



Databases

NetApp Solutions

NetApp
July 31, 2024

Table of Contents

- NetApp Enterprise Database Solutions 1
 - Oracle Database 1
 - Microsoft SQL Server 582
 - Open Source Databases 678
 - SnapCenter for Databases 688
 - DB Automation Toolkits 916
 - DB Sizing Toolkits 935

NetApp Enterprise Database Solutions

Oracle Database

AWS Cloud

TR-4986: Simplified, Automated Oracle Deployment on Amazon FSx ONTAP with iSCSI

Allen Cao, Niyaz Mohamed, NetApp

This solution provides overview and details for automated Oracle deployment and protection in Amazon FSx ONTAP as primary database storage with iSCSI protocol and Oracle database configured in standalone ReStart using Oracle asm as volume manager.

Purpose

Amazon FSx for NetApp ONTAP is a storage service that allows you to launch and run fully managed NetApp ONTAP file systems in the AWS Cloud. It provides the familiar features, performance, capabilities, and APIs of NetApp file systems with the agility, scalability, and simplicity of a fully managed AWS service. It empowers you to run the most demanding database workload, such as Oracle, in the AWS cloud with peace of mind.

This documentation demonstrates the simplified deployment of Oracle databases in an Amazon FSx ONTAP file system using Ansible automation. The Oracle database is deployed in a standalone ReStart configuration with iSCSI protocol for data access and Oracle ASM for database storage disks management. It also provides information on Oracle database backup, restore, and clone using the NetApp SnapCenter UI tool for storage-efficient database operation in AWS Cloud.

This solution addresses the following use cases:

- Automated Oracle database deployment on Amazon FSx ONTAP file system
- Oracle database backup and restore on Amazon FSx ONTAP file system using NetApp SnapCenter tool
- Oracle database clone for dev/test or other use cases on Amazon FSx ONTAP file system using NetApp SnapCenter tool

Audience

This solution is intended for the following people:


- A DBA who would like to deploy Oracle on Amazon FSx ONTAP file system.
- A database solution architect who would like to test Oracle workloads on Amazon FSx ONTAP file system.
- A storage administrator who would like to deploy and manage an Oracle database on Amazon FSx ONTAP file system.
- An application owner who would like to stand up an Oracle database on Amazon FSx ONTAP file system.

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

Simplified, automated Oracle deployment on Amazon FSx ONTAP with iSCSI



Hardware and software components

Hardware

Amazon FSx ONTAP storage	Current version offered by AWS	One FSx HA cluster in the same VPC and availability zone
EC2 instance for compute	t2.xlarge/4vCPU/16G	Two EC2 T2 xlarge EC2 instances for concurrent deployment

Software

RedHat Linux	RHEL-8.6, 4.18.0-372.9.1.el8.x86_64 kernel	Deployed RedHat subscription for testing
Windows Server	2022 Standard, 10.0.20348 Build 20348	Hosting SnapCenter server
Oracle Grid Infrastructure	Version 19.18	Applied RU patch p34762026_190000_Linux-x86-64.zip
Oracle Database	Version 19.18	Applied RU patch p34765931_190000_Linux-x86-64.zip
Oracle OPatch	Version 12.2.0.1.36	Latest patch p6880880_190000_Linux-x86-64.zip
SnapCenter Server	Version 4.9P1	Workgroup deployment

Open JDK	Version java-1.8.0-openjdk.x86_64	SnapCenter plugin requirement on DB VMs
----------	-----------------------------------	---

Oracle database configuration in the lab environment

Server	Database	DB Storage
ora_01	NTAP1(NTAP1_PDB1,NTAP1_PDB2,NTAP1_PDB3)	iSCSI luns on Amazon FSx ONTAP file system
ora_02	NTAP2(NTAP2_PDB1,NTAP2_PDB2,NTAP2_PDB3)	iSCSI luns on Amazon FSx ONTAP file system

