_LEVEL_0 RECID=8  
STAMP=1164136113  
channel c4: datafile copy complete, elapsed time: 01:06:36  
channel c4: starting datafile copy  
input datafile file number=00030  
name=/u02/oradata/NTAP1/NTAP1_pdb1/soe_10.dbf  
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-  
SOE_FNO-27_0l2m6frc tag=ORACOPYBKUPONANF_LEVEL_0 RECID=10  
STAMP=1164136293  
channel c3: datafile copy complete, elapsed time: 01:08:10  
channel c3: starting datafile copy  
input datafile file number=00031  
name=/u02/oradata/NTAP1/NTAP1_pdb1/soe_11.dbf  
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-  
SOE_FNO-28_0m2m6fsu tag=ORACOPYBKUPONANF_LEVEL_0 RECID=11  
STAMP=1164136333  
channel c1: datafile copy complete, elapsed time: 01:07:52
```

```
channel c1: starting datafile copy
input datafile file number=00032
name=/u02/oradata/NTAP1/NTAP1_pdb1/soe_12.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SOE_FNO-29_0n2m6jlr tag=ORACOPYBKUPONANF_LEVEL_0 RECID=12
STAMP=1164140082
channel c2: datafile copy complete, elapsed time: 01:06:01
channel c2: starting datafile copy
input datafile file number=00001
name=/u02/oradata/NTAP1/system01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SOE_FNO-30_0o2m6jlr tag=ORACOPYBKUPONANF_LEVEL_0 RECID=13
STAMP=1164140190
channel c4: datafile copy complete, elapsed time: 01:07:49
channel c4: starting datafile copy
input datafile file number=00003
name=/u02/oradata/NTAP1/sysaux01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSTEM_FNO-1_0r2m6nhk tag=ORACOPYBKUPONANF_LEVEL_0 RECID=14
STAMP=1164140240
channel c2: datafile copy complete, elapsed time: 00:02:38
channel c2: starting datafile copy
input datafile file number=00004
name=/u02/oradata/NTAP1/undotbs01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
UNDOTBS1_FNO-4_0t2m6nml tag=ORACOPYBKUPONANF_LEVEL_0 RECID=15
STAMP=1164140372
channel c2: datafile copy complete, elapsed time: 00:02:15
channel c2: starting datafile copy
input datafile file number=00011
name=/u02/oradata/NTAP1/NTAP1_pdb1/undotbs01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSAux_FNO-3_0s2m6nl1 tag=ORACOPYBKUPONANF_LEVEL_0 RECID=16
STAMP=1164140377
channel c4: datafile copy complete, elapsed time: 00:03:01
channel c4: starting datafile copy
input datafile file number=00010
name=/u02/oradata/NTAP1/NTAP1_pdb1/sysaux01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SOE_FNO-32_0q2m6jsi tag=ORACOPYBKUPONANF_LEVEL_0 RECID=17
STAMP=1164140385
channel c1: datafile copy complete, elapsed time: 01:07:29
channel c1: starting datafile copy
input datafile file number=00014
name=/u02/oradata/NTAP1/NTAP1_pdb2/sysaux01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
```

```

SOE_FNO-31_0p2m6jrb tag=ORACOPYBKUPONANF_LEVEL_0 RECID=18
STAMP=1164140406
channel c3: datafile copy complete, elapsed time: 01:08:31
channel c3: starting datafile copy
input datafile file number=00018
name=/u02/oradata/NTAP1/NTAP1_pdb3/sysaux01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSAUX_FNO-10_0v2m6nqs tag=ORACOPYBKUPONANF_LEVEL_0 RECID=19
STAMP=1164140459
channel c4: datafile copy complete, elapsed time: 00:01:26
channel c4: starting datafile copy
input datafile file number=00006
name=/u02/oradata/NTAP1/pdbseed/sysaux01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSAUX_FNO-14_102m6nr3 tag=ORACOPYBKUPONANF_LEVEL_0 RECID=20
STAMP=1164140468
channel c1: datafile copy complete, elapsed time: 00:01:22
channel c1: starting datafile copy
input datafile file number=00009
name=/u02/oradata/NTAP1/NTAP1_pdb1/system01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
UNDOTBS1_FNO-11_0u2m6nqs tag=ORACOPYBKUPONANF_LEVEL_0 RECID=21
STAMP=1164140471
channel c2: datafile copy complete, elapsed time: 00:01:33
channel c2: starting datafile copy
input datafile file number=00013
name=/u02/oradata/NTAP1/NTAP1_pdb2/system01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSAUX_FNO-18_112m6nrt tag=ORACOPYBKUPONANF_LEVEL_0 RECID=22
STAMP=1164140476
channel c3: datafile copy complete, elapsed time: 00:00:57
channel c3: starting datafile copy
input datafile file number=00017
name=/u02/oradata/NTAP1/NTAP1_pdb3/system01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSAUX_FNO-6_122m6nti tag=ORACOPYBKUPONANF_LEVEL_0 RECID=23
STAMP=1164140488
channel c4: datafile copy complete, elapsed time: 00:00:25
channel c4: starting datafile copy
input datafile file number=00005
name=/u02/oradata/NTAP1/pdbseed/system01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSTEM_FNO-13_142m6ntp tag=ORACOPYBKUPONANF_LEVEL_0 RECID=24
STAMP=1164140532
channel c2: datafile copy complete, elapsed time: 00:01:06
channel c2: starting datafile copy

```

```

input datafile file number=00008
name=/u02/oradata/NTAP1/pdbseed/undotbs01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSTEM_FNO-17_152m6nts tag=ORACOPYBKUPONANF_LEVEL_0 RECID=25
STAMP=1164140539
channel c3: datafile copy complete, elapsed time: 00:01:03
channel c3: starting datafile copy
input datafile file number=00015
name=/u02/oradata/NTAP1/NTAP1_pdb2/undotbs01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSTEM_FNO-9_132m6ntm tag=ORACOPYBKUPONANF_LEVEL_0 RECID=26
STAMP=1164140541
channel c1: datafile copy complete, elapsed time: 00:01:13
channel c1: starting datafile copy
input datafile file number=00019
name=/u02/oradata/NTAP1/NTAP1_pdb3/undotbs01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSTEM_FNO-5_162m6nuc tag=ORACOPYBKUPONANF_LEVEL_0 RECID=27
STAMP=1164140541
channel c4: datafile copy complete, elapsed time: 00:00:41
channel c4: starting datafile copy
input datafile file number=00007 name=/u02/oradata/NTAP1/users01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
UNDOTBS1_FNO-8_172m6nvr tag=ORACOPYBKUPONANF_LEVEL_0 RECID=28
STAMP=1164140552
channel c2: datafile copy complete, elapsed time: 00:00:16
channel c2: starting datafile copy
input datafile file number=00012
name=/u02/oradata/NTAP1/NTAP1_pdb1/users01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
UNDOTBS1_FNO-15_182m6nvs tag=ORACOPYBKUPONANF_LEVEL_0 RECID=30
STAMP=1164140561
channel c3: datafile copy complete, elapsed time: 00:00:24
channel c3: starting datafile copy
input datafile file number=00016
name=/u02/oradata/NTAP1/NTAP1_pdb2/users01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
USERS_FNO-7_1a2m6o01 tag=ORACOPYBKUPONANF_LEVEL_0 RECID=29
STAMP=1164140560
channel c4: datafile copy complete, elapsed time: 00:00:16
channel c4: starting datafile copy
input datafile file number=00020
name=/u02/oradata/NTAP1/NTAP1_pdb3/users01.dbf
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
UNDOTBS1_FNO-19_192m6nvv tag=ORACOPYBKUPONANF_LEVEL_0 RECID=31
STAMP=1164140564

```

```

channel c1: datafile copy complete, elapsed time: 00:00:21
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
USERS_FNO-12_1b2m6o0e tag=ORACOPYBKUPONANF_LEVEL_0 RECID=32
STAMP=1164140564
channel c2: datafile copy complete, elapsed time: 00:00:02
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
USERS_FNO-16_1c2m6o0k tag=ORACOPYBKUPONANF_LEVEL_0 RECID=34
STAMP=1164140565
channel c3: datafile copy complete, elapsed time: 00:00:01
output file name=/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
USERS_FNO-20_1d2m6o0k tag=ORACOPYBKUPONANF_LEVEL_0 RECID=33
STAMP=1164140565
channel c4: datafile copy complete, elapsed time: 00:00:01
Finished backup at 20-MAR-24

Starting Control File and SPFILE Autobackup at 20-MAR-24
piece
handle=/nfsanf/archlog/NTAP1/autobackup/2024_03_20/o1_mf_s_116414056
5__5g56ypks_.bkp comment=NONE
Finished Control File and SPFILE Autobackup at 20-MAR-24
released channel: c1
released channel: c2
released channel: c3
released channel: c4

RMAN> **end-of-file**

RMAN>

```

7. List database image copy after backup to observe that a database image copy has been created in ANF NFS mount point.

```

RMAN> list copy of database tag 'OraCopyBKUPonANF_level_0';

List of Datafile Copies
=====

Key       File S Completion Time Ckp SCN      Ckp Time      Sparse
-----
14        1    A 20-MAR-24      4161498      20-MAR-24      NO
      Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
      SYSTEM_FNO-1_0r2m6nhk
      Tag: ORACOPYBKUPONANF_LEVEL_0

16        3    A 20-MAR-24      4161568      20-MAR-24      NO
      Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-

```



```

SYS_AUX_FNO-3_0s2m6nl1
    Tag: ORACOPYBKUPONANF_LEVEL_0

15      4      A 20-MAR-24      4161589      20-MAR-24      NO
    Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
UNDOTBS1_FNO-4_0t2m6nml
    Tag: ORACOPYBKUPONANF_LEVEL_0

27      5      A 20-MAR-24      2379694      18-MAR-24      NO
    Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSTEM_FNO-5_162m6nuc
    Tag: ORACOPYBKUPONANF_LEVEL_0
    Container ID: 2, PDB Name: PDB$SEED

23      6      A 20-MAR-24      2379694      18-MAR-24      NO
    Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYS_AUX_FNO-6_122m6nti
    Tag: ORACOPYBKUPONANF_LEVEL_0
    Container ID: 2, PDB Name: PDB$SEED

29      7      A 20-MAR-24      4161872      20-MAR-24      NO
    Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
USERS_FNO-7_1a2m6o01
    Tag: ORACOPYBKUPONANF_LEVEL_0

28      8      A 20-MAR-24      2379694      18-MAR-24      NO
    Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
UNDOTBS1_FNO-8_172m6nvr
    Tag: ORACOPYBKUPONANF_LEVEL_0
    Container ID: 2, PDB Name: PDB$SEED

26      9      A 20-MAR-24      4161835      20-MAR-24      NO
    Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSTEM_FNO-9_132m6ntm
    Tag: ORACOPYBKUPONANF_LEVEL_0
    Container ID: 3, PDB Name: NTAP1_PDB1

19      10     A 20-MAR-24      4161784      20-MAR-24      NO
    Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYS_AUX_FNO-10_0v2m6nqs
    Tag: ORACOPYBKUPONANF_LEVEL_0
    Container ID: 3, PDB Name: NTAP1_PDB1

21      11     A 20-MAR-24      4161780      20-MAR-24      NO
    Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
UNDOTBS1_FNO-11_0u2m6nqs

```

```

Tag: ORACOPYBKUPONANF_LEVEL_0
Container ID: 3, PDB Name: NTAP1_PDB1

32      12      A 20-MAR-24      4161880      20-MAR-24      NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
USERS_FNO-12_1b2m6o0e
Tag: ORACOPYBKUPONANF_LEVEL_0
Container ID: 3, PDB Name: NTAP1_PDB1

24      13      A 20-MAR-24      4161838      20-MAR-24      NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSTEM_FNO-13_142m6ntp
Tag: ORACOPYBKUPONANF_LEVEL_0
Container ID: 4, PDB Name: NTAP1_PDB2

20      14      A 20-MAR-24      4161785      20-MAR-24      NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSAUX_FNO-14_102m6nr3
Tag: ORACOPYBKUPONANF_LEVEL_0
Container ID: 4, PDB Name: NTAP1_PDB2

30      15      A 20-MAR-24      4161863      20-MAR-24      NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
UNDOTBS1_FNO-15_182m6nvs
Tag: ORACOPYBKUPONANF_LEVEL_0
Container ID: 4, PDB Name: NTAP1_PDB2

34      16      A 20-MAR-24      4161884      20-MAR-24      NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
USERS_FNO-16_1c2m6o0k
Tag: ORACOPYBKUPONANF_LEVEL_0
Container ID: 4, PDB Name: NTAP1_PDB2

25      17      A 20-MAR-24      4161841      20-MAR-24      NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSTEM_FNO-17_152m6nts
Tag: ORACOPYBKUPONANF_LEVEL_0
Container ID: 5, PDB Name: NTAP1_PDB3

22      18      A 20-MAR-24      4161810      20-MAR-24      NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSAUX_FNO-18_112m6nrt
Tag: ORACOPYBKUPONANF_LEVEL_0
Container ID: 5, PDB Name: NTAP1_PDB3

31      19      A 20-MAR-24      4161869      20-MAR-24      NO

```

Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
 UNDOTBS1_FNO-19_192m6nvv
 Tag: ORACOPYBKUPONANF_LEVEL_0
 Container ID: 5, PDB Name: NTAP1_PDB3

33 20 A 20-MAR-24 4161887 20-MAR-24 NO
 Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
 USERS_FNO-20_1d2m6o0k
 Tag: ORACOPYBKUPONANF_LEVEL_0
 Container ID: 5, PDB Name: NTAP1_PDB3

7 21 A 20-MAR-24 4152514 20-MAR-24 NO
 Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-
 21_0f2m6brl
 Tag: ORACOPYBKUPONANF_LEVEL_0
 Container ID: 3, PDB Name: NTAP1_PDB1

4 22 A 20-MAR-24 4152518 20-MAR-24 NO
 Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-
 22_0g2m6brl
 Tag: ORACOPYBKUPONANF_LEVEL_0
 Container ID: 3, PDB Name: NTAP1_PDB1

6 23 A 20-MAR-24 4152522 20-MAR-24 NO
 Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-
 23_0h2m6brl
 Tag: ORACOPYBKUPONANF_LEVEL_0
 Container ID: 3, PDB Name: NTAP1_PDB1

5 24 A 20-MAR-24 4152529 20-MAR-24 NO
 Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-
 24_0i2m6brl
 Tag: ORACOPYBKUPONANF_LEVEL_0
 Container ID: 3, PDB Name: NTAP1_PDB1

9 25 A 20-MAR-24 4156120 20-MAR-24 NO
 Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-
 25_0j2m6fol
 Tag: ORACOPYBKUPONANF_LEVEL_0
 Container ID: 3, PDB Name: NTAP1_PDB1

8 26 A 20-MAR-24 4156130 20-MAR-24 NO
 Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-
 26_0k2m6fot
 Tag: ORACOPYBKUPONANF_LEVEL_0
 Container ID: 3, PDB Name: NTAP1_PDB1

```

10      27      A 20-MAR-24      4156159      20-MAR-24      NO
      Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-
27_012m6frc
      Tag: ORACOPYBKUPONANF_LEVEL_0
      Container ID: 3, PDB Name: NTAP1_PDB1

11      28      A 20-MAR-24      4156183      20-MAR-24      NO
      Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-
28_0m2m6fsu
      Tag: ORACOPYBKUPONANF_LEVEL_0
      Container ID: 3, PDB Name: NTAP1_PDB1

12      29      A 20-MAR-24      4158795      20-MAR-24      NO
      Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-
29_0n2m6jlr
      Tag: ORACOPYBKUPONANF_LEVEL_0
      Container ID: 3, PDB Name: NTAP1_PDB1

13      30      A 20-MAR-24      4158803      20-MAR-24      NO
      Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-
30_0o2m6jlr
      Tag: ORACOPYBKUPONANF_LEVEL_0
      Container ID: 3, PDB Name: NTAP1_PDB1

18      31      A 20-MAR-24      4158871      20-MAR-24      NO
      Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-
31_0p2m6jrb
      Tag: ORACOPYBKUPONANF_LEVEL_0
      Container ID: 3, PDB Name: NTAP1_PDB1

17      32      A 20-MAR-24      4158886      20-MAR-24      NO
      Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-
32_0q2m6jsi
      Tag: ORACOPYBKUPONANF_LEVEL_0
      Container ID: 3, PDB Name: NTAP1_PDB1

```

8. Report schema from Oracle RMAN command prompt to observe that current VLDB data files are on primary storage.

```

RMAN> report schema;

Report of database schema for database with db_unique_name NTAP1

List of Permanent Datafiles
=====
File Size(MB) Tablespace          RB segs Datafile Name

```

```

-----
1      1060      SYSTEM      YES
/u02/oradata/NTAP1/system01.dbf
3      1000      SYSAUX      NO
/u02/oradata/NTAP1/sysaux01.dbf
4      695      UNDOTBS1     YES
/u02/oradata/NTAP1/undotbs01.dbf
5      400      PDB$SEED:SYSTEM    NO
/u02/oradata/NTAP1/pdbseed/system01.dbf
6      440      PDB$SEED:SYSAUX    NO
/u02/oradata/NTAP1/pdbseed/sysaux01.dbf
7      5        USERS      NO
/u02/oradata/NTAP1/users01.dbf
8      235      PDB$SEED:UNDOTBS1    NO
/u02/oradata/NTAP1/pdbseed/undotbs01.dbf
9      410      NTAP1_PDB1:SYSTEM    YES
/u02/oradata/NTAP1/NTAP1_pdb1/system01.dbf
10     520      NTAP1_PDB1:SYSAUX    NO
/u02/oradata/NTAP1/NTAP1_pdb1/sysaux01.dbf
11     580      NTAP1_PDB1:UNDOTBS1    YES
/u02/oradata/NTAP1/NTAP1_pdb1/undotbs01.dbf
12     5        NTAP1_PDB1:USERS    NO
/u02/oradata/NTAP1/NTAP1_pdb1/users01.dbf
13     410      NTAP1_PDB2:SYSTEM    YES
/u02/oradata/NTAP1/NTAP1_pdb2/system01.dbf
14     500      NTAP1_PDB2:SYSAUX    NO
/u02/oradata/NTAP1/NTAP1_pdb2/sysaux01.dbf
15     235      NTAP1_PDB2:UNDOTBS1    YES
/u02/oradata/NTAP1/NTAP1_pdb2/undotbs01.dbf
16     5        NTAP1_PDB2:USERS    NO
/u02/oradata/NTAP1/NTAP1_pdb2/users01.dbf
17     410      NTAP1_PDB3:SYSTEM    YES
/u02/oradata/NTAP1/NTAP1_pdb3/system01.dbf
18     500      NTAP1_PDB3:SYSAUX    NO
/u02/oradata/NTAP1/NTAP1_pdb3/sysaux01.dbf
19     235      NTAP1_PDB3:UNDOTBS1    YES
/u02/oradata/NTAP1/NTAP1_pdb3/undotbs01.dbf
20     5        NTAP1_PDB3:USERS    NO
/u02/oradata/NTAP1/NTAP1_pdb3/users01.dbf
21     31744    NTAP1_PDB1:SOE      NO
/u02/oradata/NTAP1/NTAP1_pdb1/soe_01.dbf
22     31744    NTAP1_PDB1:SOE      NO
/u02/oradata/NTAP1/NTAP1_pdb1/soe_02.dbf
23     31744    NTAP1_PDB1:SOE      NO
/u02/oradata/NTAP1/NTAP1_pdb1/soe_03.dbf
24     31744    NTAP1_PDB1:SOE      NO

```

```

/u02/oradata/NTAP1/NTAP1_pdb1/soe_04.dbf
25    31744    NTAP1_PDB1:SOE        NO
/u02/oradata/NTAP1/NTAP1_pdb1/soe_05.dbf
26    31744    NTAP1_PDB1:SOE        NO
/u02/oradata/NTAP1/NTAP1_pdb1/soe_06.dbf
27    31744    NTAP1_PDB1:SOE        NO
/u02/oradata/NTAP1/NTAP1_pdb1/soe_07.dbf
28    31744    NTAP1_PDB1:SOE        NO
/u02/oradata/NTAP1/NTAP1_pdb1/soe_08.dbf
29    31744    NTAP1_PDB1:SOE        NO
/u02/oradata/NTAP1/NTAP1_pdb1/soe_09.dbf
30    31744    NTAP1_PDB1:SOE        NO
/u02/oradata/NTAP1/NTAP1_pdb1/soe_10.dbf
31    31744    NTAP1_PDB1:SOE        NO
/u02/oradata/NTAP1/NTAP1_pdb1/soe_11.dbf
32    31744    NTAP1_PDB1:SOE        NO
/u02/oradata/NTAP1/NTAP1_pdb1/soe_12.dbf

```

List of Temporary Files

=====

File	Size (MB)	Tablespace	Maxsize (MB)	Tempfile Name
1	123	TEMP	32767	/u02/oradata/NTAP1/temp01.dbf
2	123	PDB\$SEED:TEMP	32767	/u02/oradata/NTAP1/pdbseed/temp012024-03-18_16-07-32-463-PM.dbf
3	31744	NTAP1_PDB1:TEMP	32767	/u02/oradata/NTAP1/NTAP1_pdb1/temp01.dbf
4	123	NTAP1_PDB2:TEMP	32767	/u02/oradata/NTAP1/NTAP1_pdb2/temp01.dbf
5	123	NTAP1_PDB3:TEMP	32767	/u02/oradata/NTAP1/NTAP1_pdb3/temp01.dbf
6	31744	NTAP1_PDB1:TEMP	31744	/u02/oradata/NTAP1/NTAP1_pdb1/temp02.dbf

RMAN>

9. Validate database image copy from OS NFS mount point.