Key factors for deployment consideration

- **Oracle database storage layout.** In this automated Oracle deployment, we provision four database volumes to host Oracle binary, data, and logs by default. A single lun in a volume allocates to Oracle binary. We then create two ASM disk groups from data and logs luns. Within the +DATA asm disk group, we provision two data volumes with two luns in a volume. Within the +LOGS asm disk group, we create two luns in a log volume. Multiple luns laid out within an ONTAP volume provides better performance in general.
- **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 EC2 instance with different database instance IDs (Oracle SID). But ensure there is sufficient memory on the host to support deployed databases.
- **iSCSI configuration.** The EC2 instance database server connects to FSx storage with the iSCSI protocol. EC2 instances generally deploy with a single network interface or ENI. The single NIC interface carries both iSCSI and application traffic. It is important to gauge the Oracle database peak I/O throughput requirement by carefully analyzing the Oracle AWR report in order to choose the right EC2 compute instance that meets both application and iSCSI traffic-throughput requirements. Also, AWS EC2 generally limits each TCP flow to 5 Gbps. Each iSCSI path provides 5 Gbps (625 MBps) of bandwidth, and multiple iSCSI connections may be required to support higher throughput requirements.
- **Oracle ASM redundancy level to use for each Oracle ASM disk group that you create.** Because the Amazon FSx ONTAP is HA enabled for data protection at the cluster disk level, you should use `External Redundancy`, which means that the option does not allow Oracle ASM to mirror the contents of the disk group.
- **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 protection on Amazon FSx ONTAP file system with directly mounted database luns via iSCSI to EC2 instance VM in a single node Restart configuration with Oracle ASM as database volume manager.

Prerequisites for deployment

Deployment requires the following prerequisites.

1. An AWS account has been set up, and the necessary VPC and network segments have been created within your AWS account.
2. From the AWS EC2 console, deploy EC2 Linux instances as Oracle DB servers. Enable SSH private/public key authentication for ec2-user. See the architecture diagram in the previous section for details about the environment setup. Also review the [User Guide for Linux instances](#) for more information.
3. From the AWS FSx console, provision an Amazon FSx ONTAP file system that meets the requirements. Review the documentation [Creating FSx for ONTAP file systems](#) for step-by-step instructions.
4. Steps 2 and 3 can be performed using the following Terraform automation toolkit, which creates an EC2 instance named `ora_01` and an FSx file system named `fsx_01`. Review the instruction carefully and change the variables to suit your environment before execution. The template can be easily revised for your own deployment requirements.

```
git clone https://github.com/NetApp-
Automation/na_aws_fsx_ec2_deploy.git
```

5. Provision an EC2 Linux instance 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.
6. 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](#)
7. Clone a copy of the NetApp Oracle deployment automation toolkit for iSCSI.

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

8. Stage following Oracle 19c installation files on EC2 instances `/tmp/archive` directory.

```
installer_archives:
- "LINUX.X64_193000_grid_home.zip"
- "p34762026_190000_Linux-x86-64.zip"
- "LINUX.X64_193000_db_home.zip"
- "p34765931_190000_Linux-x86-64.zip"
- "p6880880_190000_Linux-x86-64.zip"
```



Ensure that you have allocated at least 50G in Oracle VM root volume to have sufficient space to stage Oracle installation files.

9. Watch the following video:

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


```
#####
#####

# Enter RHEL subscription to enable repo
redhat_sub_username: xxxxxxxx
redhat_sub_password: "xxxxxxx"

#####
#####
###           Oracle DB env specific config variables
###
#####
#####

# Enter Database domain name
db_domain: solutions.netapp.com

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

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
```

Playbook execution

There are a total of six 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.
3-ontap_config.yml - configure ONTAP svm/volumes/luns for Oracle
database and grant DB server access to luns.
4-oracle_config.yml - install and configure Oracle on DB servers for
grid infrastructure 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 ec2-user -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 ec2-user -e
@vars/vars.yml
```

```
ansible-playbook -i hosts 2-linux_config.yml -u ec2-user -e
@vars/vars.yml
```

```
ansible-playbook -i hosts 3-ontap_config.yml -u ec2-user -e
@vars/vars.yml
```

```
ansible-playbook -i hosts 4-oracle_config.yml -u ec2-user -e
@vars/vars.yml
```

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

```
ansible-playbook -i hosts 0-all_playbook.yml -u ec2-user -e
@vars/vars.yml -t ansible_requirements
```

```
ansible-playbook -i hosts 0-all_playbook.yml -u ec2-user -e
@vars/vars.yml -t linux_config
```

```
ansible-playbook -i hosts 0-all_playbook.yml -u ec2-user -e
@vars/vars.yml -t ontap_config
```

```
ansible-playbook -i hosts 0-all_playbook.yml -u ec2-user -e
@vars/vars.yml -t oracle_config
```

4. Undo the environment

```
ansible-playbook -i hosts 5-destroy.yml -u ec2-user -e
@vars/vars.yml
```

Post execution validation

After the playbook run, login to the Oracle DB server as oracle user to validate that Oracle grid infrastructure and database are created successfully. Following is an example of Oracle database validation on host ora_01.

1. Validate Oracle container database on EC2 instance