```

[oracle@ora-01 ~]$ ls -l /nfsanf/oracopy
total 399482176
-rw-r----- 1 oracle oinstall 11600384 Mar 20 21:44 bct_ntap1.ctf
-rw-r----- 1 oracle oinstall 33286004736 Mar 20 18:03 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-21_0f2m6brl
-rw-r----- 1 oracle oinstall 33286004736 Mar 20 18:01 data_D-

```

```

NTAP1_I-2441823937_TS-SOE_FNO-22_0g2m6brl
-rw-r----- 1 oracle oinstall 33286004736 Mar 20 18:03 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-23_0h2m6brl
-rw-r----- 1 oracle oinstall 33286004736 Mar 20 18:02 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-24_0i2m6brl
-rw-r----- 1 oracle oinstall 33286004736 Mar 20 19:08 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-25_0j2m6fol
-rw-r----- 1 oracle oinstall 33286004736 Mar 20 19:08 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-26_0k2m6fot
-rw-r----- 1 oracle oinstall 33286004736 Mar 20 19:11 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-27_0l2m6frc
-rw-r----- 1 oracle oinstall 33286004736 Mar 20 19:12 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-28_0m2m6fsu
-rw-r----- 1 oracle oinstall 33286004736 Mar 20 20:14 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-29_0n2m6jlr
-rw-r----- 1 oracle oinstall 33286004736 Mar 20 20:16 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-30_0o2m6jlr
-rw-r----- 1 oracle oinstall 33286004736 Mar 20 20:20 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-31_0p2m6jrb
-rw-r----- 1 oracle oinstall 33286004736 Mar 20 20:19 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-32_0q2m6jsi
-rw-r----- 1 oracle oinstall 545267712 Mar 20 20:20 data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-10_0v2m6nqs
-rw-r----- 1 oracle oinstall 524296192 Mar 20 20:21 data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-14_102m6nr3
-rw-r----- 1 oracle oinstall 524296192 Mar 20 20:21 data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-18_112m6nrt
-rw-r----- 1 oracle oinstall 1048584192 Mar 20 20:19 data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-3_0s2m6nl1
-rw-r----- 1 oracle oinstall 461381632 Mar 20 20:21 data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-6_122m6nti
-rw-r----- 1 oracle oinstall 1111498752 Mar 20 20:17 data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-1_0r2m6nhk
-rw-r----- 1 oracle oinstall 429924352 Mar 20 20:22 data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-13_142m6ntp
-rw-r----- 1 oracle oinstall 429924352 Mar 20 20:22 data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-17_152m6nts
-rw-r----- 1 oracle oinstall 419438592 Mar 20 20:22 data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-5_162m6nuc
-rw-r----- 1 oracle oinstall 429924352 Mar 20 20:22 data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-9_132m6ntm
-rw-r----- 1 oracle oinstall 608182272 Mar 20 20:21 data_D-
NTAP1_I-2441823937_TS-UNDOTBS1_FNO-11_0u2m6nqs
-rw-r----- 1 oracle oinstall 246423552 Mar 20 20:22 data_D-
NTAP1_I-2441823937_TS-UNDOTBS1_FNO-15_182m6nvs
-rw-r----- 1 oracle oinstall 246423552 Mar 20 20:22 data_D-

```

```
NTAP1_I-2441823937_TS-UNDOTBS1_FNO-19_192m6nvv
-rw-r----- 1 oracle oinstall 728768512 Mar 20 20:19 data_D-
NTAP1_I-2441823937_TS-UNDOTBS1_FNO-4_0t2m6nml
-rw-r----- 1 oracle oinstall 246423552 Mar 20 20:22 data_D-
NTAP1_I-2441823937_TS-UNDOTBS1_FNO-8_172m6nvr
-rw-r----- 1 oracle oinstall 5251072 Mar 20 20:22 data_D-
NTAP1_I-2441823937_TS-USERS_FNO-12_1b2m6o0e
-rw-r----- 1 oracle oinstall 5251072 Mar 20 20:22 data_D-
NTAP1_I-2441823937_TS-USERS_FNO-16_1c2m6o0k
-rw-r----- 1 oracle oinstall 5251072 Mar 20 20:22 data_D-
NTAP1_I-2441823937_TS-USERS_FNO-20_1d2m6o0k
-rw-r----- 1 oracle oinstall 5251072 Mar 20 20:22 data_D-
NTAP1_I-2441823937_TS-USERS_FNO-7_1a2m6o01
[oracle@ora-01 ~]$
```

This completes the setup of an Oracle VLDB standby image copy backup and merge.

Switch Oracle VLDB to image copy for quick recovery

In the event of a failure due to primary storage issue such as data loss or corruption, database can be quickly switched over to image copy on ANF NFS mount and recovered to current state without database restore. Eliminating media restoration speeds up the database recovery tremendously for a VLDB. This use case assumes that the Oracle VLDB DB server is intact and database control file, archived and current logs are all available for recovery.

1. Login to Azure primary VLDB server host as oracle user and create a test table before switch over.

```
[oracle@ora-01 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Thu Mar 21 15:13:52
2024
Version 19.18.0.0.0

Copyright (c) 1982, 2022, Oracle. All rights reserved.

Connected to:
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 -
Production
Version 19.18.0.0.0

SQL> show pdbs

          CON_ID CON_NAME                                OPEN MODE  RESTRICTED
          -----
          2  PDB$SEED                                     READ ONLY  NO
          3  NTAP1_PDB1                                   READ WRITE NO
          4  NTAP1_PDB2                                   READ WRITE NO
          5  NTAP1_PDB3                                   READ WRITE NO
SQL> alter session set container=ntap1_pdb1;

Session altered.

SQL> create table test (id integer, dt timestamp, event
varchar(100));

Table created.

SQL> insert into test values(1, sysdate, 'test oracle incremental
merge switch to copy');

1 row created.

SQL> commit;
```

```
Commit complete.
```

```
SQL> select * from test;
```

```
          ID
-----
DT
-----
EVENT
-----
          1
21-MAR-24 03.15.03.000000 PM
test oracle incremental merge switch to copy
```

2. Simulate a failure by shutdown abort database, then start up oracle in mount stage.

```
SQL> shutdown abort;
ORACLE instance shut down.
SQL> startup mount;
ORACLE instance started.
```

```
Total System Global Area 6442449688 bytes
Fixed Size                  9177880 bytes
Variable Size              1325400064 bytes
Database Buffers           5100273664 bytes
Redo Buffers                7598080 bytes
Database mounted.
SQL> exit
```

3. As oracle user, connect to Oracle database via RMAN to switch database to copy.

```
[oracle@ora-01 ~]$ rman target / nocatalog

Recovery Manager: Release 19.0.0.0.0 - Production on Thu Mar 21
15:20:58 2024
Version 19.18.0.0.0

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

connected to target database: NTAP1 (DBID=2441823937, not open)
using target database control file instead of recovery catalog
```

```
RMAN> switch database to copy;
```

```
datafile 1 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-1_0r2m6nhk"
datafile 3 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-3_0s2m6nl1"
datafile 4 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-UNDOTBS1_FNO-4_0t2m6nml"
datafile 5 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-5_162m6nuc"
datafile 6 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-6_122m6nti"
datafile 7 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-USERS_FNO-7_1a2m6o01"
datafile 8 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-UNDOTBS1_FNO-8_172m6nvr"
datafile 9 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-9_132m6ntm"
datafile 10 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-10_0v2m6nqs"
datafile 11 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-UNDOTBS1_FNO-11_0u2m6nqs"
datafile 12 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-USERS_FNO-12_1b2m6o0e"
datafile 13 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-13_142m6ntp"
datafile 14 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-14_102m6nr3"
datafile 15 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-UNDOTBS1_FNO-15_182m6nvs"
datafile 16 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-USERS_FNO-16_1c2m6o0k"
datafile 17 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-17_152m6nts"
datafile 18 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-18_112m6nrt"
datafile 19 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-UNDOTBS1_FNO-19_192m6nvv"
datafile 20 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-USERS_FNO-20_1d2m6o0k"
datafile 21 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-21_0f2m6brl"
datafile 22 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-22_0g2m6brl"
datafile 23 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-23_0h2m6brl"
```

```

datafile 24 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-24_0i2m6brl"
datafile 25 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-25_0j2m6fol"
datafile 26 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-26_0k2m6fot"
datafile 27 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-27_0l2m6frc"
datafile 28 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-28_0m2m6fsu"
datafile 29 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-29_0n2m6jlr"
datafile 30 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-30_0o2m6jlr"
datafile 31 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-31_0p2m6jrb"
datafile 32 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-32_0q2m6jsi"

```

4. Recover and open database to bring it up to current from last incremental backup.

```

RMAN> recover database;

```

```

Starting recover at 21-MAR-24
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=392 device type=DISK
channel ORA_DISK_1: starting incremental datafile backup set restore
channel ORA_DISK_1: specifying datafile(s) to restore from backup
set
destination for restore of datafile 00009: /nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-9_0qlsd7cm
destination for restore of datafile 00023: /nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-23_04lsd6s5
destination for restore of datafile 00027: /nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-27_08lsd70i
destination for restore of datafile 00031: /nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-31_0c1sd74u
destination for restore of datafile 00034: /nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-34_0f1sd788
channel ORA_DISK_1: reading from backup piece
/nfsanf/oracopy/321sfous_98_1_1
channel ORA_DISK_1: piece handle=/nfsanf/oracopy/321sfous_98_1_1
tag=ORACOPYBKUPONANF_LEVEL_0
channel ORA_DISK_1: restored backup piece 1
channel ORA_DISK_1: restore complete, elapsed time: 00:00:01

```

```

channel ORA_DISK_1: starting incremental datafile backup set restore
channel ORA_DISK_1: specifying datafile(s) to restore from backup
set
destination for restore of datafile 00010: /nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-10_0k1sd7bb
destination for restore of datafile 00021: /nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-21_021sd6pv
destination for restore of datafile 00025: /nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-25_061sd6uc
.
.
.
channel ORA_DISK_1: starting incremental datafile backup set restore
channel ORA_DISK_1: specifying datafile(s) to restore from backup
set
destination for restore of datafile 00016: /nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-USERS_FNO-16_121sd7dn
channel ORA_DISK_1: reading from backup piece
/nfsanf/oracopy/3i1sfov0_114_1_1
channel ORA_DISK_1: piece handle=/nfsanf/oracopy/3i1sfov0_114_1_1
tag=ORACOPYBKUPONANF_LEVEL_0
channel ORA_DISK_1: restored backup piece 1
channel ORA_DISK_1: restore complete, elapsed time: 00:00:01
channel ORA_DISK_1: starting incremental datafile backup set restore
channel ORA_DISK_1: specifying datafile(s) to restore from backup
set
destination for restore of datafile 00020: /nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-USERS_FNO-20_131sd7do
channel ORA_DISK_1: reading from backup piece
/nfsanf/oracopy/3j1sfov0_115_1_1
channel ORA_DISK_1: piece handle=/nfsanf/oracopy/3j1sfov0_115_1_1
tag=ORACOPYBKUPONANF_LEVEL_0
channel ORA_DISK_1: restored backup piece 1
channel ORA_DISK_1: restore complete, elapsed time: 00:00:01

starting media recovery
media recovery complete, elapsed time: 00:00:01

Finished recover at 21-MAR-24

RMAN> alter database open;

Statement processed

RMAN>

```

5. Check database structure from sqlplus after recovery to observe that all VLDB data files with exception of control, temp, and current log files are now switched over to copy on ANF NFS file system.

```
SQL> select name from v$datafile
2 union
3 select name from v$tempfile
4 union
5 select name from v$controlfile
6 union
7* select member from v$logfile
SQL> /
```

NAME

```
-----
-----
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-21_0f2m6brl
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-22_0g2m6brl
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-23_0h2m6brl
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-24_0i2m6brl
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-25_0j2m6fol
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-26_0k2m6fot
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-27_0l2m6frc
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-28_0m2m6fsu
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-29_0n2m6jlr
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-30_0o2m6jlr
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-31_0p2m6jrb
```

NAME

```
-----
-----
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-32_0q2m6jsi
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSAUX_FNO-10_0v2m6nqs
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSAUX_FNO-14_102m6nr3
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSAUX_FNO-18_112m6nrt
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSAUX_FNO-3_0s2m6nl1
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSAUX_FNO-6_122m6nti
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSTEM_FNO-13_142m6ntp
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSTEM_FNO-17_152m6nts
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSTEM_FNO-1_0r2m6nhk
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSTEM_FNO-5_162m6nuc
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSTEM_FNO-9_132m6ntm
```

NAME

```

/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-UNDOTBS1_FNO-
11_0u2m6nqs
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-UNDOTBS1_FNO-
15_182m6nvs
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-UNDOTBS1_FNO-
19_192m6nvv
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-UNDOTBS1_FNO-4_0t2m6nml
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-UNDOTBS1_FNO-8_172m6nvr
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-USERS_FNO-12_1b2m6o0e
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-USERS_FNO-16_1c2m6o0k
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-USERS_FNO-20_1d2m6o0k
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-USERS_FNO-7_1a2m6o01
/u02/oradata/NTAP1/NTAP1_pdb1/temp01.dbf
/u02/oradata/NTAP1/NTAP1_pdb1/temp02.dbf

```

NAME

```

-----
-----
/u02/oradata/NTAP1/NTAP1_pdb2/temp01.dbf
/u02/oradata/NTAP1/NTAP1_pdb3/temp01.dbf
/u02/oradata/NTAP1/control01.ctl
/u02/oradata/NTAP1/pdbseed/temp012024-03-18_16-07-32-463-PM.dbf
/u02/oradata/NTAP1/temp01.dbf
/u03/orareco/NTAP1/control02.ctl
/u03/orareco/NTAP1/onlinelog/redo01.log
/u03/orareco/NTAP1/onlinelog/redo02.log
/u03/orareco/NTAP1/onlinelog/redo03.log

```

42 rows selected.

6. From SQL plus, check the content of test table we have inserted before the switch over to copy.

```
SQL> alter session set container=ntap1_pdb1;
```

```
Session altered.
```

```
SQL> select * from test;
```

```
          ID
-----
DT
-----
EVENT
-----
          1
21-MAR-24 03.15.03.000000 PM
test oracle incremental merge switch to copy
```

```
SQL>
```

7. You could run the Oracle VLDB in ANF NFS mount for an extended period of time while maintaining expected performance level. When the primary storage issue is fixed, you can swing back to it by reversing the incremental backup merge processes with minimal downtime.

Oracle VLDB recovery from image copy to a standby DB server

In the event of a failure where both the primary storage and primary DB server host are lost, recovery cannot be performed from the original server. However, your Oracle database backup image copy available on the ANF NFS file system comes in handy. You can quickly recover the primary database to a standby DB server if one is available, using the backup image copy. In this section, we will demonstrate the step-by-step procedures for such recovery.

1. Insert a row to test table we have created previously for Oracle VLDB restoring to alternative host validation.

```
SQL> insert into test values(2, sysdate, 'test recovery on a new
Azure VM host with image copy on ANF');
```

```
1 row created.
```

```
SQL> commit;
```

```
Commit complete.
```

```
SQL> select * from test;
```

```
          ID
-----
DT
-----
EVENT
-----
          1
21-MAR-24 03.15.03.000000 PM
test oracle incremental merge switch to copy

          2
22-MAR-24 02.22.06.000000 PM
test recovery on a new Azure VM host with image copy on ANF
```

```
          ID
-----
DT
-----
EVENT
-----
```

```
SQL>
```

2. As oracle user, run RMAN incremental backup and merge to flush the transaction to backup set on ANF NFS mount.

```
[oracle@ip-172-30-15-99 ~]$ rman target / nocatalog

Recovery Manager: Release 19.0.0.0.0 - Production on Tue May 30
17:26:03 2023
Version 19.18.0.0.0

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

connected to target database: NTAP1 (DBID=2441823937)
using target database control file instead of recovery catalog

RMAN> @rman_bkup_merge.cmd
```

3. Shutdown primary VLDB server host to simulate a total failure of storage and DB server host.
4. On the standby DB server ora-02 with same OS and version, OS kernel should be patched up as primary VLDB server host. Also, the same version and patches of Oracle has been installed and configured on standby DB server with software only option.
5. Configure oracle environment similarly to primary VLDB server ora_01, such as oratab, and oracle user .bash_profile etc. It is a good practice to backup those files to ANF NFS mount point.
6. The Oracle database backup image copy on ANF NFS file system is then mounted on the standby DB server for recovery. The following procedures demonstrate the process details.

As azueruser, create the mount point.

```
sudo mkdir /nfsanf
```

As azureuser, mount the NFS volume that stored Oracle VLDB backup image copy.

```
sudo mount 172.30.136.68:/ora-01-u02-copy /nfsanf -o
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsiz=262144,wsiz=262144,noi
tr
```

7. Validate the Oracle database backup image copy on ANF NFS mount point.

```
[oracle@ora-02 ~]$ ls -ltr /nfsanf/oracopy/
total 400452728
-rw-r-----. 1 oracle oinstall 461381632 Mar 21 23:47 data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-6_242m9oan
-rw-r-----. 1 oracle oinstall 419438592 Mar 21 23:49 data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-5_282m9oem
-rw-r-----. 1 oracle oinstall 246423552 Mar 21 23:49 data_D-
NTAP1_I-2441823937_TS-UNDOTBS1_FNO-8_292m9oem
```

```

-rw-r-----. 1 oracle oinstall      21438464 Mar 22 14:35
2h2mbccv_81_1_1
-rw-r-----. 1 oracle oinstall      17956864 Mar 22 14:35
2i2mbcd0_82_1_1
-rw-r-----. 1 oracle oinstall      17956864 Mar 22 14:35
2j2mbcd1_83_1_1
-rw-r-----. 1 oracle oinstall      15245312 Mar 22 14:35
2k2mbcd3_84_1_1
-rw-r-----. 1 oracle oinstall      1638400 Mar 22 14:35
2m2mbcdn_86_1_1
-rw-r-----. 1 oracle oinstall      40042496 Mar 22 14:35
2l2mbcdn_85_1_1
-rw-r-----. 1 oracle oinstall      21856256 Mar 22 14:35
2n2mbcdo_87_1_1
-rw-r-----. 1 oracle oinstall      3710976 Mar 22 14:35
2o2mbcdv_88_1_1
-rw-r-----. 1 oracle oinstall      3416064 Mar 22 14:35
2p2mbcdv_89_1_1
-rw-r-----. 1 oracle oinstall      2596864 Mar 22 14:35
2r2mbce0_91_1_1
-rw-r-----. 1 oracle oinstall      2531328 Mar 22 14:35
2s2mbce1_92_1_1
-rw-r-----. 1 oracle oinstall      4718592 Mar 22 14:35
2v2mbce2_95_1_1
-rw-r-----. 1 oracle oinstall      4243456 Mar 22 14:35
302mbce2_96_1_1
-rw-r-----. 1 oracle oinstall      57344 Mar 22 14:35
312mbce3_97_1_1
-rw-r-----. 1 oracle oinstall      57344 Mar 22 14:35
322mbce3_98_1_1
-rw-r-----. 1 oracle oinstall      57344 Mar 22 14:35
332mbce3_99_1_1
-rw-r-----. 1 oracle oinstall    608182272 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-UNDOTBS1_FNO-11_202m9o22
-rw-r-----. 1 oracle oinstall 33286004736 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-30_1q2m9k7a
-rw-r-----. 1 oracle oinstall   555753472 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-10_212m9o52
-rw-r-----. 1 oracle oinstall 33286004736 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-26_1m2m9g9j
-rw-r-----. 1 oracle oinstall 33286004736 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-27_1n2m9gcg
-rw-r-----. 1 oracle oinstall   429924352 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-9_252m9oc5
-rw-r-----. 1 oracle oinstall 33286004736 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-22_1i2m9cap

```

```

-rw-r-----. 1 oracle oinstall 33286004736 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-23_1j2m9cap
-rw-r-----. 1 oracle oinstall      5251072 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-USERS_FNO-12_2d2m9ofs
-rw-r-----. 1 oracle oinstall 33286004736 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-28_1o2m9gd4
-rw-r-----. 1 oracle oinstall 33286004736 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-31_1r2m9kfk
-rw-r-----. 1 oracle oinstall 33286004736 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-29_1p2m9ju6
-rw-r-----. 1 oracle oinstall 33286004736 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-32_1s2m9kgg
-rw-r-----. 1 oracle oinstall 33286004736 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-25_1l2m9g3u
-rw-r-----. 1 oracle oinstall 33286004736 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-24_1k2m9cap
-rw-r-----. 1 oracle oinstall 33286004736 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SOE_FNO-21_1h2m9cap
-rw-r-----. 1 oracle oinstall 1121984512 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-1_1t2m9nij
-rw-r-----. 1 oracle oinstall 1142956032 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-3_1u2m9nog
-rw-r-----. 1 oracle oinstall 728768512 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-UNDOTBS1_FNO-4_1v2m9nu6
-rw-r-----. 1 oracle oinstall 534781952 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-14_222m9o53
-rw-r-----. 1 oracle oinstall 534781952 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-18_232m9oa8
-rw-r-----. 1 oracle oinstall 429924352 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-13_262m9oca
-rw-r-----. 1 oracle oinstall 246423552 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-UNDOTBS1_FNO-15_2a2m9of6
-rw-r-----. 1 oracle oinstall 429924352 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-17_272m9oel
-rw-r-----. 1 oracle oinstall      5251072 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-USERS_FNO-7_2c2m9ofn
-rw-r-----. 1 oracle oinstall      5251072 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-USERS_FNO-16_2e2m9og8
-rw-r-----. 1 oracle oinstall 246423552 Mar 22 15:31 data_D-
NTAP1_I-2441823937_TS-UNDOTBS1_FNO-19_2b2m9ofn
-rw-r-----. 1 oracle oinstall      5251072 Mar 22 15:32 data_D-
NTAP1_I-2441823937_TS-USERS_FNO-20_2f2m9og8
-rw-r-----. 1 oracle oinstall 76546048 Mar 22 15:37
362mbft5_102_1_1
-rw-r-----. 1 oracle oinstall 14671872 Mar 22 15:37
392mbgli_105_1_1

```

```

-rw-r-----. 1 oracle oinstall      79462400 Mar 22 15:37
372mbftb_103_1_1
-rw-r-----. 1 oracle oinstall        917504 Mar 22 15:37
3a2mbg23_106_1_1
-rw-r-----. 1 oracle oinstall    428498944 Mar 22 15:37
352mbfst_101_1_1
-rw-r-----. 1 oracle oinstall    88702976 Mar 22 15:37
382mbftm_104_1_1
-rw-r-----. 1 oracle oinstall    5021696 Mar 22 15:37
3b2mbg2b_107_1_1
-rw-r-----. 1 oracle oinstall    278528 Mar 22 15:38
3c2mbg2f_108_1_1
-rw-r-----. 1 oracle oinstall    278528 Mar 22 15:38
3d2mbg2i_109_1_1
-rw-r-----. 1 oracle oinstall    425984 Mar 22 15:38
3f2mbg2m_111_1_1
-rw-r-----. 1 oracle oinstall    442368 Mar 22 15:38
3g2mbg2q_112_1_1
-rw-r-----. 1 oracle oinstall    278528 Mar 22 15:38
3j2mbg37_115_1_1
-rw-r-----. 1 oracle oinstall    270336 Mar 22 15:38
3k2mbg3a_116_1_1
-rw-r-----. 1 oracle oinstall    57344 Mar 22 15:38
3l2mbg3f_117_1_1
-rw-r-----. 1 oracle oinstall    57344 Mar 22 15:38
3n2mbg3k_119_1_1
-rw-r-----. 1 oracle oinstall    57344 Mar 22 15:38
3m2mbg3g_118_1_1
-rw-r-----. 1 oracle oinstall    11600384 Mar 22 15:52 bct_ntap1.ctf
[oracle@ora-02 ~]$

```

8. Verify the available Oracle archived logs on the ANF NFS mount for recovery and note the last log file log sequency number. In this case, it is 10. Our recovery point is up to log sequency number 11.

```
[oracle@ora-02 ~]$ ls -ltr
/nfsanf/archlog/NTAP1/archivelog/2024_03_22
total 1429548
-r--r-----. 1 oracle oinstall 176650752 Mar 22 12:00
o1_mf_1_2__9m198x6t_.arc
-r--r-----. 1 oracle oinstall 17674752 Mar 22 14:34
o1_mf_1_3__9vn701r5_.arc
-r--r-----. 1 oracle oinstall 188782080 Mar 22 15:20
o1_mf_1_4__9y6gn5co_.arc
-r--r-----. 1 oracle oinstall 183638016 Mar 22 15:21
o1_mf_1_5__9y7p68s6_.arc
-r--r-----. 1 oracle oinstall 193106944 Mar 22 15:21
o1_mf_1_6__9y8ygtss_.arc
-r--r-----. 1 oracle oinstall 179439104 Mar 22 15:22
o1_mf_1_7__9ybjdp55_.arc
-r--r-----. 1 oracle oinstall 198815232 Mar 22 15:23
o1_mf_1_8__9yctxjgy_.arc
-r--r-----. 1 oracle oinstall 185494528 Mar 22 15:24
o1_mf_1_9__9yfrj0b1_.arc
-r--r-----. 1 oracle oinstall 134470144 Mar 22 15:29
o1_mf_1_10__9yomybbc_.arc
[oracle@ora-02 ~]$
```

9. As oracle user, set ORACLE_HOME variable to current Oracle installation on standby DB server ora-02, ORACLE_SID to primary Oracle instance SID. In this case, it is NTAP1.

```
[oracle@ora-02 ~]$ export
ORACLE_HOME=/u01/app/oracle/product/19.0.0/NTAP2
[oracle@ora-02 ~]$ export ORACLE_SID=NTAP1
[oracle@ora-02 ~]$ export PATH=$PATH:$ORACLE_HOME/bin
```

10. As oracle user, create a generic Oracle init file in \$ORACLE_HOME/dbs directory with proper admin directories configured. Most importantly, have Oracle flash recovery area point to ANF NFS mount path as defined in primary Oracle VLDB server. flash recovery area configuration is demonstrated in section Setup Oracle RMAN incremental merge to image copy on ANF. Set the Oracle control file to ANF NFS file system.

```
vi $ORACLE_HOME/dbs/initNTAP1.ora
```

With following example entries:

```

*.audit_file_dest='/u01/app/oracle/admin/NTAP1/adump'
*.audit_trail='db'
*.compatible='19.0.0'
*.control_files=('/nfsanf/oracopy/NTAP1.ctl')
*.db_block_size=8192
*.db_create_file_dest='/nfsanf/oracopy/'
*.db_domain='solutions.netapp.com'
*.db_name='NTAP1'
*.db_recovery_file_dest_size=85899345920
*.db_recovery_file_dest='/nfsanf/archlog/'
*.diagnostic_dest='/u01/app/oracle'
*.dispatchers='(PROTOCOL=TCP) (SERVICE=NTAP1XDB) '
*.enable_pluggable_database=true
*.local_listener='LISTENER'
*.nls_language='AMERICAN'
*.nls_territory='AMERICA'
*.open_cursors=300
*.pga_aggregate_target=1024m
*.processes=320
*.remote_login_passwordfile='EXCLUSIVE'
*.sga_target=10240m
*.undo_tablespace='UNDOTBS1'

```

The above init file should be replaced by restored backup init file from primary Oracle VLDB server in the case of discrepancy.

11. As oracle user, launch RMAN to run Oracle recovery on the standby DB server host. First, start the Oracle instance in nomount state.


```
[oracle@ora-02 ~]$ rman target / nocatalog

Recovery Manager: Release 19.0.0.0.0 - Production on Fri Mar 22
16:02:55 2024
Version 19.18.0.0.0

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

connected to target database (not started)

RMAN> startup nomount;

Oracle instance started

Total System Global Area      10737418000 bytes

Fixed Size                     9174800 bytes
Variable Size                  1577058304 bytes
Database Buffers               9126805504 bytes
Redo Buffers                    24379392 bytes
```

12. Set database ID. The database ID can be retrieved from Oracle file name of image copy on ANF NFS mount point.

```
RMAN> set dbid = 2441823937;

executing command: SET DBID
```

13. Restore controlfile from autobackup. If Oracle controlfile and spfile autobackup is enabled, they are backed up in every incremental backup and merge cycle. The latest backup will be restored if multiple copies are available.

```

RMAN> restore controlfile from autobackup;

Starting restore at 22-MAR-24
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=2 device type=DISK

recovery area destination: /nfsanf/archlog/
database name (or database unique name) used for search: NTAP1
channel ORA_DISK_1: AUTOBACKUP
/nfsanf/archlog/NTAP1/autobackup/2024_03_22/o1_mf_s_1164296325__9z77
zyxb_.bkp found in the recovery area
channel ORA_DISK_1: looking for AUTOBACKUP on day: 20240322
channel ORA_DISK_1: restoring control file from AUTOBACKUP
/nfsanf/archlog/NTAP1/autobackup/2024_03_22/o1_mf_s_1164296325__9z77
zyxb_.bkp
channel ORA_DISK_1: control file restore from AUTOBACKUP complete
output file name=/nfsanf/oracopy/NTAP1.ctl
Finished restore at 22-MAR-24

```

14. Restore init file from spfile to a /tmp folder for updating parameter file later to match with primary VLDB.

```

RMAN> restore spfile to pfile '/tmp/archive/initNTAP1.ora' from
autobackup;

Starting restore at 22-MAR-24
using channel ORA_DISK_1

recovery area destination: /nfsanf/archlog/
database name (or database unique name) used for search: NTAP1
channel ORA_DISK_1: AUTOBACKUP
/nfsanf/archlog/NTAP1/autobackup/2024_03_22/o1_mf_s_1164296325__9z77
zyxb_.bkp found in the recovery area
channel ORA_DISK_1: looking for AUTOBACKUP on day: 20240322
channel ORA_DISK_1: restoring spfile from AUTOBACKUP
/nfsanf/archlog/NTAP1/autobackup/2024_03_22/o1_mf_s_1164296325__9z77
zyxb_.bkp
channel ORA_DISK_1: SPFILE restore from AUTOBACKUP complete
Finished restore at 22-MAR-24

```

15. Mount control file and validate the database backup image copy.

```

RMAN> alter database mount;

```

released channel: ORA_DISK_1

Statement processed

RMAN> list copy of database tag 'ORACOPYBKUPONANF_LEVEL_0';

List of Datafile Copies

=====

Key	File	S	Completion Time	Ckp SCN	Ckp Time	Sparse
-----	----	-	-----	-----	-----	-----
82	1	A	22-MAR-24	4598427	22-MAR-24	NO
	Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSTEM_FNO-1_1t2m9nij					
	Tag: ORACOPYBKUPONANF_LEVEL_0					
83	3	A	22-MAR-24	4598423	22-MAR-24	NO
	Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSAUX_FNO-3_1u2m9nog					
	Tag: ORACOPYBKUPONANF_LEVEL_0					
84	4	A	22-MAR-24	4598431	22-MAR-24	NO
	Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-UNDOTBS1_FNO-4_1v2m9nu6					
	Tag: ORACOPYBKUPONANF_LEVEL_0					
58	5	A	21-MAR-24	2379694	18-MAR-24	NO
	Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSTEM_FNO-5_282m9oem					
	Tag: ORACOPYBKUPONANF_LEVEL_0					
	Container ID: 2, PDB Name: PDB\$SEED					
52	6	A	21-MAR-24	2379694	18-MAR-24	NO
	Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSAUX_FNO-6_242m9oan					
	Tag: ORACOPYBKUPONANF_LEVEL_0					
	Container ID: 2, PDB Name: PDB\$SEED					
90	7	A	22-MAR-24	4598462	22-MAR-24	NO
	Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-USERS_FNO-7_2c2m9ofn					
	Tag: ORACOPYBKUPONANF_LEVEL_0					
59	8	A	21-MAR-24	2379694	18-MAR-24	NO
	Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-UNDOTBS1_FNO-8_292m9oem					
	Tag: ORACOPYBKUPONANF_LEVEL_0					

Container ID: 2, PDB Name: PDB\$SEED

```
71      9      A 22-MAR-24      4598313      22-MAR-24      NO
      Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSTEM_FNO-9_252m9oc5
      Tag: ORACOPYBKUPONANF_LEVEL_0
      Container ID: 3, PDB Name: NTAP1_PDB1

68     10     A 22-MAR-24      4598308      22-MAR-24      NO
      Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYS_AUX_FNO-10_212m9o52
      Tag: ORACOPYBKUPONANF_LEVEL_0
      Container ID: 3, PDB Name: NTAP1_PDB1

66     11     A 22-MAR-24      4598304      22-MAR-24      NO
      Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
UNDOTBS1_FNO-11_202m9o22
      Tag: ORACOPYBKUPONANF_LEVEL_0
      Container ID: 3, PDB Name: NTAP1_PDB1

74     12     A 22-MAR-24      4598318      22-MAR-24      NO
      Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
USERS_FNO-12_2d2m9ofs
      Tag: ORACOPYBKUPONANF_LEVEL_0
      Container ID: 3, PDB Name: NTAP1_PDB1

86     13     A 22-MAR-24      4598445      22-MAR-24      NO
      Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSTEM_FNO-13_262m9oca
      Tag: ORACOPYBKUPONANF_LEVEL_0
      Container ID: 4, PDB Name: NTAP1_PDB2

85     14     A 22-MAR-24      4598437      22-MAR-24      NO
      Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYS_AUX_FNO-14_222m9o53
      Tag: ORACOPYBKUPONANF_LEVEL_0
      Container ID: 4, PDB Name: NTAP1_PDB2

87     15     A 22-MAR-24      4598454      22-MAR-24      NO
      Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
UNDOTBS1_FNO-15_2a2m9of6
      Tag: ORACOPYBKUPONANF_LEVEL_0
      Container ID: 4, PDB Name: NTAP1_PDB2

89     16     A 22-MAR-24      4598466      22-MAR-24      NO
      Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
```

USERS_FNO-16_2e2m9og8
Tag: ORACOPYBKUPONANF_LEVEL_0
Container ID: 4, PDB Name: NTAP1_PDB2

91 17 A 22-MAR-24 4598450 22-MAR-24 NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSTEM_FNO-17_272m9oel
Tag: ORACOPYBKUPONANF_LEVEL_0
Container ID: 5, PDB Name: NTAP1_PDB3

88 18 A 22-MAR-24 4598441 22-MAR-24 NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
SYSAUX_FNO-18_232m9oa8
Tag: ORACOPYBKUPONANF_LEVEL_0
Container ID: 5, PDB Name: NTAP1_PDB3

92 19 A 22-MAR-24 4598458 22-MAR-24 NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
UNDOTBS1_FNO-19_2b2m9ofn
Tag: ORACOPYBKUPONANF_LEVEL_0
Container ID: 5, PDB Name: NTAP1_PDB3

93 20 A 22-MAR-24 4598470 22-MAR-24 NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-
USERS_FNO-20_2f2m9og8
Tag: ORACOPYBKUPONANF_LEVEL_0
Container ID: 5, PDB Name: NTAP1_PDB3

81 21 A 22-MAR-24 4598318 22-MAR-24 NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-
21_1h2m9cap
Tag: ORACOPYBKUPONANF_LEVEL_0
Container ID: 3, PDB Name: NTAP1_PDB1

72 22 A 22-MAR-24 4598304 22-MAR-24 NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-
22_1i2m9cap
Tag: ORACOPYBKUPONANF_LEVEL_0
Container ID: 3, PDB Name: NTAP1_PDB1

73 23 A 22-MAR-24 4598308 22-MAR-24 NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-
23_1j2m9cap
Tag: ORACOPYBKUPONANF_LEVEL_0
Container ID: 3, PDB Name: NTAP1_PDB1

80	24	A	22-MAR-24	4598313	22-MAR-24	NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-						
24_1k2m9cap						
Tag: ORACOPYBKUPONANF_LEVEL_0						
Container ID: 3, PDB Name: NTAP1_PDB1						
79	25	A	22-MAR-24	4598318	22-MAR-24	NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-						
25_112m9g3u						
Tag: ORACOPYBKUPONANF_LEVEL_0						
Container ID: 3, PDB Name: NTAP1_PDB1						
69	26	A	22-MAR-24	4598304	22-MAR-24	NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-						
26_1m2m9g9j						
Tag: ORACOPYBKUPONANF_LEVEL_0						
Container ID: 3, PDB Name: NTAP1_PDB1						
70	27	A	22-MAR-24	4598308	22-MAR-24	NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-						
27_1n2m9gcg						
Tag: ORACOPYBKUPONANF_LEVEL_0						
Container ID: 3, PDB Name: NTAP1_PDB1						
75	28	A	22-MAR-24	4598313	22-MAR-24	NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-						
28_1o2m9gd4						
Tag: ORACOPYBKUPONANF_LEVEL_0						
Container ID: 3, PDB Name: NTAP1_PDB1						
77	29	A	22-MAR-24	4598318	22-MAR-24	NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-						
29_1p2m9ju6						
Tag: ORACOPYBKUPONANF_LEVEL_0						
Container ID: 3, PDB Name: NTAP1_PDB1						
67	30	A	22-MAR-24	4598304	22-MAR-24	NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-						
30_1q2m9k7a						
Tag: ORACOPYBKUPONANF_LEVEL_0						
Container ID: 3, PDB Name: NTAP1_PDB1						
76	31	A	22-MAR-24	4598308	22-MAR-24	NO
Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-						
31_1r2m9kfk						
Tag: ORACOPYBKUPONANF_LEVEL_0						

Container ID: 3, PDB Name: NTAP1_PDB1

```
78      32    A 22-MAR-24      4598313      22-MAR-24      NO
      Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-
32_1s2m9kgg
      Tag: ORACOPYBKUPONANF_LEVEL_0
      Container ID: 3, PDB Name: NTAP1_PDB1
```

16. Switch database to copy to run recovery without database restore.

```
RMAN> switch database to copy;
```

```
Starting implicit crosscheck backup at 22-MAR-24
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=12 device type=DISK
Crosschecked 33 objects
Finished implicit crosscheck backup at 22-MAR-24
```

```
Starting implicit crosscheck copy at 22-MAR-24
using channel ORA_DISK_1
Crosschecked 31 objects
Finished implicit crosscheck copy at 22-MAR-24
```

```
searching for all files in the recovery area
cataloging files...
cataloging done
```

```
List of Cataloged Files
```

```
=====
```

```
File Name:
```

```
/nfsanf/archlog/NTAP1/autobackup/2024_03_20/o1_mf_s_1164140565__5g56
ypks_.bkp
```

```
File Name:
```

```
/nfsanf/archlog/NTAP1/autobackup/2024_03_22/o1_mf_s_1164296325__9z77
zyxb_.bkp
```

```
datafile 1 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-1_1t2m9nij"
datafile 3 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-3_1u2m9nog"
datafile 4 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-UNDOTBS1_FNO-4_1v2m9nu6"
datafile 5 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SYSTEM_FNO-5_282m9oem"
datafile 6 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SYSAUX_FNO-6_242m9oan"
```

datafile 7 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-USERS_FNO-7_2c2m9ofn"
datafile 8 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-UNDOTBS1_FNO-8_292m9oem"
datafile 9 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSTEM_FNO-9_252m9oc5"
datafile 10 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSAUX_FNO-10_212m9o52"
datafile 11 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-UNDOTBS1_FNO-11_202m9o22"
datafile 12 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-USERS_FNO-12_2d2m9ofs"
datafile 13 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSTEM_FNO-13_262m9oca"
datafile 14 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSAUX_FNO-14_222m9o53"
datafile 15 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-UNDOTBS1_FNO-15_2a2m9of6"
datafile 16 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-USERS_FNO-16_2e2m9og8"
datafile 17 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSTEM_FNO-17_272m9oel"
datafile 18 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSAUX_FNO-18_232m9oa8"
datafile 19 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-UNDOTBS1_FNO-19_2b2m9ofn"
datafile 20 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-USERS_FNO-20_2f2m9og8"
datafile 21 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-21_1h2m9cap"
datafile 22 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-22_1i2m9cap"
datafile 23 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-23_1j2m9cap"
datafile 24 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-24_1k2m9cap"
datafile 25 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-25_1l2m9g3u"
datafile 26 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-26_1m2m9g9j"
datafile 27 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-27_1n2m9gcg"
datafile 28 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-28_1o2m9gd4"
datafile 29 switched to datafile copy "/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-29_1p2m9ju6"


```
datafile 30 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-30_1q2m9k7a"
datafile 31 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-31_1r2m9kfk"
datafile 32 switched to datafile copy "/nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-32_1s2m9kgg"
```

17. Run Oracle recovery up to last available archive log in flash recovery area.

```
RMAN> run {
2> set until sequence=11;
3> recover database;
4> }
```

executing command: SET until clause

Starting recover at 22-MAR-24
using channel ORA_DISK_1

starting media recovery

```
archived log for thread 1 with sequence 4 is already on disk as file
/nfsanf/archlog/NTAP1/archivelog/2024_03_22/o1_mf_1_4__9y6gn5co_.arc
archived log for thread 1 with sequence 5 is already on disk as file
/nfsanf/archlog/NTAP1/archivelog/2024_03_22/o1_mf_1_5__9y7p68s6_.arc
archived log for thread 1 with sequence 6 is already on disk as file
/nfsanf/archlog/NTAP1/archivelog/2024_03_22/o1_mf_1_6__9y8ygtss_.arc
archived log for thread 1 with sequence 7 is already on disk as file
/nfsanf/archlog/NTAP1/archivelog/2024_03_22/o1_mf_1_7__9ybjdp55_.arc
archived log for thread 1 with sequence 8 is already on disk as file
/nfsanf/archlog/NTAP1/archivelog/2024_03_22/o1_mf_1_8__9yctxjgy_.arc
archived log for thread 1 with sequence 9 is already on disk as file
/nfsanf/archlog/NTAP1/archivelog/2024_03_22/o1_mf_1_9__9yfrj0b1_.arc
archived log for thread 1 with sequence 10 is already on disk as
file
/nfsanf/archlog/NTAP1/archivelog/2024_03_22/o1_mf_1_10__9yomybbc_.ar
c
archived log file
name=/nfsanf/archlog/NTAP1/archivelog/2024_03_22/o1_mf_1_4__9y6gn5co
_.arc thread=1 sequence=4
archived log file
name=/nfsanf/archlog/NTAP1/archivelog/2024_03_22/o1_mf_1_5__9y7p68s6
_.arc thread=1 sequence=5
archived log file
name=/nfsanf/archlog/NTAP1/archivelog/2024_03_22/o1_mf_1_6__9y8ygtss
```

```

_.arc thread=1 sequence=6
archived log file
name=/nfsanf/archlog/NTAP1/archivelog/2024_03_22/o1_mf_1_7__9ybjdp55
_.arc thread=1 sequence=7
archived log file
name=/nfsanf/archlog/NTAP1/archivelog/2024_03_22/o1_mf_1_8__9yctxjgy
_.arc thread=1 sequence=8
archived log file
name=/nfsanf/archlog/NTAP1/archivelog/2024_03_22/o1_mf_1_9__9yfrj0b1
_.arc thread=1 sequence=9
archived log file
name=/nfsanf/archlog/NTAP1/archivelog/2024_03_22/o1_mf_1_10__9yomybb
c_.arc thread=1 sequence=10
media recovery complete, elapsed time: 00:01:17
Finished recover at 22-MAR-24

RMAN> exit

```

Recovery Manager complete.



For faster recovery, enable parallel sessions with `recovery_parallelism` parameter or specify degree of parallel in recovery command for database recovery: `RECOVER DATABASE PARALLEL (DEGREE d INSTANCES DEFAULT) ;`. In general, degrees of parallelism should be equal to number of CPU cores on the host.

18. Exit RMAN, login to Oracle as oracle user via sqlplus to open database and reset log after an incomplete recovery.

```

SQL> select name, open_mode from v$database;

NAME          OPEN_MODE
-----
NTAP1         MOUNTED

SQL> select instance_name, host_name from v$instance;

INSTANCE_NAME
-----
HOST_NAME
-----
NTAP1
ora-02

SQL>

```

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

```
MEMBER
```

```
-----
```

```
/u03/orareco/NTAP1/onlineelog/redo03.log  
/u03/orareco/NTAP1/onlineelog/redo02.log  
/u03/orareco/NTAP1/onlineelog/redo01.log
```

```
SQL> alter database rename file  
'/u03/orareco/NTAP1/onlineelog/redo01.log' to  
'/nfsanf/oracopy/redo01.log';
```

```
Database altered.
```

```
SQL> alter database rename file  
'/u03/orareco/NTAP1/onlineelog/redo02.log' to  
'/nfsanf/oracopy/redo02.log';
```

```
Database altered.
```

```
SQL> alter database rename file  
'/u03/orareco/NTAP1/onlineelog/redo03.log' to  
'/nfsanf/oracopy/redo03.log';
```

```
Database altered.
```

```
SQL> alter database open resetlogs;
```

```
Database altered.
```

```
SQL> show pdbs
```

CON_ID	CON_NAME	OPEN	MODE	RESTRICTED
2	PDB\$SEED	READ	ONLY	NO
3	NTAP1_PDB1	READ	WRITE	NO
4	NTAP1_PDB2	READ	WRITE	NO
5	NTAP1_PDB3	READ	WRITE	NO

19. Validate the database structure restored to new host as well as the test row we have inserted before primary VLDB failure.

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

NAME

/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSTEM_FNO-1_1t2m9nij
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSAUX_FNO-3_1u2m9nog
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-UNDOTBS1_FNO-4_1v2m9nu6
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSTEM_FNO-5_282m9oem
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSAUX_FNO-6_242m9oan
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-USERS_FNO-7_2c2m9ofn
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-UNDOTBS1_FNO-8_292m9oem
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSTEM_FNO-9_252m9oc5
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSAUX_FNO-10_212m9o52
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-UNDOTBS1_FNO-
11_202m9o22
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-USERS_FNO-12_2d2m9ofs

NAME

/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSTEM_FNO-13_262m9oca
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSAUX_FNO-14_222m9o53
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-UNDOTBS1_FNO-
15_2a2m9of6
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-USERS_FNO-16_2e2m9og8
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSTEM_FNO-17_272m9oel
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SYSAUX_FNO-18_232m9oa8
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-UNDOTBS1_FNO-
19_2b2m9ofn
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-USERS_FNO-20_2f2m9og8
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-21_1h2m9cap
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-22_1i2m9cap
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-23_1j2m9cap

NAME

/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-24_1k2m9cap
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-25_1l2m9g3u
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-26_1m2m9g9j
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-27_1n2m9gcg
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-28_1o2m9gd4
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-29_1p2m9ju6
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-30_1q2m9k7a
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-31_1r2m9kfk
/nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-32_1s2m9kgg

31 rows selected.

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

MEMBER

```
-----  
-----  
/nfsanf/oracopy/redo03.log  
/nfsanf/oracopy/redo02.log  
/nfsanf/oracopy/redo01.log
```

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

NAME

```
-----  
-----  
/nfsanf/oracopy/NTAP1.ctl
```

```
SQL> alter session set container=ntap1_pdb1;
```

Session altered.

```
SQL> select * from test;
```

```
          ID  
-----  
DT  
-----  
-----  
EVENT  
-----  
-----  
          1  
21-MAR-24 03.15.03.000000 PM  
test oracle incremental merge switch to copy  
  
          2  
22-MAR-24 02.22.06.000000 PM  
test recovery on a new Azure VM host with image copy on ANF
```

20. Drop invalid tempfiles and add new tempfiles to temp tablespaces.

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

NAME

```
-----  
/u02/oradata/NTAP1/NTAP1_pdb1/temp01.dbf  
/u02/oradata/NTAP1/NTAP1_pdb1/temp02.dbf
```

```
SQL> alter tablespace temp add tempfile  
'/nfsanf/oracopy/ntap1_pdb1_temp01.dbf' size 100M;
```

Tablespace altered.

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

```
NAME  
-----  
-----
```

```
/u02/oradata/NTAP1/NTAP1_pdb1/temp01.dbf  
/u02/oradata/NTAP1/NTAP1_pdb1/temp02.dbf  
/nfsanf/oracopy/ntap1_pdb1_temp01.dbf
```

```
SQL> alter database tempfile  
'/u02/oradata/NTAP1/NTAP1_pdb1/temp01.dbf' offline;
```

Database altered.

```
SQL> alter database tempfile  
'/u02/oradata/NTAP1/NTAP1_pdb1/temp01.dbf' drop;
```

Database altered.

```
SQL> alter database tempfile  
'/u02/oradata/NTAP1/NTAP1_pdb1/temp02.dbf' offline;
```

Database altered.

```
SQL> alter database tempfile  
'/u02/oradata/NTAP1/NTAP1_pdb1/temp02.dbf' drop;
```

Database altered.

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

```
NAME  
-----  
-----
```

```
/nfsanf/oracopy/ntap1_pdb1_temp01.dbf
```

```
SQL>
```

21. Other post recovery tasks

- Add ANF NFS mount to fstab so that the NFS file system will be mounted when DB server host rebooted.

As azureuser, sudo vi /etc/fstab and add following entry:

```
172.30.136.68:/ora-01-u02-copy          /nfsanf          nfs
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsiz=262144,wsiz=262144,noi
tr 0          0
```

- Update the Oracle init file from primary database init file backup that is restored to /tmp/archive and create spfile as needed.

This completes the Oracle VLDB database recovery from backup image copy on ANF NFS file system to a standby DB server host.

Where to find additional information

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

- RMAN: Merged Incremental Backup Strategies (Doc ID 745798.1)

https://support.oracle.com/knowledge/Oracle%20Database%20Products/745798_1.html

- RMAN Backup and Recovery User's Guide

<https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/getting-started-rman.html>

- Azure NetApp Files

<https://azure.microsoft.com/en-us/products/netapp>

TR-4987: Simplified, Automated Oracle Deployment on Azure NetApp Files with NFS

Allen Cao, Niyaz Mohamed, NetApp

This solution provides overview and details for automated Oracle deployment in Microsoft Azure NetApp Files as primary database storage with NFS protocol and Oracle database is deployed as container database with dNFS enabled.

Purpose

Running performance-intensive and latency-sensitive Oracle workloads in the cloud can be challenging. Azure NetApp Files (ANF) makes it easy for enterprise line-of-business (LOB) and storage professionals to migrate

and run demanding Oracle workloads without code change. Azure NetApp Files is widely used as the underlying shared file-storage service in various scenarios, such as new deployment or migration (lift and shift) of Oracle databases from on-premises to Azure.

This documentation demonstrates the simplified deployment of Oracle databases in Azure NetApp files via NFS mounts using Ansible automation. The Oracle database deploys in a container database (CDB) and pluggable databases (PDB) configuration with Oracle dNFS protocol enabled to boost performance. Furthermore, the on-premises Oracle single instance database or PDB can be migrated into a newly deployed container database in Azure using automated PDB relocation methodology with minimal service interruption. It also provides information on fast Oracle database backup, restore, and clone with NetApp SnapCenter UI tool in Azure Cloud.

This solution addresses the following use cases:

- Automated Oracle container database deployment on Azure NetApp files
- Automated Oracle database migration between on-premises and Azure cloud

Audience


This solution is intended for the following people:

- A DBA who would like to deploy Oracle on Azure NetApp Files.
- A database solution architect who would like to test Oracle workloads on Azure NetApp Files.
- A storage administrator who would like to deploy and manage an Oracle database on Azure NetApp Files.
- An application owner who would like to stand up an Oracle database on Azure NetApp Files.

Solution test and validation environment

The testing and validation of this solution were performed in a lab setting that might not match the final deployment environment. See the section [Key factors for deployment consideration](#) for more information.

Architecture



Hardware and software components

Hardware		
Azure NetApp Files	Current offering in Azure by Microsoft	One capacity pool with Premium service level
Azure VM for DB server	Standard_B4ms - 4 vCPUs, 16GiB	Two Linux virtual machine instances for concurrent deployment
Azure VM for SnapCenter	Standard_B4ms - 4 vCPUs, 16GiB	One Windows virtual machine instance
Software		
RedHat Linux	RHEL Linux 8.6 (LVM) - x64 Gen2	Deployed RedHat subscription for testing
Windows Server	2022 DataCenter; Azure Edition Hotpatch - x64 Gen2	Hosting SnapCenter server
Oracle Database	Version 19.18	Applied RU patch p34765931_190000_Linux-x86-64.zip
Oracle OPatch	Version 12.2.0.1.36	Latest patch p6880880_190000_Linux-x86-64.zip
SnapCenter Server	Version 5.0	Workgroup deployment
Open JDK	Version java-11-openjdk	SnapCenter plugin requirement on DB VMs

NFS	Version 3.0	Oracle dNFS enabled
Ansible	core 2.16.2	Python 3.6.8

Oracle database configuration in the lab environment

Server	Database	DB Storage
ora-01	NTAP1(NTAP1_PDB1,NTAP1_PD B2,NTAP1_PDB3)	/u01, /u02, /u03 NFS mounts on ANF capacity pool
ora-02	NTAP2(NTAP2_PDB1,NTAP2_PD B2,NTAP2_PDB3)	/u01, /u02, /u03 NFS mounts on ANF capacity pool

Key factors for deployment consideration

- **Oracle database storage layout.** In this automated Oracle deployment, we provision three database volumes for each database to host Oracle binary, data, and logs by default. The volumes are mounted on Oracle DB server as /u01 - binary, /u02 - data, /u03 - logs via NFS. Dual control files are configured on /u02 and /u03 mount points for redundancy.
- **Multiple DB servers deployment.** The automation solution can deploy an Oracle container database to multiple DB servers in a single Ansible playbook run. Regardless of the number of DB servers, the playbook execution remains the same. You can deploy multiple container databases to a single VM instance by repeating the deployment with different database instance IDs (Oracle SID). But ensure there is sufficient memory on the host to support deployed databases.
- **dNFS configuration.** By using dNFS (available since Oracle 11g), an Oracle database running on an Azure Virtual Machine can drive significantly more I/O than the native NFS client. Automated Oracle deployment configures dNFS on NFSv3 by default.
- **Allocate large size volume to speed up deployment.** ANF file system IO throughput is regulated based on the size of volume. For initial deployment, allocate large size volumes can speed up the deployment. The volumes subsequently can be downsized dynamically without application impact.
- **Database backup.** NetApp provides a SnapCenter software suite for database backup, restore, and cloning with a user-friendly UI interface. NetApp recommends implementing such a management tool to achieve fast (under a minute) snapshot backup, quick (minutes) database restore, and database clone.

Solution deployment

The following sections provide step-by-step procedures for automated Oracle 19c deployment and database migration on Azure NetApp Files with directly mounted database volumes via NFS to Azure VMs.

Prerequisites for deployment

Deployment requires the following prerequisites.

1. An Azure account has been set up, and the necessary VNet and network segments have been created within your Azure account.
2. From the Azure cloud portal, deploy Azure Linux VMs as Oracle DB servers. Create an Azure NetApp Files capacity pool and database volumes for Oracle database. Enable VM SSH private/public key authentication for azureuser to DB servers. See the architecture diagram in the previous section for details about the environment setup. Also referred to [Step-by-Step Oracle deployment procedures on Azure VM and Azure NetApp Files](#) for detailed information.



For Azure VMs deployed with local disk redundancy, ensure that you have allocated at least 128G in the VM root disk to have sufficient space to stage Oracle installation files and add OS swap file. Expand /tmplv and /rootlv OS partition accordingly. Ensure the database volume naming follows the VMname-u01, VMname-u02, and VMname-u03 convention.

```
sudo lvresize -r -L +20G /dev/mapper/rootvg-rootlv
```

```
sudo lvresize -r -L +10G /dev/mapper/rootvg-tmplv
```

3. From the Azure cloud portal, provision a Windows server to run the NetApp SnapCenter UI tool with the latest version. Refer to the following link for details: [Install the SnapCenter Server](#)
4. Provision a Linux VM as the Ansible controller node with the latest version of Ansible and Git installed. Refer to the following link for details: [Getting Started with NetApp solution automation](#) in section -
Setup the Ansible Control Node for CLI deployments on RHEL / CentOS or
Setup the Ansible Control Node for CLI deployments on Ubuntu / Debian.



The Ansible controller node can locate either on-premises or in Azure cloud as far as it can reach Azure DB VMs via ssh port.

5. Clone a copy of the NetApp Oracle deployment automation toolkit for NFS.

```
git clone https://bitbucket.ngage.netapp.com/scm/ns-bb/na_oracle_deploy_nfs.git
```

6. Stage following Oracle 19c installation files on Azure DB VM /tmp/archive directory with 777 permission.

```
installer_archives:
  - "LINUX.X64_193000_db_home.zip"
  - "p34765931_190000_Linux-x86-64.zip"
  - "p6880880_190000_Linux-x86-64.zip"
```

7. Watch the following video:

[Simplified and automated Oracle deployment on Azure NetApp Files with NFS](#)

Automation parameter files

Ansible playbook executes database installation and configuration tasks with predefined parameters. For this Oracle automation solution, there are three user-defined parameter files that need user input before playbook execution.

- `hosts` - define targets that the automation playbook is running against.
- `vars/vars.yml` - the global variable file that defines variables that apply to all targets.
- `host_vars/host_name.yml` - the local variable file that defines variables that apply only to a named target. In our use case, these are the Oracle DB servers.

In addition to these user-defined variable files, there are several default variable files that contain default parameters that do not require change unless necessary. The following sections show how to configure the user-defined variable files.

Parameter files configuration

1. Ansible target hosts file configuration:

```
# Enter Oracle servers names to be deployed one by one, follow by  
each Oracle server public IP address, and ssh private key of admin  
user for the server.
```

```
[oracle]
```

```
ora_01 ansible_host=10.61.180.21 ansible_ssh_private_key_file  
=ora_01.pem
```

```
ora_02 ansible_host=10.61.180.23 ansible_ssh_private_key_file  
=ora_02.pem
```

2. Global vars/vars.yml file configuration

```
#####
##
##### Oracle 19c deployment user configuration variables
#####
##### Consolidate all variables from ONTAP, linux and oracle
#####
#####
#####

#####
### ONTAP env specific config variables ###
#####

# Prerequisite to create three volumes in NetApp ONTAP storage from
System Manager or cloud dashboard with following naming convention:
# db_hostname_u01 - Oracle binary
# db_hostname_u02 - Oracle data
# db_hostname_u03 - Oracle redo
# It is important to strictly follow the name convention or the
automation will fail.

#####
### Linux env specific config variables ###
#####

redhat_sub_username: XXXXXXXXX
redhat_sub_password: XXXXXXXXX

#####
### DB env specific install and config variables ###
#####

# Database domain name
db_domain: solutions.netapp.com

# Set initial password for all required Oracle passwords. Change
them after installation.
initial_pwd_all: XXXXXXXXX
```

3. Local DB server host_vars/host_name.yml configuration such as ora_01.yml, ora_02.yml ...

```
# User configurable Oracle host specific parameters

# Enter container database SID. By default, a container DB is
created with 3 PDBs within the CDB
oracle_sid: NTAP1

# Enter database shared memory size or SGA. CDB is created with SGA
at 75% of memory_limit, MB. The grand total of SGA should not exceed
75% available RAM on node.
memory_limit: 8192

# Local NFS lif ip address to access database volumes
nfs_lif: 172.30.136.68
```

Playbook execution

There are a total of five playbooks in the automation toolkit. Each performs different task blocks and serves different purposes.

```
0-all_playbook.yml - execute playbooks from 1-4 in one playbook run.
1-ansible_requirements.yml - set up Ansible controller with required
libs and collections.
2-linux_config.yml - execute Linux kernel configuration on Oracle DB
servers.
4-oracle_config.yml - install and configure Oracle on DB servers and
create a container database.
5-destroy.yml - optional to undo the environment to dismantle all.
```

There are three options to run the playbooks with the following commands.

1. Execute all deployment playbooks in one combined run.

```
ansible-playbook -i hosts 0-all_playbook.yml -u azureuser -e
@vars/vars.yml
```

2. Execute playbooks one at a time with the number sequence from 1-4.

```
ansible-playbook -i hosts 1-ansible_requirements.yml -u azureuser -e
@vars/vars.yml
```

```
ansible-playbook -i hosts 2-linux_config.yml -u azureuser -e
@vars/vars.yml
```

```
ansible-playbook -i hosts 4-oracle_config.yml -u azureuser -e
@vars/vars.yml
```

3. Execute 0-all_playbook.yml with a tag.

```
ansible-playbook -i hosts 0-all_playbook.yml -u azureuser -e
@vars/vars.yml -t ansible_requirements
```

```
ansible-playbook -i hosts 0-all_playbook.yml -u azureuser -e
@vars/vars.yml -t linux_config
```



```
ansible-playbook -i hosts 0-all_playbook.yml -u azureuser -e  
@vars/vars.yml -t oracle_config
```

4. Undo the environment

```
ansible-playbook -i hosts 5-destroy.yml -u azureuser -e  
@vars/vars.yml
```

Post execution validation

After the playbook run, login to the Oracle DB server VM to validate that Oracle is installed and configured and a container database is created successfully. Following is an example of Oracle database validation on host ora-01.

1. Validate NFS mounts

```
[azureuser@ora-01 ~]$ cat /etc/fstab

#
# /etc/fstab
# Created by anaconda on Thu Sep 14 11:04:01 2023
#
# Accessible filesystems, by reference, are maintained under
# '/dev/disk/'.
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for
# more info.
#
# After editing this file, run 'systemctl daemon-reload' to update
# systemd
# units generated from this file.
#
/dev/mapper/rootvg-rootlv / xfs defaults
0 0
UUID=268633bd-f9bb-446d-9a1d-8fca4609a1e1 /boot
xfs defaults 0 0
UUID=89D8-B037 /boot/efi vfat
defaults,uid=0,gid=0,umask=077,shortname=winnt 0 2
/dev/mapper/rootvg-homelv /home xfs defaults
0 0
/dev/mapper/rootvg-tmplv /tmp xfs defaults
0 0
/dev/mapper/rootvg-usrlv /usr xfs defaults
0 0
/dev/mapper/rootvg-varlv /var xfs defaults
0 0
/mnt/swapfile swap swap defaults 0 0
172.30.136.68:/ora-01-u01 /u01 nfs
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsiz=65536,wsiz=65536 0 0
172.30.136.68:/ora-01-u02 /u02 nfs
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsiz=65536,wsiz=65536 0 0
172.30.136.68:/ora-01-u03 /u03 nfs
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsiz=65536,wsiz=65536 0 0

[azureuser@ora-01 ~]$ df -h
Filesystem                Size      Used Avail Use% Mounted on
devtmpfs                  7.7G         0  7.7G   0% /dev
```

tmpfs	7.8G	0	7.8G	0%	/dev/shm
tmpfs	7.8G	8.6M	7.7G	1%	/run
tmpfs	7.8G	0	7.8G	0%	/sys/fs/cgroup
/dev/mapper/rootvg-rootlv	22G	17G	5.8G	74%	/
/dev/mapper/rootvg-usrlv	10G	2.0G	8.1G	20%	/usr
/dev/mapper/rootvg-varlv	8.0G	890M	7.2G	11%	/var
/dev/sda1	496M	106M	390M	22%	/boot
/dev/mapper/rootvg-homelv	1014M	40M	975M	4%	/home
/dev/sda15	495M	5.9M	489M	2%	/boot/efi
/dev/mapper/rootvg-tmplv	12G	8.4G	3.7G	70%	/tmp
tmpfs	1.6G	0	1.6G	0%	/run/user/54321
172.30.136.68:/ora-01-u01	500G	11G	490G	3%	/u01
172.30.136.68:/ora-01-u03	250G	1.2G	249G	1%	/u03
172.30.136.68:/ora-01-u02	250G	7.1G	243G	3%	/u02
tmpfs	1.6G	0	1.6G	0%	/run/user/1000

2. Validate Oracle listener

```
[azureuser@ora-01 ~]$ sudo su
[root@ora-01 azureuser]# su - oracle
Last login: Thu Feb  1 16:13:44 UTC 2024
[oracle@ora-01 ~]$ lsnrctl status listener.ntap1

LSNRCTL for Linux: Version 19.0.0.0.0 - Production on 01-FEB-2024
16:25:37

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

Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=TCP) (HOST=ora-
01.internal.cloudapp.net) (PORT=1521)))
STATUS of the LISTENER
-----
Alias                     LISTENER.NTAP1
Version                   TNSLSNR for Linux: Version 19.0.0.0.0 -
Production
Start Date                01-FEB-2024 16:13:49
Uptime                    0 days 0 hr. 11 min. 49 sec
Trace Level               off
Security                  ON: Local OS Authentication
SNMP                      OFF
Listener Parameter File   /u01/app/oracle/product/19.0.0/NTAP1/network/admin/listener.ora
Listener Log File         /u01/app/oracle/diag/tnslsnr/ora-
01/listener.ntap1/alert/log.xml
Listening Endpoints Summary...
```

```

      (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp) (HOST=ora-
01.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net) (PORT=1521)))
      (DESCRIPTION=(ADDRESS=(PROTOCOL=ipc) (KEY=EXTPROC1521)))
      (DESCRIPTION=(ADDRESS=(PROTOCOL=tcps) (HOST=ora-
01.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net) (PORT=5500)) (
Security=(my_wallet_directory=/u01/app/oracle/product/19.0.0/NTAP1/a
dmin/NTAP1/xdb_wallet)) (Presentation=HTTP) (Session=RAW))
Services Summary...
Service "104409ac02da6352e063bb891eacf34a.solutions.netapp.com" has
1 instance(s).
    Instance "NTAP1", status READY, has 1 handler(s) for this
service...
Service "104412c14c2c63cae063bb891eacf64d.solutions.netapp.com" has
1 instance(s).
    Instance "NTAP1", status READY, has 1 handler(s) for this
service...
Service "1044174670ad63ffe063bb891eac6b34.solutions.netapp.com" has
1 instance(s).
    Instance "NTAP1", status READY, has 1 handler(s) for this
service...
Service "NTAP1.solutions.netapp.com" has 1 instance(s).
    Instance "NTAP1", status READY, has 1 handler(s) for this
service...
Service "NTAP1XDB.solutions.netapp.com" has 1 instance(s).
    Instance "NTAP1", status READY, has 1 handler(s) for this
service...
Service "ntap1_pdb1.solutions.netapp.com" has 1 instance(s).
    Instance "NTAP1", status READY, has 1 handler(s) for this
service...
Service "ntap1_pdb2.solutions.netapp.com" has 1 instance(s).
    Instance "NTAP1", status READY, has 1 handler(s) for this
service...
Service "ntap1_pdb3.solutions.netapp.com" has 1 instance(s).
    Instance "NTAP1", status READY, has 1 handler(s) for this
service...
The command completed successfully

```

3. Validate Oracle database and dNFS

```

[oracle@ora-01 ~]$ 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.
#
#
NTAP1:/u01/app/oracle/product/19.0.0/NTAP1:Y
```

```
[oracle@ora-01 ~]$ sqlplus / as sysdba
```

```
SQL*Plus: Release 19.0.0.0.0 - Production on Thu Feb 1 16:37:51 2024
Version 19.18.0.0.0
```

```
Copyright (c) 1982, 2022, Oracle.  All rights reserved.
```

```
Connected to:
```

```
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 -
Production
Version 19.18.0.0.0
```

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

NAME	OPEN_MODE	LOG_MODE
NTAP1	READ WRITE	ARCHIVELOG

```
SQL> show pdbs
```

CON_ID	CON_NAME	OPEN MODE	RESTRICTED
2	PDB\$SEED	READ ONLY	NO
3	NTAP1_PDB1	READ WRITE	NO
4	NTAP1_PDB2	READ WRITE	NO

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

```
NAME
```

```
-----
-----
/u02/oradata/NTAP1/system01.dbf
/u02/oradata/NTAP1/sysaux01.dbf
/u02/oradata/NTAP1/undotbs01.dbf
/u02/oradata/NTAP1/pdbseed/system01.dbf
/u02/oradata/NTAP1/pdbseed/sysaux01.dbf
/u02/oradata/NTAP1/users01.dbf
/u02/oradata/NTAP1/pdbseed/undotbs01.dbf
/u02/oradata/NTAP1/NTAP1_pdb1/system01.dbf
/u02/oradata/NTAP1/NTAP1_pdb1/sysaux01.dbf
/u02/oradata/NTAP1/NTAP1_pdb1/undotbs01.dbf
/u02/oradata/NTAP1/NTAP1_pdb1/users01.dbf
```

```
NAME
```

```
-----
-----
/u02/oradata/NTAP1/NTAP1_pdb2/system01.dbf
/u02/oradata/NTAP1/NTAP1_pdb2/sysaux01.dbf
/u02/oradata/NTAP1/NTAP1_pdb2/undotbs01.dbf
/u02/oradata/NTAP1/NTAP1_pdb2/users01.dbf
/u02/oradata/NTAP1/NTAP1_pdb3/system01.dbf
/u02/oradata/NTAP1/NTAP1_pdb3/sysaux01.dbf
/u02/oradata/NTAP1/NTAP1_pdb3/undotbs01.dbf
/u02/oradata/NTAP1/NTAP1_pdb3/users01.dbf
```

```
19 rows selected.
```

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

```
NAME
```

```
-----
-----
/u02/oradata/NTAP1/control01.ctl
/u03/orareco/NTAP1/control02.ctl
```

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

```
MEMBER
```

```
-----
-----
/u03/orareco/NTAP1/onlineelog/redo03.log
```

```
/u03/orareco/NTAP1/onlineelog/redo02.log
```

```
/u03/orareco/NTAP1/onlineelog/redo01.log
```

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

```
SVRNAME
```

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

```
DIRNAME
```

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

```
NFSVERSION
```

```
-----
```

```
172.30.136.68
```

```
/ora-01-u02
```

```
NFSv3.0
```

```
172.30.136.68
```

```
/ora-01-u03
```

```
NFSv3.0
```

```
SVRNAME
```

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

```
DIRNAME
```

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

```
NFSVERSION
```

```
-----
```

```
172.30.136.68
```

```
/ora-01-u01
```

```
NFSv3.0
```

4. Login to Oracle Enterprise Manager Express to validate database.

ORACLE ENTERPRISE MANAGER DATABASE EXPRESS

Username:

Password:

Container Name:

Log In

ORACLE

Copyright 2013, 2020, Oracle and/or its affiliates. All rights reserved.

ORACLE Enterprise Manager Database Express

NTAP1 (19.18.0.0.0) Performance Storage

Database Home

Time Zone:

1 min Auto-Refresh Refresh

Status

Up Time: 34 minutes, 43 seconds
 Type: Single Instance (NTAP1)
 CDB (3 PDB(s))
 Version: 19.18.0.0.0 Enterprise Edition
 Platform Name: Linux x86 64-bit
 Thread: 1
 Archiver: Started
 Last Backup Time: N/A
 Incident(s): 1

Performance

Activity Services Containers

4.0
3.0
2.0
1.0
0.0

03:49:00 PM 03:54:20 PM 03:59:40 PM 04:05:00 PM 04:10:20 PM 04:15:40 PM 04:21:00 PM 04:26:20 PM 04:31:40 PM 04:37:00 PM 04:42:20 PM 04:47:40 PM

Feb 1, 2024 GMT-00:00

Other Network System I/O User I/O CPU

Resources

1% 1% 1% 0% 0% 0% 0% 0% 0% 0%

Host CPU

60 50 40 30 20 10 0

Active Sessions

Wait User I/O CPU

14 GB 11.2 GB 8.4 GB 5.6 GB 2.8 GB 0.8

Memory

total_sga total_lpga shared_pool large_pool buffer cache Shared IO P...

3.7 GB 2.8 GB 1.9 GB 953.7 MB 0.8

Data Storage

NTAP1_PDB3 NTAP1_PDB2 NTAP1_PDB1

SQL Monitor - Last Hour (20 max)

Top 20 by Last Active Time

Filter by Status, SQL ID or User Name

Status	Duration	SQL ID	SQL Plan Hash	User Name	Parallel	Database Time	I/O Requests	SQL Text
--------	----------	--------	---------------	-----------	----------	---------------	--------------	----------

Migrate Oracle database to Azure

Oracle database migration from on-premises to the cloud is a heavy-lifting. Using the right strategy and automation can smooth the process and minimize service interruption and downtime. Follow this detailed instruction [Database migration from on-premises to Azure cloud](#) to guide your database migration journey.

Oracle backup, restore, and clone with SnapCenter

NetApp recommends SnapCenter UI tool to manage Oracle database deployed in Azure cloud. Please refer to TR-4988: [Oracle Database Backup, Recovery, and Clone on ANF with SnapCenter](#) for details.

Where to find additional information

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

- Oracle Database Backup, Recovery, and Clone on ANF with SnapCenter

[Oracle Database Backup, Recovery, and Clone on ANF with SnapCenter](#)

- Azure NetApp Files

<https://azure.microsoft.com/en-us/products/netapp>

- Deploying Oracle Direct NFS

<https://docs.oracle.com/en/database/oracle/oracle-database/19/ldb/ deploying-dnfs.html#GUID-D06079DB-8C71-4F68-A1E3-A75D7D96DCE2>

- Installing and Configuring Oracle Database Using Response Files

<https://docs.oracle.com/en/database/oracle/oracle-database/19/ldb/ installing-and-configuring-oracle-database-using-response-files.html#GUID-D53355E9-E901-4224-9A2A-B882070EDDF7>

Oracle Database Deployment and Protection on Azure NetApp Files

TR-4954: Oracle Database Deployment and Protection on Azure NetApp Files

This best practice guide provides details of a solution for deploying and protecting Oracle database on Azure NetApp file storage and Azure VM.

Author(s): Allen Cao, Niyaz Mohamed, NetApp

Overview

Many mission-critical Oracle enterprise databases are still hosted on-premises, and many enterprises are looking to migrate these Oracle databases to a public cloud. Often, these Oracle databases are application centric and thus require user-specific configurations, a capability that is missing from many database-as-a-service public-cloud offerings. Therefore, the current database landscape calls for a public-cloud-based Oracle database solution built from a high-performance, scalable compute and storage service that can accommodate unique requirements. Azure virtual machine compute instances and the Azure NetApp Files storage service might be the missing pieces of this puzzle that you can leverage to build and migrate your mission-critical Oracle database workloads to a public cloud.

Azure Virtual Machine

Azure virtual machines are one of several types of on-demand, scalable computing resources that Azure offers. Typically, you choose a virtual machine when you need more control over the computing environment than the other choices offer. Azure virtual machines offer a quick and easy way to create a computer with specific configurations required to run your Oracle database, whether it is for compute- or memory-intensive workloads. Virtual machines in an Azure virtual network can easily be connected to your organization's network, for example through a secured VPN tunnel.

Azure NetApp Files (ANF)

Azure NetApp Files is a fully managed Microsoft service that will take your database workload to the cloud faster and more securely than ever before. It was designed to meet the core requirements of running high-performance workloads such as Oracle databases in the cloud, and it provides performance tiers that reflect the real-world range of IOPS demands, low latency, high availability, high durability, manageability at scale, and fast and efficient backup, recovery, and cloning. These capabilities are possible because Azure NetApp Files is based on physical all-flash NetApp ONTAP systems running within the Azure data center environment. Azure NetApp Files is completely integrated into the Azure DCs and portal, and customers can use the same comfortable graphical interface and APIs for creating and managing shared files as with any other Azure object. With Azure NetApp file, you can unlock the full capabilities of Azure without extra risk, cost, or time and trust the only enterprise file service native to Azure.

Conclusion

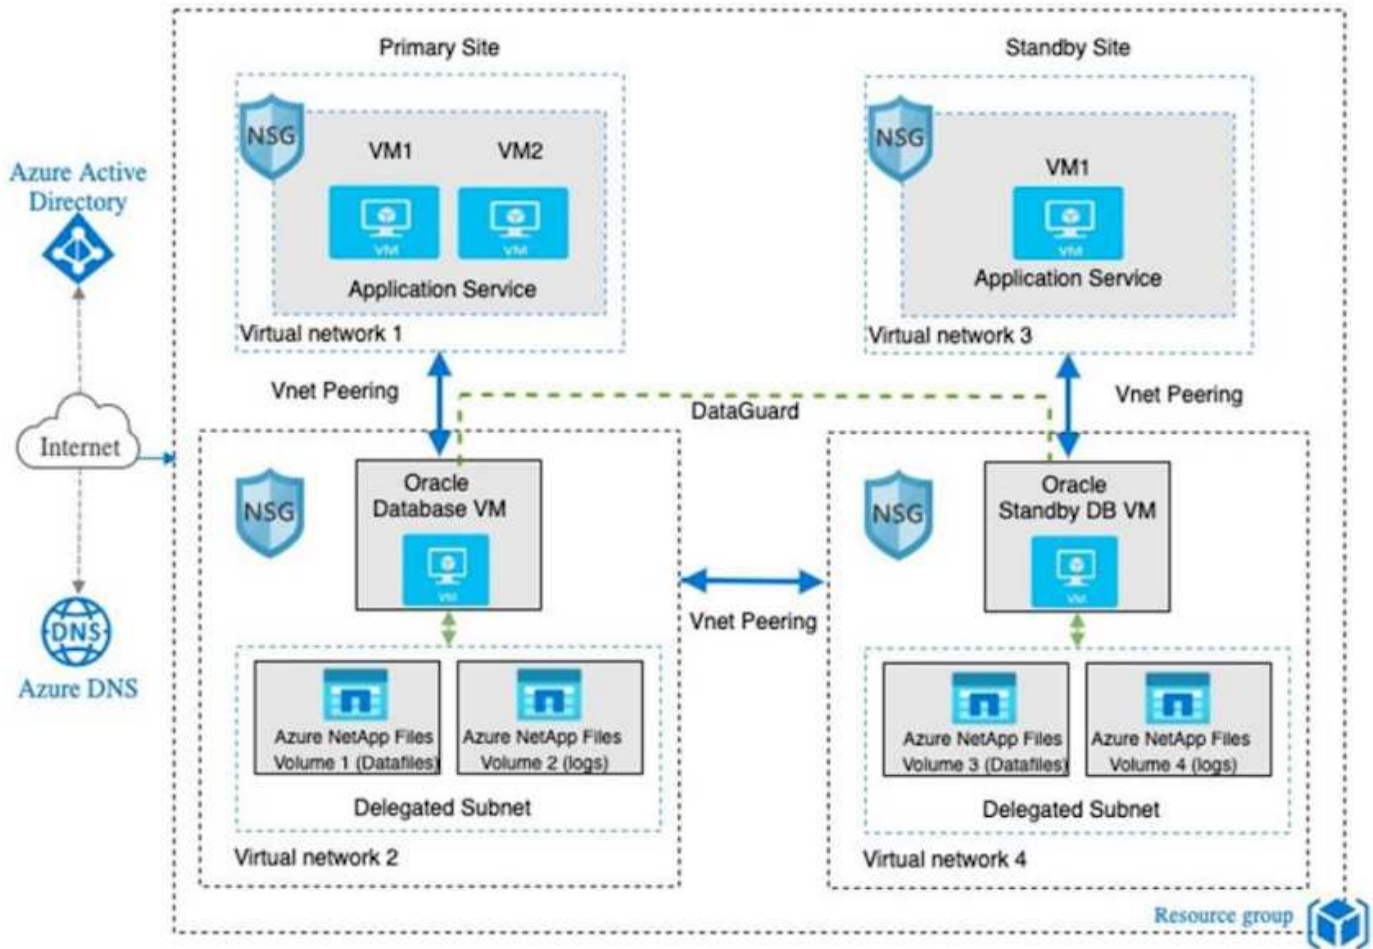
This documentation describes in detail how to deploy, configure, and protect an Oracle database with an Azure virtual machine and Azure NetApp Files storage service that delivers performance and durability similar to an on-premises system. For best-practices guidance, see TR-4780 [Oracle Databases on Microsoft Azure](#). More importantly, NetApp also provides automation toolkits that automate most of the tasks that are required for the deployment, configuration, data protection, migration, and management of your Oracle database workload in the Azure public cloud. The automation toolkits are available for download at NetApp public GitHub site: [NetApp-Automation](#).

Solution Architecture

The following architecture diagram illustrates a highly available Oracle database deployment on Azure VM instances and the Azure NetApp Files storage.


Within the environment, the Oracle compute instance is deployed via an Azure services VM console. There are multiple Azure instance types available from the console. NetApp recommends deploying a database-oriented Azure VM instance that meets your expected workload.

Oracle database storage on the other hand is deployed with the Azure NetApp Files service available from Azure console. The Oracle binary, data, or log volumes are subsequently presented and mounted on an Azure VM instance Linux host.



In many respects, the implementation of Azure NetApp Files in Azure cloud is very similar to an on-premises ONTAP data storage architecture with many built-in redundancies, such as RAID and dual controllers. For disaster recovery, a standby site can be setup in different regions and database can be synced up with the primary site using application-level replication (for example, Oracle Data Guard).

In our test validation for Oracle database deployment and data protection, the Oracle database is deployed on a single Azure VM as illustrated in the following diagram:



The Azure Oracle environment can be managed with an Ansible controller node for automation using tool kits provided by NetApp for database deployment, backup, recovery, and database migration. Any updates to the Oracle Azure VM instance operating-system kernel or Oracle patching can be performed in parallel to keep the primary and standby in sync. In fact, the initial toolkits can be easily expanded to perform daily Oracle tasks if needed. If you need help to set up a CLI Ansible controller, see [NetApp Solution Automation](#) to get started.

Factors to consider for Oracle database deployment

A public cloud provides many choices for compute and storage, and using the correct type of compute instance and storage engine is a good place to start for database deployment. You should also select compute and storage configurations that are optimized for Oracle databases.

The following sections describe the key considerations when deploying an Oracle database in the Azure public cloud on an Azure virtual machine instance with Azure NetApp Files storage.

VM type and sizing

Selecting the right VM type and size is important for optimal performance of a relational database in a public cloud. An Azure virtual machine provides a variety of compute instances that can be used to host Oracle database workloads. See the Microsoft documentation [Sizes for virtual machines in Azure](#) for different types of Azure virtual machines and their sizing. In general, NetApp recommends using a general-purpose Azure virtual machine for the deployment of small- and medium-sized Oracle databases. For the deployment of larger Oracle databases, a memory-optimized Azure VM is appropriate. With more available RAM, a larger Oracle SGA or smart flash cache can be configured to reduce the physical I/O, which in turn improves database performance.

Azure NetApp Files works as an NFS mount attached to an Azure virtual machine, which offers higher throughput and overcomes the storage-optimized VM throughput limit with local storage. Therefore, running Oracle on Azure NetApp Files could reduce the licensable Oracle CPU core count and licensing costs. See [TR-4780: Oracle Databases on Microsoft Azure](#), Section 7 - How Does Oracle Licensing Work?

Other factors to consider include the following:

- Choose the correct vCPU and RAM combination based on workload characteristics. As the RAM size increases on the VM, so does the number of vCPU cores. There should be a balance at some point as the Oracle license fees are charged on the number of vCPU cores.
- Add swap space to a VM. The default Azure VM deployment does not create a swap space, which is not optimal for a database.

Azure NetApp Files performance

Azure NetApp Files volumes are allocated from a capacity pool the customer must provision in their Azure NetApp Files storage account. Each capacity pool is assigned as follows:

- To a service level that defines the overall performance capability.
- The initially provisioned storage capacity or tiering for that capacity pool. A quality of service (QoS) level that defines the overall maximum throughput per provisioned space.

The service level and initially provisioned storage capacity determines the performance level for a particular Oracle database volume.

1. Service Levels for Azure NetApp Files

Azure NetApp Files supports three service levels: Ultra, Premium, and Standard.

- **Ultra storage.** This tier provides up to 128MiBps of throughput per 1TiB of volume quota assigned.
- **Premium storage.** This tier provides up to 64MiBps of throughput per 1TiB of volume quota assigned.

- **Standard storage.** This tier provides up to 16MiBps of throughput per 1TiB of volume quota assigned.

2. Capacity pool and quality of service

Each of the desired service levels has an associated cost for provisioned capacity and includes a quality-of-service (QoS) level that defines the overall maximum throughput for provisioned space.


For example, a 10TiB-provisioned single-capacity pool with the premium service level provides an overall available throughput for all volumes in this capacity pool of 10x 64MBps, so 640MBps with 40,000 (16K) IOPs or 80,000 (8K) IOPs.

The minimum capacity pool size is 4TiB. You can change the size of a capacity pool in 1TiB increments in response to changes in your workload requirements to manage storage needs and costs.

3. Calculate the service level at a database volume

The throughput limit for an Oracle database volume is determined by a combination of the following factors: The service level of the capacity pool to which the volume belongs and The quota assigned to the volume.

The following diagram shows how the throughput limit for an Oracle database volume is calculated.



In example 1, a volume from a capacity pool with the Premium storage tier that is assigned 2TiB of quota is assigned a throughput limit of 128MiBps (2TiB * 64MiBps). This scenario applies regardless of the capacity pool size or the actual volume consumption.

In example 2, a volume from a capacity pool with the Premium storage tier that is assigned 100GiB of quota is assigned a throughput limit of 6.25MiBps (0.09765625TiB * 64MiBps). This scenario applies regardless of the capacity pool size or the actual volume consumption.

Please note that the minimum volume size is 100GiB.

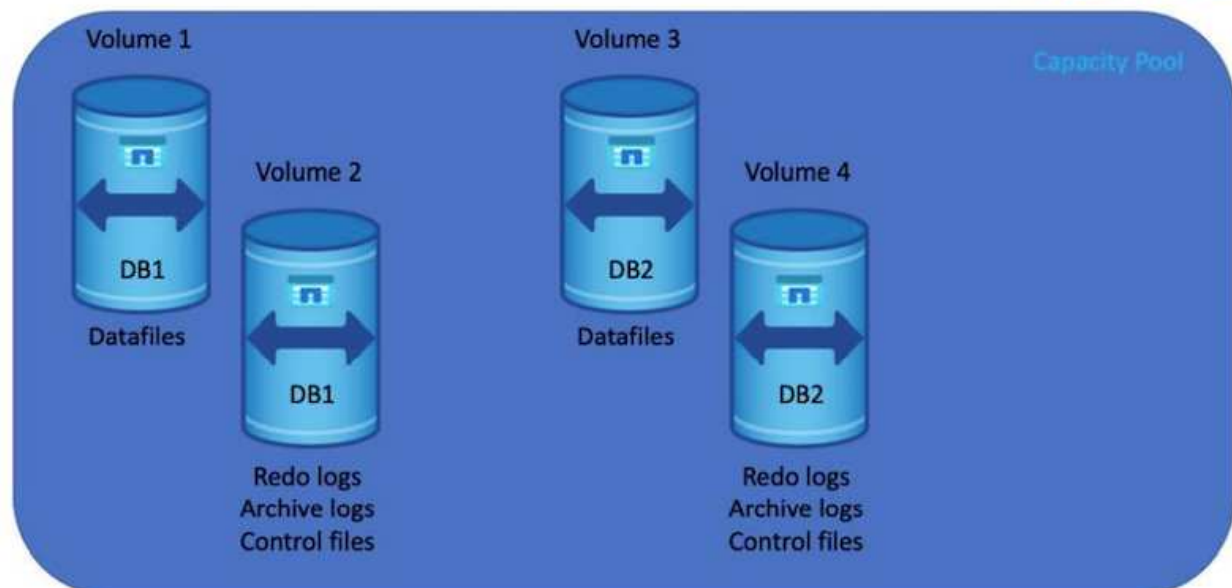
Storage layout and settings

NetApp recommends the following storage layout:

- For small databases, using single volume layout for all Oracle files.



- For large databases, the recommended volume layout is multiple volumes: one for Oracle data and a duplicate control file and one for the Oracle active log, archived log, and control file. NetApp highly recommends allocating a volume for the Oracle binary instead of the local drive so that the database can be relocated to a new host and quickly restored.




NFS configuration

Linux, the most common operating system, includes native NFS capabilities. Oracle offers a direct NFS (dNFS) client natively integrated into Oracle. Oracle dNFS bypasses the OS cache and enables parallel processing to

improve database performance. Oracle has supported NFSv3 for over 20 years, and NFSv4 is supported with Oracle 12.1.0.2 and later.

By using dNFS (available since Oracle 11g), an Oracle database running on an Azure Virtual Machine can drive significantly more I/O than the native NFS client. Automated Oracle deployment using the NetApp automation toolkit automatically configures dNFS on NFSv3.

The following diagram demonstrates the SLOB benchmark on Azure NetApp Files with Oracle dNFS.



Other factors to consider:

- TCP slot tables are the NFS equivalent of host-bus-adapter (HBA) queue depth. These tables control the number of NFS operations that can be outstanding at any one time. The default value is usually 16, which is far too low for optimum performance. The opposite problem occurs on newer Linux kernels, which can automatically increase the TCP slot table limit to a level that saturates the NFS server with requests.

For optimum performance and to prevent performance problems, adjust the kernel parameters that control TCP slot tables to 128.

```
sysctl -a | grep tcp.*.slot_table
```

- The following table provides recommended NFS mount options for a single instance of Linux NFSv3.

File Type	Mount Options
<ul style="list-style-type: none"> Control files Data files Redo logs 	<code>rw,bg,hard,vers=3,proto=tcp,timeo=600,rsiz=65536,wsiz=65536</code>
<ul style="list-style-type: none"> ORACLE_HOME ORACLE_BASE 	<code>rw,bg,hard,vers=3,proto=tcp,timeo=600,rsiz=65536,wsiz=65536</code>



Before using dNFS, verify that the patches described in Oracle Doc 1495104.1 are installed. The NetApp Support matrix for NFSv3 and NFSv4 do not include specific operating systems. All OSs that obey the RFC are supported. When searching the online IMT for NFSv3 or NFSv4 support, do not select a specific OS because no matches will be displayed. All OSs are implicitly supported by the general policy.

Step-by-Step Oracle deployment procedures on Azure VM and Azure NetApp Files

This section describes the deployment procedures of deploying Oracle RDS custom database with FSx storage.

Deploy an Azure VM with ANF for Oracle via Azure portal console

If you are new to Azure, you first need to set up an Azure account environment. This includes signing up your organization to use Azure Active Directory. The following section is a summary of these steps. For details, see the linked Azure-specific documentation.

Create and consume Azure resources

After your Azure environment is set up and an account is created and associated with a subscription, you can log into Azure portal with the account to create the necessary resources to run Oracle.

1. Create a virtual network or VNet

Azure Virtual Network (VNet) is the fundamental building block for your private network in Azure. VNet enables many types of Azure resources, such as Azure Virtual Machines (VMs), to securely communicate with each other, the internet, and on-premises networks. Before provisioning an Azure VM, a VNet (where a VM is deployed) must first be configured.

See [Create a virtual network using the Azure portal](#) to create a VNet.

2. Create a NetApp storage account and capacity pool for ANF

In this deployment scenario, an Azure VM OS is provisioned using regular Azure storage, but ANF volumes are provisioned to run Oracle database via NFS. First, you need to create a NetApp storage account and a capacity pool to host the storage volumes.

See [Set up Azure NetApp Files and create an NFS volume](#) to set up an ANF capacity pool.

3. Provision Azure VM for Oracle

Based on your workload, determine what type of Azure VM you need and the size of the VM vCPU and RAM to deploy for Oracle. Then, from the Azure console, click the VM icon to launch the VM deployment workflow.

1. From the Azure VM page, click **Create** and then choose **Azure virtual machine**.

Microsoft Azure									
Virtual machines									
Hybrid Cloud TME									
Create Switch to classic Reservations Manage view Refresh Export to CSV Open query Assign tags Start Restart Stop Delete Services Maintenance									
Filter for any field... Subscription equals all Type equals all Resource group equals all Location equals all Add filter									
No grouping List view									
Name	Type	Subscription	Resource group	Location	Status	Operating system	Size	Public IP address	Disks
acao-ora01	Virtual machine	Hybrid Cloud TME Onprem	TMEtstres	South Central US	Stopped (deallocated)	Linux	Standard_B4ms	13.65.63.157	1
ANFAV5val2JH	Virtual machine	Hybrid Cloud TME Onprem	ANFAV5VAL2	West Europe	Running	Windows	Standard_DS2_v2	20.229.80.88	1
ANFAV5fio01	Virtual machine	Hybrid Cloud TME Onprem	anfav5rg	South Central US	Stopped (deallocated)	Linux	Standard_D32ds_v4	-	1
ANFAV5fioA21	Virtual machine	Hybrid Cloud TME Onprem	anfav5rg	South Central US	Running	Linux	Standard_E32as_v4	40.124.74.246	1
ANFAV5fioA22	Virtual machine	Hybrid Cloud TME Onprem	anfav5rg	South Central US	Stopped (deallocated)	Linux	Standard_E32as_v4	40.124.178.111	1
ANFAV5fioA23	Virtual machine	Hybrid Cloud TME Onprem	anfav5rg	South Central US	Stopped (deallocated)	Linux	Standard_E32as_v4	40.124.194.32	1
ANFAV5valDC	Virtual machine	Hybrid Cloud TME Onprem	anfav5rg	South Central US	Stopped (deallocated)	Windows	Standard_B4ms	-	1
ANFAV5valIH	Virtual machine	Hybrid Cloud TME Onprem	anfav5rg	South Central US	Running	Windows	Standard_B2ms	70.37.66.218	1
ANFAV5valIH2	Virtual machine	Hybrid Cloud TME Onprem	anfav5rg	South Central US	Running	Windows	Standard_B2s	20.225.210.195	1
ANFCVOCM	Virtual machine	Hybrid Cloud TME Onprem	anfav5val2	West Europe	Running	Linux	Standard_DS3_v2	-	1
ANFCVODRD2	Virtual machine	Hybrid Cloud TME Onprem	anfav5val2	West Europe	Running	Windows	Standard_B2s	-	1
ANFCVODRDemo	Virtual machine	Hybrid Cloud TME Onprem	anfcvodrdemo-rg	West Europe	Running	Linux	Standard_E4s_v3	-	5
AVSCVOPerfinguest	Virtual machine	Hybrid Cloud TME Onprem	avscvoperfinguest-rg	West Europe	Stopped (deallocated)	Linux	Standard_DS15_v2	-	5

2. Choose the subscription ID for the deployment, and then choose the resource group, region, host name, VM image, size, and authentication method. Go to the Disk page.

[Home](#) > [Virtual machines](#) >

Create a virtual machine ...

[Basics](#) [Disks](#) [Networking](#) [Management](#) [Advanced](#) [Tags](#) [Review + create](#)

Create a virtual machine that runs Linux or Windows. Select an image from Azure marketplace or use your own customized image. Complete the Basics tab then Review + create to provision a virtual machine with default parameters or review each tab for full customization. [Learn more](#)

Project details

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription *	<div>Hybrid Cloud TME Onprem</div> <div>▼</div>
Resource group *	<div>ANFAVSRG</div> <div>▼</div>

[Create new](#)

Instance details

Virtual machine name *	<div>acao-ora01</div> <div>✓</div>
Region *	<div>(US) South Central US</div> <div>▼</div>
Availability options	<div>No infrastructure redundancy required</div> <div>▼</div>
Security type	<div>Standard</div> <div>▼</div>
Image *	<div> Red Hat Enterprise Linux 8.0 (LVM) - Gen2</div> <div>▼</div>
See all images Configure VM generation	
Run with Azure Spot discount	<div><input type="checkbox"/></div>
Size *	<div>Standard_D8s_v3 - 8 vcpus, 32 GiB memory (\$273.02/month)</div> <div>▼</div>
See all sizes	

Administrator account

Authentication type	<div><input type="radio"/> SSH public key</div> <div><input checked="" type="radio"/> Password</div>
---------------------	--

[Review + create](#)[< Previous](#)[Next : Disks >](#)

Create a virtual machine ...

Size * ⓘ

Standard_D8s_v3 - 8 vcpus, 32 GiB memory (\$273.02/month) ▼

[See all sizes](#)

Administrator account

Authentication type ⓘ

☐ SSH public key

☒ Password

Username * ⓘ

azureuser ✓

Password * ⓘ

..... ✓

Confirm password * ⓘ

..... ✓

Inbound port rules

Select which virtual machine network ports are accessible from the public internet. You can specify more limited or granular network access on the Networking tab.


Public inbound ports * ⓘ

☐ None

☒ Allow selected ports

Select inbound ports *

SSH (22) ▼

 **This will allow all IP addresses to access your virtual machine.** This is only recommended for testing. Use the Advanced controls in the Networking tab to create rules to limit inbound traffic to known IP addresses.

Licensing

If you have eligible Red Hat Enterprise Linux subscriptions that are enabled for Red Hat Cloud Access, you can use Azure Hybrid Benefit to attach your Red Hat subscriptions to this VM and save money on compute costs [Learn more](#)

Your Azure subscription is currently not a part of Red Hat Cloud Access. In order to enable AHB for this VM, you must add this Azure subscription to Cloud Access. [Learn more](#)

Review + create

< Previous

Next : Disks >

- Choose **premium SSD** for OS local redundancy and leave the data disk blank because the data disks are mounted from ANF storage. Go to the Networking page.

Create a virtual machine ...

[Basics](#)
[Disks](#)
[Networking](#)
[Management](#)
[Advanced](#)
[Tags](#)
[Review + create](#)

Azure VMs have one operating system disk and a temporary disk for short-term storage. You can attach additional data disks. The size of the VM determines the type of storage you can use and the number of data disks allowed. [Learn more](#)

Disk options

OS disk type *

Premium SSD (locally-redundant storage)

Delete with VM

☒

Enable encryption at host

☐

Encryption at host is not registered for the selected subscription. [Learn more about enabling this feature](#)

Encryption type *

(Default) Encryption at-rest with a platform-managed key

Enable Ultra Disk compatibility

☐

Data disks for acao-ora01

You can add and configure additional data disks for your virtual machine or attach existing disks. This VM also comes with a temporary disk.

LUN	Name	Size (GiB)	Disk type	Host caching	Delete with VM
<div> Create and attach a new disk Attach an existing disk </div>					

Advanced

Review + create

< Previous

Next : Networking >

- Choose the VNet and subnet. Allocate a public IP for external VM access. Then go to the Management page.


Home > Virtual machines >

Create a virtual machine ...

Network interface

When creating a virtual machine, a network interface will be created for you.

Virtual network * ⓘ	<div>ANFAVSVal</div> <div>Create new</div>
Subnet * ⓘ	<div>VM_Sub (172.30.137.128/25)</div> <div>Manage subnet configuration</div>
Public IP ⓘ	<div>(new) acao-ora01-ip</div> <div>Create new</div>
NIC network security group ⓘ	<div><input type="radio"/> None</div> <div><input checked="" type="radio"/> Basic</div> <div><input type="radio"/> Advanced</div>
Public inbound ports * ⓘ	<div><input type="radio"/> None</div> <div><input checked="" type="radio"/> Allow selected ports</div>
Select inbound ports *	<div>SSH (22)</div>

 **This will allow all IP addresses to access your virtual machine.** This is only recommended for testing. Use the Advanced controls in the Networking tab to create rules to limit inbound traffic to known IP addresses.

Delete public IP and NIC when VM is deleted ⓘ	<input checked="" type="checkbox"/>
Enable accelerated networking ⓘ	<input checked="" type="checkbox"/>

Load balancing

You can place this virtual machine in the backend pool of an existing Azure load balancing solution. [Learn more](#)

Place this virtual machine behind an existing load balancing solution?	<input type="checkbox"/>
--	--------------------------

Review + create

< Previous

Next : Management >

5. Keep all defaults for Management and move to the Advanced page.

Microsoft Azure

Search resources, services, and docs (G+/I)

Home > Virtual machines >

Create a virtual machine ...

BasicsDisksNetworkingManagementAdvancedTagsReview + create

Configure monitoring and management options for your VM.

Microsoft Defender for Cloud

Microsoft Defender for Cloud provides unified security management and advanced threat protection across hybrid cloud workloads. [Learn more](#)

✔ Your subscription is protected by Microsoft Defender for Cloud basic plan.

Monitoring

Boot diagnostics ⓘ
☒ Enable with managed storage account (recommended)
☐ Enable with custom storage account
☐ Disable

Enable OS guest diagnostics ⓘ ☐

Identity

Enable system assigned managed identity ⓘ ☐

Azure AD

Login with Azure AD ⓘ ☐

ℹ RBAC role assignment of Virtual Machine Administrator Login or Virtual Machine User Login is required when using Azure AD login. [Learn more](#)

ℹ Azure AD login now uses SSH certificate-based authentication. You will need to use an SSH client that supports OpenSSH certificates. You can use Azure CLI or Cloud Shell from the Azure Portal. [Learn more](#)

Auto-shutdown

Enable auto-shutdown ⓘ ☐

Backup

Review + create

< Previous

Next : Advanced >

6. Keep all defaults for the Advanced page unless you need to customize a VM after deployment with custom scripts. Then go to Tags page.

Home > Virtual machines >

Create a virtual machine ...

Basics Disks Networking Management Advanced Tags Review + create

Add additional configuration, agents, scripts or applications via virtual machine extensions or cloud-init.

Extensions

Extensions provide post-deployment configuration and automation.

Extensions ⓘ [Select an extension to install](#)

VM applications

VM applications contain application files that are securely and reliably downloaded on your VM after deployment. In addition to the application files, an install and uninstall script are included in the application. You can easily add or remove applications on your VM after create. [Learn more](#) ⓘ

[Select a VM application to install](#)

Custom data

Pass a script, configuration file, or other data into the virtual machine **while it is being provisioned**. The data will be saved on the VM in a known location. [Learn more about custom data for VMs](#) ⓘ

Custom data

i Your image must have a code to support consumption of custom data. If your image supports cloud-init, custom-data will be processed by cloud-init. [Learn more about custom data for VMs](#) ⓘ

User data

Pass a script, configuration file, or other data that will be accessible to your applications **throughout the lifetime of the virtual machine**. Don't use user data for storing your secrets or passwords. [Learn more about user data for VMs](#) ⓘ

Enable user data

☐

[Review + create](#)

[< Previous](#)

[Next : Tags >](#)

7. Add a tag for the VM if desired. Then, go to the Review + create page.


[Home](#) > [Virtual machines](#) >

Create a virtual machine ...

Basics Disks Networking Management Advanced Tags Review + create

Tags are name/value pairs that enable you to categorize resources and view consolidated billing by applying the same tag to multiple resources and resource groups. [Learn more about tags](#)

Note that if you create tags and then change resource settings on other tabs, your tags will be automatically updated.

Name ⓘ	Value ⓘ	Resource
<input type="text" value="database"/>	<input type="text" value="oracle"/>	12 selected  
<input type="text"/>	<input type="text"/>	12 selected 

Review + create

< Previous

Next : Review + create >

- The deployment workflow runs a validation on the configuration, and, if the validation passes, click **Create** to create the VM.

4. Provision ANF database volumes for Oracle

You must create three NFS volumes for an ANF capacity pool for the Oracle binary, data, and log volumes respectively.

1. From the Azure console, under the list of Azure services, click Azure NetApp Files to open a volume creation workflow. If you have more than one ANF storage account, click the account that you would like to provision volumes from.

The screenshot shows the Microsoft Azure portal interface. At the top, there's a search bar and a navigation menu. Below the navigation menu, the 'Azure services' section displays various service icons, including 'Create a resource', 'Azure NetApp Files' (highlighted), 'Virtual networks', 'Virtual machines', 'Storage accounts', 'Users', 'Subscriptions', 'Azure Active Directory', 'Quickstart Center', and 'More services'. Below this, the 'Resources' section shows a table of recent resources.

Name	Type	Last Viewed
ANFAVSAcct	NetApp account	a few seconds ago
ANFAVVal	Virtual network	3 hours ago
acao-ora01	Virtual machine	5 days ago
Hybrid Cloud TME Onprem	Subscription	2 weeks ago
WEANFAVSAcct	NetApp account	2 weeks ago
ANFAVSAcct/CapPool/acao-ora01-u03	Volume	2 weeks ago
ANFAVSAcct/CapPool/acao-ora01-u02	Volume	2 weeks ago
ANFAVSAcct/CapPool/acao-ora01-u01	Volume	2 weeks ago
acao-ora01_OsDisk_1_673bad70ccce4709afc81278e2bc97cb	Disk	2 weeks ago
acao-ora0166	Network Interface	3 weeks ago
TMEtstres	Resource group	3 weeks ago

2. Under your NetApp storage account, click **Volumes**, and then **Add volume** to create new Oracle volumes.

The screenshot shows the Microsoft Azure portal interface for the 'ANFAVSAcct' NetApp account. The left sidebar shows the 'Azure NetApp Files' section with 'Volumes' highlighted. The main content area shows the 'Essentials' section with details about the resource group, location, subscription, and tags. Below this, there's a section titled 'Enterprise files storage, powered by NetApp' with three cards: 'Connect to Active Directory', 'Capacity pools', and 'Volumes'. The 'Volumes' card is highlighted, and it shows a 'View volumes' button.

Enterprise files storage, powered by NetApp

Azure NetApp Files makes it easy for enterprise line-of-business (LOB) and storage professionals to migrate and run complex, file-based applications with no code change. [Learn more](#)

Connect to Active Directory
Connect your NetApp to Active Directory. [Learn more](#)

Capacity pools
Purchase pools of capacity with a service level in which you provision volumes. [Learn more](#)

Volumes
Container for active file system, associated meta-data, and snapshots. [Learn more](#)

Microsoft Azure

Home > Azure NetApp Files > ANFAVSAcct

Azure NetApp Files

Hybrid Cloud TIME

+ Create Manage view

Filter for any field...

Name ↑

ANFAVSAcct

WEANFAVSAcct

ANFAVSAcct | Volumes

NetApp account

Search (Ctrl+/)

+ Add volume + Add data replication Refresh

Search volumes

Name	Quota	Throughput	Protocol type	Mount path	Service level	Network features	Capacity pool
anf2-z1-stds01	200 GiB	25.6 MiB/s	NFSv3	172.30.136.70/anf2-z1-std	Ultra	Standard	capool
anf2-z1-stds02	200 GiB	25.6 MiB/s	NFSv3	172.30.136.70/anf2-z1-std	Ultra	Standard	capool
anf2-z1-stds03	100 GiB	12.8 MiB/s	NFSv3	172.30.136.70/anf2-z1-std	Ultra	Standard	capool
anf2-z1-stds04	100 GiB	12.8 MiB/s	NFSv3	172.30.136.70/anf2-z1-std	Ultra	Standard	capool
anf2-z1-stds05	100 GiB	12.8 MiB/s	NFSv3	172.30.136.70/anf2-z1-std	Ultra	Standard	capool
anf2-z1-stds06	100 GiB	12.8 MiB/s	NFSv3	172.30.136.70/anf2-z1-std	Ultra	Standard	capool
anf2-z1-stds07	100 GiB	12.8 MiB/s	NFSv3	172.30.136.70/anf2-z1-std	Ultra	Standard	capool
anf2-z1-stds08	100 GiB	12.8 MiB/s	NFSv3	172.30.136.70/anf2-z1-std	Ultra	Standard	capool
anf-z1-stds01	6 TiB	786.432 MiB/s	NFSv3	172.30.136.70/anf-z1-std	Ultra	Standard	capool
anf-z1-stds02	200 GiB	25.6 MiB/s	NFSv3	172.30.136.70/anf-z1-std	Ultra	Standard	capool

- As a good practice, identify Oracle volumes with the VM hostname as a prefix and then followed by the mount point on the host, such as u01 for Oracle binary, u02 for Oracle data, and u03 for Oracle log. Choose the same VNet for the volume as for the VM. Click **Next: Protocol**.

Microsoft Azure

Home > Azure NetApp Files > ANFAVSAcct | Volumes

ANFAVSAcct | Volumes

NetApp account

Search (Ctrl+/)

+ Add volume

Search volumes

Overview

Activity log

Access control (IAM)

Tags

Settings

Quota

Properties

Locks

Azure NetApp Files

Active Directory connections

Storage service

Capacity pools

Volumes

Data protection

Snapshot policies

Storage service add-ons

NetApp add-ons

Automation

Tasks (preview)

Export template

Support + troubleshooting

New Support Request

Create a volume

Basics Protocol Tags Review + create

This page will help you create an Azure NetApp Files volume in your subscription and enable you to access the volume from within your virtual network. [Learn more about Azure NetApp Files](#)

Volume details

Volume name * acao-ora01_u01

Capacity pool * CapPool

Available quota (GiB) 572

Quota (GiB) * 100

Available throughput (MiB/s) 73.22

Throughput (MiB/s) 12.5

Enable Cool Access

Coolness Period 31

Virtual network * ANFAVSAcct (172.30.136.64/26,172.30.137.128/25,172.30.152.0/27)

Delegated subnet * ANF_Sub (172.30.136.64/26)

Network features Basic Standard

Availability Zone None

Show advanced section

Review + create

< Previous

Next: Protocol >

- Choose the NFS protocol, add the Oracle host IP address to the allowed client, and remove the default policy that allows all IP addresses 0.0.0.0/0. Then click **Next: Tags**.

Microsoft Azure

Search resources, services, and docs (G+)

Home > Azure NetApp Files > ANFAVSAcct | Volumes >

ANFAVSAcct | Volumes

NetApp account

Search (Ctrl+J)

+ Add volume

Search volumes

Overview

Activity log

Access control (IAM)

Tags

Settings

Quota

Properties

Locks

Azure NetApp Files

Active Directory connections

Storage service

Capacity pools

Volumes

Data protection

Snapshot policies

Storage service add-ons

NetApp add-ons

Automation

Tasks (preview)

Export template

Support + troubleshooting

New Support Request

Name

Quota

anf2-z1-stdds01

200 GiB

anf2-z1-stdds02

200 GiB

anf2-z1-stdds03

100 GiB

anf2-z1-stdds04

100 GiB

anf2-z1-stdds05

100 GiB

anf2-z1-stdds06

100 GiB

anf2-z1-stdds07

100 GiB

anf2-z1-stdds08

100 GiB

anf-z1-stdds01

6 TiB

anf-z1-stdds02

200 GiB

anf-z1-stdds03

1 TiB

anf-z1-stdds04

200 GiB

anf-z1-stdds06

200 GiB

anf-z1-stdds07

200 GiB

anf-z1-stdds08

200 GiB

anf-zq-stdds05

1 TiB

vol1

1 TiB

vol3basic

100 GiB

volnfsbasic

100 GiB

volnfsstd

100 GiB

volnfsstdnew

100 GiB

zone1basic

6 TiB

zone2basic

100 GiB

Create a volume

Basics

Protocol

Tags

Review + create

Configure access to your volume.

Access

Protocol type

NFS

SMB

Dual-protocol

Configuration

File path *

aca0-ora01_u01

Versions *

NFSv3

Kerberos

Enabled

Disabled

LDAP

Enabled

Disabled

Azure VMware Solution DataStore

Export policy

Configure the volume's export policy. This can be edited later. [Learn more](#)

Move up

Move down

Move to top

Move to bottom

Delete

Index

Allowed clients

Access

Root Access

1

172.30.137.142

Read & Write

On

2

172.30.137.142

Read & Write

On

Review + create

< Previous

Next : Tags >

5. Add a volume tag if desired. Then click **Review + Create**.

Microsoft Azure

Search resources, services, and docs (G+)

Home > Azure NetApp Files > ANFAVSAcct | Volumes >

ANFAVSAcct | Volumes

NetApp account

Search (Ctrl+/)

+ Add volume

...

Overview

Activity log

Access control (IAM)

Tags

Settings

Quota

Properties

Locks

Azure NetApp Files

Active Directory connections

Storage service

Capacity pools

Volumes

Data protection

Snapshot policies

Storage service add-ons

NetApp add-ons

Automation

Tasks (preview)

Export template

Support + troubleshooting

New Support Request

Search volumes

Name	Quota
anf2-z1-stds01	200 GiB
anf2-z1-stds02	200 GiB
anf2-z1-stds03	100 GiB
anf2-z1-stds04	100 GiB
anf2-z1-stds05	100 GiB
anf2-z1-stds06	100 GiB
anf2-z1-stds07	100 GiB
anf2-z1-stds08	100 GiB
anf-z1-stds01	6 TiB
anf-z1-stds02	200 GiB
anf-z1-stds03	1 TiB
anf-z1-stds04	200 GiB
anf-z1-stds06	200 GiB
anf-z1-stds07	200 GiB
anf-z1-stds08	200 GiB
anf-zq-stds05	1 TiB
vol1	1 TiB
vol3basic	100 GiB
volnfsbasic	100 GiB
volnfsstd	100 GiB
volnfsstdnew	100 GiB
zone1basic	6 TiB
zone2basic	100 GiB

Create a volume

Basics

Protocol

Tags

Review + create

Tags are name/value pairs that enable you to categorize resources and view consolidated billing by applying the same tag to multiple resources and resource groups. [Learn more about tags](#)

Note that if you create tags and then change resource settings on other tabs, your tags will be automatically updated.

Name

Value

database : oracle

Review + create

< Previous

Next : Review + create >

6. If the validation passes, click **Create** to create the volume.

Microsoft Azure

Search resources, services, and docs (G+)

Home > Azure NetApp Files > ANFAVSAcct | Volumes >

ANFAVSAcct | Volumes

NetApp account

Search (Ctrl+J)

Overview

Activity log

Access control (IAM)

Tags

Settings

Quota

Properties

Locks

Azure NetApp Files

Active Directory connections

Storage service

Capacity pools

Volumes

Data protection

Snapshot policies

Storage service add-ons

NetApp add-ons

Automation

Tasks (preview)

Export template

Support + troubleshooting

New Support Request

Search volumes

Name

Quota

anf2-z1-stdds01

200 GiB

anf2-z1-stdds02

200 GiB

anf2-z1-stdds03

100 GiB

anf2-z1-stdds04

100 GiB

anf2-z1-stdds05

100 GiB

anf2-z1-stdds06

100 GiB

anf2-z1-stdds07

100 GiB

anf2-z1-stdds08

100 GiB

anf-z1-stdds01

6 TiB

anf-z1-stdds02

200 GiB

anf-z1-stdds03

1 TiB

anf-z1-stdds04

200 GiB

anf-z1-stdds06

200 GiB

anf-z1-stdds07

200 GiB

anf-z1-stdds08

200 GiB

anf-zq-stdds05

1 TiB

vol1

1 TiB

vol3basic

100 GiB

volInfsbasic

100 GiB

volInfsstd

100 GiB

volInfsstdnew

100 GiB

zone1basic

6 TiB

zone2basic

100 GiB

Create a volume

Validation passed

Basics

Protocol

Tags

Review + create

Basics

Subscription

Resource group

Region

Volume name

Capacity pool

Service level

Quota

Encryption key source

Availability Zone

Networking

Virtual network

Delegated subnet

Network features

Protocol

Protocol

File path

Tags

database

Hybrid Cloud TME Onprem

ANFAVSRG

South Central US

acao-ora01-u01

CapPool

Ultra

100 GiB

Microsoft.NetApp

None

ANFAVSVol (172.30.136.64/26,172.30.137.128/25,172.30.152.0/27)

ANF_Sub (172.30.136.64/26)

Standard

NFSv3

acao-ora01-u01

oracle

Create

< Previous

Next >

Download a template for automation

Install and configure Oracle on Azure VM with ANF

The NetApp solutions team has created many Ansible-based automation toolkits to help you deploy Oracle in Azure smoothly. Follow these steps to deploy Oracle on an Azure VM.

Set up an Ansible controller

If you have not set up an Ansible controller, see [NetApp Solution Automation](#), which has detailed instructions on how to setup an Ansible controller.

Obtain Oracle deployment automation toolkit

Clone a copy of the Oracle deployment toolkit in your home directory under the user ID that you use to log into the Ansible controller.

```
git clone https://github.com/NetApp-Automation/na_oracle19c_deploy.git
```

Execute the toolkit with your configuration

See the [CLI deployment Oracle 19c Database](#) to execute the playbook with the CLI. You can ignore the ONTAP portion of the variables configuration in the global VARS file when you create database volumes from

the Azure console rather than the CLI.



The toolkit default deploys Oracle 19c with RU 19.8. It can be easily adapted for any other patch level with minor default configuration changes. Also default seed-database active log files are deployed into the data volume. If you need active log files on the log volume, it should be relocated after initial deployment. Reach out to the NetApp Solution team for help if needed.

Set up AzAcSnap backup tool for app-consistent snapshots for Oracle

The Azure Application-Consistent Snapshot tool (AzAcSnap) is a command-line tool that enables data protection for third-party databases by handling all the orchestration required to put them into an application-consistent state before taking a storage snapshot. It then returns these databases to an operational state. NetApp recommends installing the tool on the database server host. See the following installation and configuration procedures.

Install AzAcSnap tool

1. Get the most recent version of the [the AzArcSnap Installer](#).
2. Copy the downloaded self-installer to the target system.
3. Execute the self-installer as the root user with the default installation option. If necessary, make the file executable using the `chmod +x *.run` command.

```
./azacsnap_installer_v5.0.run -I
```

Configure Oracle connectivity

The snapshot tools communicate with the Oracle database and need a database user with appropriate permissions to enable or disable backup mode.

1. Set up AzAcSnap database user

The following examples show the setup of the Oracle database user and the use of sqlplus for communication to the Oracle database. The example commands set up a user (AZACSNAP) in the Oracle database and change the IP address, usernames, and passwords as appropriate.

1. From the Oracle database installation, launch sqlplus to log into the database.

```
su - oracle  
sqlplus / AS SYSDBA
```

2. Create the user.

```
CREATE USER azacsnap IDENTIFIED BY password;
```

3. Grant the user permissions. This example sets the permission for the AZACSNAP user to enable putting the database into backup mode.

```
GRANT CREATE SESSION TO azacsnap;  
GRANT SYSBACKUP TO azacsnap;
```

4. Change the default user's password expiration to unlimited.

```
ALTER PROFILE default LIMIT PASSWORD_LIFE_TIME unlimited;
```

5. Validate azacsnap connectivity for the database.

```
connect azacsnap/password  
quit;
```

2. Configure Linux-user azacsnap for DB access with Oracle wallet

The AzAcSnap default installation creates an azacsnap OS user. It's Bash shell environment must be configured for Oracle database access with the password stored in an Oracle wallet.

1. As root user, run the `cat /etc/oratab` command to identify the ORACLE_HOME and ORACLE_SID variables on the host.

```
cat /etc/oratab
```

2. Add ORACLE_HOME, ORACLE_SID, TNS_ADMIN, and PATH variables to the azacsnap user bash profile. Change the variables as needed.

```
echo "export ORACLE_SID=ORATEST" >> /home/azacsnap/.bash_profile  
echo "export ORACLE_HOME=/u01/app/oracle/product/19800/ORATST" >>  
/home/azacsnap/.bash_profile  
echo "export TNS_ADMIN=/home/azacsnap" >> /home/azacsnap/.bash_profile  
echo "export PATH=\$PATH:\$ORACLE_HOME/bin" >>  
/home/azacsnap/.bash_profile
```

3. As the Linux user azacsnap, create the wallet. You are prompted for the wallet password.

```
sudo su - azacsnap  
  
mkstore -wrl $TNS_ADMIN/.oracle_wallet/ -create
```

4. Add the connect string credentials to the Oracle Wallet. In the following example command, AZACSNAP is the ConnectString to be used by AzAcSnap, azacsnap is the Oracle Database User, and AzPasswd1 is the Oracle User's database password. You are again prompted for the wallet password.


```
mkstore -wrl $TNS_ADMIN/.oracle_wallet/ -createCredential AZACSNAP  
azacsnap AzPasswd1
```

5. Create the `tnsnames-ora` file. In the following example command, `HOST` should be set to the IP address of the Oracle Database and the `Server SID` should be set to the Oracle Database SID.

```
echo "# Connection string  
AZACSNAP=\"(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP) (HOST=172.30.137.142) (PORT=1521)) (CONNECT_DATA=(SID=ORATST)))\"  
" > $TNS_ADMIN/tnsnames.ora
```

6. Create the `sqlnet.ora` file.

```
echo "SQLNET.WALLET_OVERRIDE = TRUE  
WALLET_LOCATION=(  
    SOURCE=(METHOD=FILE)  
    (METHOD_DATA=(DIRECTORY=$TNS_ADMIN/.oracle_wallet))  
) " > $TNS_ADMIN/sqlnet.ora
```

7. Test Oracle access using the wallet.

```
sqlplus /@AZACSNAP as SYSBACKUP
```

The expected output from the command:

```
[azacsnap@acao-ora01 ~]$ sqlplus /@AZACSNAP as SYSBACKUP  
  
SQL*Plus: Release 19.0.0.0.0 - Production on Thu Sep 8 18:02:07 2022  
Version 19.8.0.0.0  
  
Copyright (c) 1982, 2019, Oracle. All rights reserved.  
  
Connected to:  
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production  
Version 19.8.0.0.0  
  
SQL>
```

Configure ANF connectivity

This section explains how to enable communication with Azure NetApp Files (with a VM).

1. Within an Azure Cloud Shell session, make sure that you are logged into the subscription that you want to be associated with the service principal by default.

```
az account show
```

2. If the subscription isn't correct, use the following command:

```
az account set -s <subscription name or id>
```

3. Create a service principal using the Azure CLI as in the following example:

```
az ad sp create-for-rbac --name "AzAcSnap" --role Contributor --scopes /subscriptions/{subscription-id} --sdk-auth
```

The expected output:

```
{
  "clientId": "00aa000a-aaaa-0000-00a0-00aa000aaa0a",
  "clientSecret": "00aa000a-aaaa-0000-00a0-00aa000aaa0a",
  "subscriptionId": "00aa000a-aaaa-0000-00a0-00aa000aaa0a",
  "tenantId": "00aa000a-aaaa-0000-00a0-00aa000aaa0a",
  "activeDirectoryEndpointUrl": "https://login.microsoftonline.com",
  "resourceManagerEndpointUrl": "https://management.azure.com/",
  "activeDirectoryGraphResourceId": "https://graph.windows.net/",
  "sqlManagementEndpointUrl":
"https://management.core.windows.net:8443/",
  "galleryEndpointUrl": "https://gallery.azure.com/",
  "managementEndpointUrl": "https://management.core.windows.net/"
}
```

4. Cut and paste the output content into a file called `oracle.json` stored in the Linux user `azacsnap` user `bin` directory and secure the file with the appropriate system permissions.



Make sure the format of the JSON file is exactly as described above, especially with the URLs enclosed in double quotes (").

Complete the setup of AzAcSnap tool

Follow these steps to configure and test the snapshot tools. After successful testing, you can perform the first database-consistent storage snapshot.

1. Change into the snapshot user account.

```
su - azacsnap
```

2. Change the location of commands.

```
cd /home/azacsnap/bin/
```

3. Configure a storage backup detail file. This creates an azacsnap.json configuration file.

```
azacsnap -c configure --configuration new
```

The expected output with three Oracle volumes:

```
[azacsnap@acao-ora01 bin]$ azacsnap -c configure --configuration new
Building new config file
Add comment to config file (blank entry to exit adding comments): Oracle
snapshot bkup
Add comment to config file (blank entry to exit adding comments):
Enter the database type to add, 'hana', 'oracle', or 'exit' (for no
database): oracle

=== Add Oracle Database details ===
Oracle Database SID (e.g. CDB1): ORATST
Database Server's Address (hostname or IP address): 172.30.137.142
Oracle connect string (e.g. /@AZACSNAP): /@AZACSNAP

=== Azure NetApp Files Storage details ===
Are you using Azure NetApp Files for the database? (y/n) [n]: y
--- DATA Volumes have the Application put into a consistent state before
they are snapshot ---
Add Azure NetApp Files resource to DATA Volume section of Database
configuration? (y/n) [n]: y
Full Azure NetApp Files Storage Volume Resource ID (e.g.
/subscriptions/.../resourceGroups/.../providers/Microsoft.NetApp/netAppA
ccounts/.../capacityPools/Premium/volumes/...): /subscriptions/0efa2dfb-
917c-4497-b56a-
b3f4eadb8111/resourceGroups/ANFAVSRG/providers/Microsoft.NetApp/netAppAc
counts/ANFAVSAcct/capacityPools/CapPool/volumes/acao-ora01-u01
Service Principal Authentication filename or Azure Key Vault Resource ID
(e.g. auth-file.json or https://...): oracle.json
Add Azure NetApp Files resource to DATA Volume section of Database
configuration? (y/n) [n]: y
Full Azure NetApp Files Storage Volume Resource ID (e.g.
/subscriptions/.../resourceGroups/.../providers/Microsoft.NetApp/netAppA
```

```

ccounts/.../capacityPools/Premium/volumes/...): /subscriptions/0efa2dfb-
917c-4497-b56a-
b3f4eadb8111/resourceGroups/ANFAVSRG/providers/Microsoft.NetApp/netAppAc
counts/ANFAVSAcct/capacityPools/CapPool/volumes/acao-ora01-u02
Service Principal Authentication filename or Azure Key Vault Resource ID
(e.g. auth-file.json or https://...): oracle.json
Add Azure NetApp Files resource to DATA Volume section of Database
configuration? (y/n) [n]: n
--- OTHER Volumes are snapshot immediately without preparing any
application for snapshot ---
Add Azure NetApp Files resource to OTHER Volume section of Database
configuration? (y/n) [n]: y
Full Azure NetApp Files Storage Volume Resource ID (e.g.
/subscriptions/.../resourceGroups/.../providers/Microsoft.NetApp/netAppA
ccounts/.../capacityPools/Premium/volumes/...): /subscriptions/0efa2dfb-
917c-4497-b56a-
b3f4eadb8111/resourceGroups/ANFAVSRG/providers/Microsoft.NetApp/netAppAc
counts/ANFAVSAcct/capacityPools/CapPool/volumes/acao-ora01-u03
Service Principal Authentication filename or Azure Key Vault Resource ID
(e.g. auth-file.json or https://...): oracle.json
Add Azure NetApp Files resource to OTHER Volume section of Database
configuration? (y/n) [n]: n

=== Azure Managed Disk details ===
Are you using Azure Managed Disks for the database? (y/n) [n]: n

=== Azure Large Instance (Bare Metal) Storage details ===
Are you using Azure Large Instance (Bare Metal) for the database? (y/n)
[n]: n

Enter the database type to add, 'hana', 'oracle', or 'exit' (for no
database): exit

Editing configuration complete, writing output to 'azacsnap.json'.

```

4. As the azacsnap Linux user, run the azacsnap test command for an Oracle backup.

```

cd ~/bin
azacsnap -c test --test oracle --configfile azacsnap.json

```

The expected output:

```
[azacsnap@acao-ora01 bin]$ azacsnap -c test --test oracle --configfile
azacsnap.json
BEGIN : Test process started for 'oracle'
BEGIN : Oracle DB tests
PASSED: Successful connectivity to Oracle DB version 1908000000
END   : Test process complete for 'oracle'
[azacsnap@acao-ora01 bin]$
```

5. Run your first snapshot backup.

```
azacsnap -c backup --volume data --prefix ora_test --retention=1
```

Protect your Oracle database in Azure cloud

Allen Cao, NetApp Solutions Engineering

This section describes how to protect your Oracle database with azacsnap tool and snapshot backup, restore and snapshots tiering to Azure blob.

Backup Oracle database with snapshot using AzAcSnap tool

The Azure Application-Consistent Snapshot tool (AzAcSnap) is a command-line tool that enables data protection for third-party databases by handling all the orchestration required to put them into an application-consistent state before taking a storage snapshot, after which it returns the databases to an operational state.

In the case of Oracle, you put the database in backup mode to take a snapshot and then take the database out of backup mode.

Backup data and log volumes

The backup can be set up on the database server host with simple shell script that executes the snapshot command. Then, the script can be scheduled to run from crontab.

Generally, the frequency of backup depends on the desired RTO and RPO. Frequent snapshot creation consumes more storage space. There is a trade off between the frequency of backup and space consumption.

Data volumes typically consume more storage space than log volumes. Therefore, you can take snapshots on data volumes every few hours and more frequent snapshots on log volumes every 15 to 30 minutes.

See the following examples of backup scripts and scheduling.

For data volume snapshots:

```
# /bin/sh
cd /home/azacsnap/bin
. ~/.bash_profile
azacsnap -c backup --volume data --prefix acao-ora01-data --retention 36
azacsnap -c backup --volume other --prefix acao-ora01-log --retention 250
```

For log volume snapshots:

```
# /bin/sh
cd /home/azacsnap/bin
. ~/.bash_profile
azacsnap -c backup --volume other --prefix acao-ora01-log --retention 250
```

Crontab schedule:

```
15,30,45 * * * * /home/azacsnap/snap_log.sh
0 */2 * * * /home/azacsnap/snap_data.sh
```



When setting up the backup `azacsnap.json` configuration file, add all data volumes, including the binary volume, to `dataVolume` and all log volumes to `otherVolume`. The maximum retention of snapshots is 250 copies.

Validate the snapshots

Go to the Azure portal > Azure NetApp Files/volumes to check if the snapshots have been successfully created.

The screenshot displays the Microsoft Azure portal interface. On the left, the 'ANFAVSAcct | Volumes' section is expanded, showing a list of volumes including 'acao-ora01-u01' through 'anf2-z1-stdds08' and 'anf-z1-stdds01'. The main pane shows the 'acao-ora01-u01 (ANFAVSAcct/CapPool/acao-ora01-u01) | Snapshots' view. It includes a search bar, a table of snapshots, and a sidebar with navigation options like Overview, Activity log, Access control (IAM), Tags, Settings, Properties, Locks, Storage service, Mount instructions, Export policy, Snapshots, Replication, and Monitoring.

Name	Location	Created
acao-ora01-data_2022-09-09T163255-92586502	South Central US	09/09/2022, 12:53:22 PM
acao-ora01-data_2022-09-12T160536-98096392	South Central US	09/12/2022, 12:05:55 PM

The screenshot displays the Microsoft Azure portal interface. On the left, the 'ANFAVSAcct | Volumes' section is active, showing a list of volumes. The 'Volumes' tab is selected in the left-hand navigation pane. The main pane shows a list of volumes with columns for Name, Quota, and Location. The volume 'acao-ora01-u03' is highlighted. On the right, the 'acao-ora01-u03 (ANFAVSAcct/CapPool/acao-ora01-u03) | Snapshots' section is active, showing a list of snapshots. The 'Snapshots' tab is selected in the right-hand navigation pane. The main pane shows a list of snapshots with columns for Name, Location, Created, and Location. The snapshot 'acao-ora01-log_2022-09-12T183002-44073932' is highlighted.

Name	Location	Created	Location
acao-ora01-data_2022-09-12T160628-67547967	South Central US	09/12/2022, 12:06:31 PM	...
acao-ora01-log_2022-09-12T164501-79692522	South Central US	09/12/2022, 12:45:04 PM	...
acao-ora01-log_2022-09-12T171501-60326612	South Central US	09/12/2022, 01:15:04 PM	...
acao-ora01-log_2022-09-12T173001-47879192	South Central US	09/12/2022, 01:30:04 PM	...
acao-ora01-log_2022-09-12T174501-50106142	South Central US	09/12/2022, 01:45:04 PM	...
acao-ora01-log_2022-09-12T180053-50298742	South Central US	09/12/2022, 02:00:53 PM	...
acao-ora01-log_2022-09-12T181502-38570272	South Central US	09/12/2022, 02:15:05 PM	...
acao-ora01-log_2022-09-12T183002-44073932	South Central US	09/12/2022, 02:30:07 PM	...

Oracle restore and recovery from local backup

One of key benefits of snapshot backup is that it coexists with source database volumes, and the primary database volumes can be rolled back almost instantly.

Restore and recovery of Oracle on the primary server

The following example demonstrates how to restore and recover an Oracle database from the Azure dashboard and CLI on the same Oracle host.

1. Create a test table in the database to be restored.

```

[oracle@acao-ora01 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Mon Sep 12 19:02:35 2022
Version 19.8.0.0.0

Copyright (c) 1982, 2019, Oracle. All rights reserved.

Connected to:
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
Version 19.8.0.0.0

SQL> create table testsnapshot(
    id integer,
    event varchar(100),
    dt timestamp);

Table created.

SQL> insert into testsnapshot values(1,'insert a data marker to validate
snapshot restore',sysdate);

1 row created.

SQL> commit;

Commit complete.

SQL> select * from testsnapshot;

   ID
-----
EVENT
-----
DT
-----
---
          1
insert a data marker to validate snapshot restore
12-SEP-22 07.07.35.000000 PM

```

2. Drop the table after the snapshot backups.


```

[oracle@acao-ora01 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Tue Sep 13 14:20:22 2022
Version 19.8.0.0.0

Copyright (c) 1982, 2019, Oracle. All rights reserved.

Connected to:
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
Version 19.8.0.0.0

SQL> drop table testsnapshot;

Table dropped.

SQL> select * from testsnapshot;
select * from testsnapshot
      *
ERROR at line 1:
ORA-00942: table or view does not exist

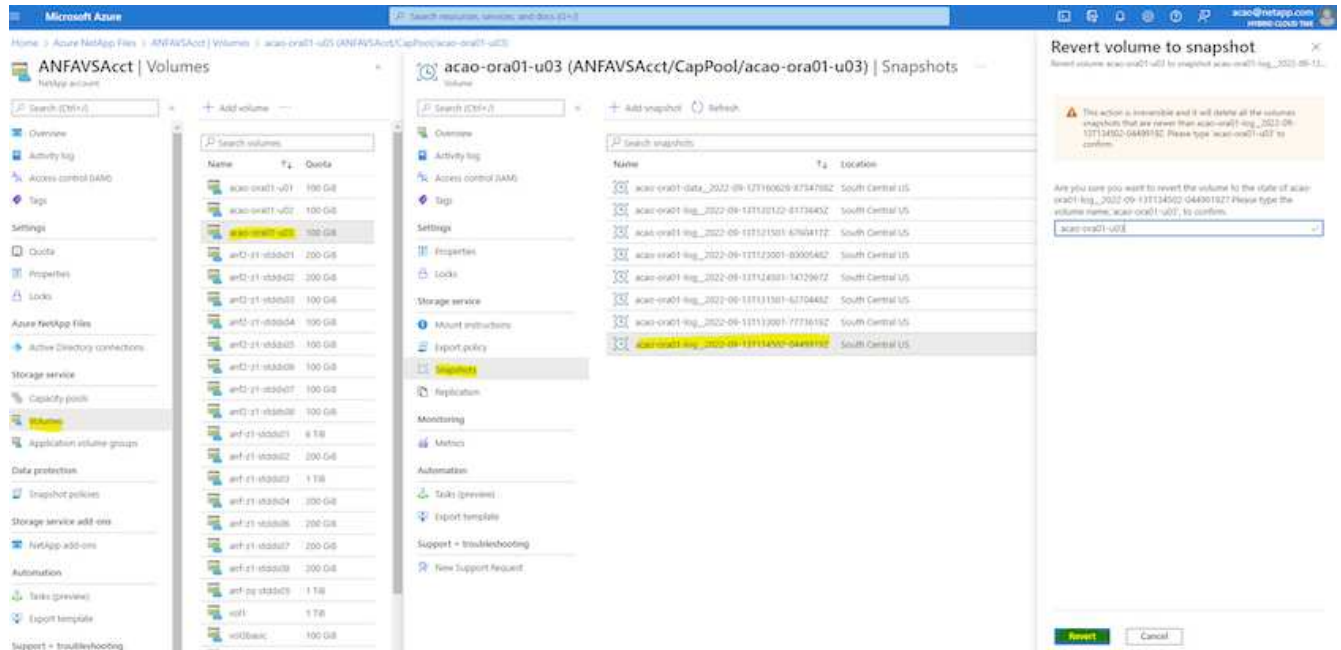
SQL> shutdown immediate;
Database closed.
Database dismounted.
ORACLE instance shut down.
SQL> exit
Disconnected from Oracle Database 19c Enterprise Edition Release
19.0.0.0.0 - Production
Version 19.8.0.0.0

```

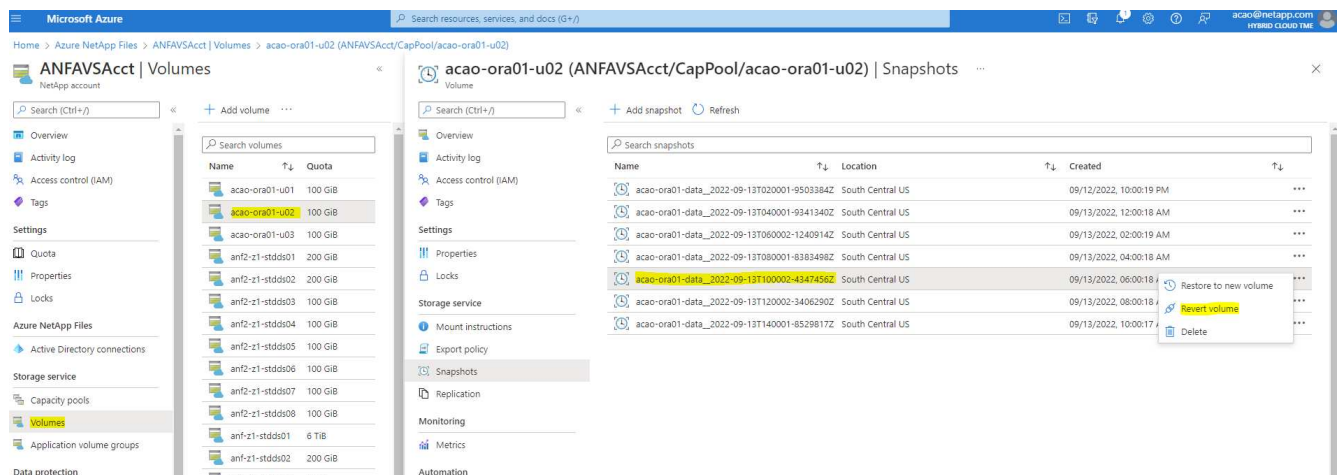
- From the Azure NetApp Files dashboard, restore the log volume to the last available snapshot. Choose **Revert volume**.

The screenshot displays the Microsoft Azure NetApp Files dashboard. On the left, the 'Volumes' section is expanded, showing a list of volumes including 'acao-ora01-u01', 'acao-ora01-u02', 'acao-ora01-u03', and several 'anf2-z1-stds' volumes. The 'Volumes' section is highlighted in yellow. In the center, the 'Snapshots' section for the selected volume 'acao-ora01-u03' is shown. It lists several snapshots, with the most recent one, 'acao-ora01-log_2022-09-13T14502-044919Z', highlighted in yellow. A context menu is open over this snapshot, showing options: 'Restore to new volume', 'Revert volume' (highlighted in yellow), and 'Delete'.


4. Confirm revert volume and click **Revert** to complete the volume reversion to the latest available backup.



5. Repeat the same steps for the data volume, and make sure that the backup contains the table to be recovered.



6. Again confirm the volume reversion, and click "Revert."



7. Resync the control files if you have multiple copies of them, and replace the old control file with the latest copy available.

```
[oracle@acao-ora01 ~]$ mv /u02/oradata/ORATST/control01.ctl
/u02/oradata/ORATST/control01.ctl.bk
[oracle@acao-ora01 ~]$ cp /u03/orareco/ORATST/control02.ctl
/u02/oradata/ORATST/control01.ctl
```

8. Log into the Oracle server VM and run database recovery with sqlplus.

```
[oracle@acao-ora01 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Tue Sep 13 15:10:17 2022
Version 19.8.0.0.0

Copyright (c) 1982, 2019, Oracle. All rights reserved.

Connected to an idle instance.

SQL> startup mount;
ORACLE instance started.

Total System Global Area 6442448984 bytes
Fixed Size 8910936 bytes
Variable Size 1090519040 bytes
Database Buffers 5335154688 bytes
Redo Buffers 7864320 bytes
Database mounted.
```

```

SQL> recover database using backup controlfile until cancel;
ORA-00279: change 3188523 generated at 09/13/2022 10:00:09 needed for
thread 1
ORA-00289: suggestion :
/u03/orareco/ORATST/archivelog/2022_09_13/o1_mf_1_43__22rnjq9q_.arc
ORA-00280: change 3188523 for thread 1 is in sequence #43

Specify log: {<RET>=suggested | filename | AUTO | CANCEL}

ORA-00279: change 3188862 generated at 09/13/2022 10:01:20 needed for
thread 1
ORA-00289: suggestion :
/u03/orareco/ORATST/archivelog/2022_09_13/o1_mf_1_44__29f2lgb5_.arc
ORA-00280: change 3188862 for thread 1 is in sequence #44
ORA-00278: log file
'/u03/orareco/ORATST/archivelog/2022_09_13/o1_mf_1_43__22rnjq9q_.arc' no
longer
needed for this recovery

Specify log: {<RET>=suggested | filename | AUTO | CANCEL}

ORA-00279: change 3193117 generated at 09/13/2022 12:00:08 needed for
thread 1
ORA-00289: suggestion :
/u03/orareco/ORATST/archivelog/2022_09_13/o1_mf_1_45__29h6qqyw_.arc
ORA-00280: change 3193117 for thread 1 is in sequence #45
ORA-00278: log file
'/u03/orareco/ORATST/archivelog/2022_09_13/o1_mf_1_44__29f2lgb5_.arc' no
longer
needed for this recovery

Specify log: {<RET>=suggested | filename | AUTO | CANCEL}

ORA-00279: change 3193440 generated at 09/13/2022 12:01:20 needed for
thread 1
ORA-00289: suggestion :
/u03/orareco/ORATST/archivelog/2022_09_13/o1_mf_1_46_%u_.arc
ORA-00280: change 3193440 for thread 1 is in sequence #46
ORA-00278: log file
'/u03/orareco/ORATST/archivelog/2022_09_13/o1_mf_1_45__29h6qqyw_.arc' no
longer
needed for this recovery

Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
cancel
Media recovery cancelled.

```

```

SQL> alter database open resetlogs;

Database altered.

SQL> select * from testsnapshot;

   ID
-----
EVENT
-----
DT
-----
---
          1
insert a data marker to validate snapshot restore
12-SEP-22 07.07.35.000000 PM

SQL> select systimestamp from dual;

SYSTIMESTAMP
-----
---
13-SEP-22 03.28.52.646977 PM +00:00

```

This screen demonstrates that the dropped table has been recovered using local snapshot backups.

Database migration from on-premises to Azure cloud

As a result of the Oracle decision to phase out single-instance databases, many organizations have converted single-instance Oracle databases to multitenant container databases. This enables the easy relocation of a subset of container databases called PDB to cloud with the maximum availability option, which minimize downtime during migration.

However, if you still have a single instance of a Oracle database, it can first be converted into a multitenant container database in place before attempting PDB relocation.

The following sections provide details for the migration of on-premises Oracle databases to Azure cloud in either scenarios.

Converting a single instance non-CDB to a PDB in a multitenant CDB

If you still have a single-instance Oracle database, it must be converted into a multitenant container database whether you wish to migrate it to the cloud or not, because Oracle will stop supporting single-instance databases some time soon.

The following procedures plug a single instance database into a container database as a pluggable database

or PDB.

1. Build a shell container database on the same host as the single-instance database in a separate ORACLE_HOME.
2. Shut down the single instance database and restart it in read-only mode.
3. Run the DBMS_PDB.DESCRIBE procedure to generate the database metadata.

```
BEGIN
  DBMS_PDB.DESCRIBE(
    pdb_descr_file => '/home/oracle/ncdb.xml');
END;
/
```

4. Shut down the single-instance database.
5. Start up the container database.
6. Run the DBMS_PDB.CHECK_PLUG_COMPATIBILITY function to determine whether the non-CDB is compatible with the CDB.

```
SET SERVEROUTPUT ON
DECLARE
  compatible CONSTANT VARCHAR2(3) :=
    CASE DBMS_PDB.CHECK_PLUG_COMPATIBILITY(
      pdb_descr_file => '/disk1/oracle/ncdb.xml',
      pdb_name        => 'NCDB')
    WHEN TRUE THEN 'YES'
    ELSE 'NO'
END;
BEGIN
  DBMS_OUTPUT.PUT_LINE(compatible);
END;
/
```

If the output is YES, then the non-CDB is compatible, and you can continue with the next step.

If the output is NO, then the non-CDB is not compatible, and you can check the PDB_PLUG_IN_VIOLATIONS view to see why it is not compatible. All violations must be corrected before you continue. For example, any version or patch mismatches should be resolved by running an upgrade or the opatch utility. After correcting the violations, run DBMS_PDB.CHECK_PLUG_COMPATIBILITY again to ensure that the non-CDB is compatible with the CDB.

7. Plug in the single instance non-CDB.

```
CREATE PLUGGABLE DATABASE ncdb USING '/home/oracle/ncdb.xml'
COPY
FILE_NAME_CONVERT = ('/disk1/oracle/dbs/', '/disk2/oracle/ncdb/')
;
```



If there is not sufficient space on the host, the `NOCOPY` option can be used to create the PDB. In that case, a single-instance non-CDB is not useable after plug in as a PDB because the original data files has been used for the PDB. Make sure to create a backup before the conversion so that there is something to fall back on if anything goes wrong.

8. Start with PDB upgrade after conversion if the version between the source single-instance non-CDB and the target CDB are different. For the same-version conversion, this step can be skipped.

```
sqlplus / as sysdba;
alter session set container=ncdb
alter pluggable database open upgrade;
exit;
dbupgrade -c ncdb -l /home/oracle
```

Review the upgrade log file in the `/home/oracle` directory.

9. Open the pluggable database, check for pdb plug-in violations, and recompile the invalid objects.

```
alter pluggable database ncdb open;
alter session set container=ncdb;
select message from pdb_plug_in_violations where type like '%ERR%' and
status <> 'RESOLVED';
$ORACLE_HOME/perl/bin/perl $ORACLE_HOME/rdbms/admin/catcon.pl -n 1 -c
'ncdb' -e -b utlrp -d $ORACLE_HOME/rdbms/admin utlrp.sql
```

10. Execute `noncdb_to_pdb.sql` to update the data dictionary.

```
sqlplus / as sysdba
alter session set container=ncdb;
@$ORACLE_HOME/rdbms/admin/noncdb_to_pdb.sql;
```

Shut down and restart the container DB. The ncdb is taken out of restricted mode.

Migrate on-premises Oracle databases to Azure with PDB relocation

Oracle PDB relocation with the maximum-availability option employs PDB hot-clone technology, which allows source PDB availability while the PDB is copying over to the target. At switchover, user connections are redirected to the target PDB automatically. Thus, downtime is minimized independent of the size of the PDB.

NetApp provides an Ansible-based toolkit that automates the migration procedure.

1. Create a CDB in the Azure public cloud on an Azure VM with the same version and patch level.
2. From the Ansible controller, clone a copy of the automation toolkit.

```
git clone https://github.com/NetApp-Automation/na_ora_aws_migration.git
```

3. Read the instruction in the README file.
4. Configure the Ansible host variable files for both the source and target Oracle servers and the DB server host's configuration file for name resolution.
5. Install the Ansible controller prerequisites on Ansible controller.

```
ansible-playbook -i hosts requirements.yml
ansible-galaxy collection install -r collections/requirements.yml
--force
```

6. Execute any pre-migration tasks against the on-premises server.

```
ansible-playbook -i hosts ora_pdb_relocate.yml -u admin -k -K -t
ora_pdb_relo_onprem
```



The admin user is the management user on the on-premises Oracle server host with sudo privileges. The admin user is authenticated with a password.

7. Execute Oracle PDB relocation from on-premises to the target Azure Oracle host.

```
ansible-playbook -i hosts ora_pdb_relocate.yml -u azureuser --private
-key db1.pem -t ora_pdb_relo_primary
```



The Ansible controller can be located either on-premises or in the Azure cloud. The controller needs connectivity to the on-premises Oracle server host and the Azure Oracle VM host. The Oracle database port (such as 1521) is open between the on-premises Oracle server host and the Azure Oracle VM host.

Additional Oracle database migration options

Please see the Microsoft documentation for additional migration options: [Oracle database migration decision process](#).

Copyright information

Copyright © 2024 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.