```
[admin@ansiblectl na_oracle_deploy_iscsi]$ ssh -i ora_01.pem ec2-
user@172.30.15.40
Last login: Fri Dec  8 17:14:21 2023 from 10.61.180.18
[ec2-user@ip-172-30-15-40 ~]$ uname -a
Linux ip-172-30-15-40.ec2.internal 4.18.0-372.9.1.el8.x86_64 #1 SMP
Fri Apr 15 22:12:19 EDT 2022 x86_64 x86_64 x86_64 GNU/Linux

[ec2-user@ip-172-30-15-40 ~]$ sudo su
[root@ip-172-30-15-40 ec2-user]# su - oracle
Last login: Fri Dec  8 16:25:52 UTC 2023 on pts/0
[oracle@ip-172-30-15-40 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Fri Dec 8 18:18:20 2023
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
          5 NTAP1_PDB3      READ WRITE NO

SQL> select name from v$datafile;

NAME
```

```
-----  
+DATA/NTAP1/DATAFILE/system.257.1155055419  
+DATA/NTAP1/DATAFILE/sysaux.258.1155055463  
+DATA/NTAP1/DATAFILE/undotbs1.259.1155055489  
+DATA/NTAP1/86B637B62FE07A65E053F706E80A27CA/DATAFILE/system.266.115  
5056241  
+DATA/NTAP1/86B637B62FE07A65E053F706E80A27CA/DATAFILE/sysaux.267.115  
5056241  
+DATA/NTAP1/DATAFILE/users.260.1155055489  
+DATA/NTAP1/86B637B62FE07A65E053F706E80A27CA/DATAFILE/undotbs1.268.1  
155056241  
+DATA/NTAP1/0C03AAFA7C6FD2E5E063280F1EACFBE0/DATAFILE/system.272.115  
5057059  
+DATA/NTAP1/0C03AAFA7C6FD2E5E063280F1EACFBE0/DATAFILE/sysaux.273.115  
5057059  
+DATA/NTAP1/0C03AAFA7C6FD2E5E063280F1EACFBE0/DATAFILE/undotbs1.271.1  
155057059  
+DATA/NTAP1/0C03AAFA7C6FD2E5E063280F1EACFBE0/DATAFILE/users.275.1155  
057075
```

NAME

```
-----  
+DATA/NTAP1/0C03AC0089ACD352E063280F1EAC12BD/DATAFILE/system.277.115  
5057075  
+DATA/NTAP1/0C03AC0089ACD352E063280F1EAC12BD/DATAFILE/sysaux.278.115  
5057075  
+DATA/NTAP1/0C03AC0089ACD352E063280F1EAC12BD/DATAFILE/undotbs1.276.1  
155057075  
+DATA/NTAP1/0C03AC0089ACD352E063280F1EAC12BD/DATAFILE/users.280.1155  
057091  
+DATA/NTAP1/0C03ACEABA54D386E063280F1EACE573/DATAFILE/system.282.115  
5057091  
+DATA/NTAP1/0C03ACEABA54D386E063280F1EACE573/DATAFILE/sysaux.283.115  
5057091  
+DATA/NTAP1/0C03ACEABA54D386E063280F1EACE573/DATAFILE/undotbs1.281.1  
155057091  
+DATA/NTAP1/0C03ACEABA54D386E063280F1EACE573/DATAFILE/users.285.1155  
057105
```

19 rows selected.

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

NAME

```
-----  
+DATA/NTAP1/CONTROLFILE/current.261.1155055529  
+LOGS/NTAP1/CONTROLFILE/current.256.1155055529
```

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

```
MEMBER  
-----
```

```
-----  
+DATA/NTAP1/ONLINELOG/group_3.264.1155055531  
+LOGS/NTAP1/ONLINELOG/group_3.259.1155055539  
+DATA/NTAP1/ONLINELOG/group_2.263.1155055531  
+LOGS/NTAP1/ONLINELOG/group_2.257.1155055539  
+DATA/NTAP1/ONLINELOG/group_1.262.1155055531  
+LOGS/NTAP1/ONLINELOG/group_1.258.1155055539
```

```
6 rows selected.
```

```
SQL> exit
```

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

2. Validate Oracle listener.

```
[oracle@ip-172-30-15-40 ~]$ lsnrctl status listener
```

```
LSNRCTL for Linux: Version 19.0.0.0.0 - Production on 08-DEC-2023  
18:20:24
```

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

```
Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=TCP) (HOST=ip-172-30-  
15-40.ec2.internal) (PORT=1521)))
```

```
STATUS of the LISTENER
```

```
-----  
Alias                LISTENER  
Version              TNSLSNR for Linux: Version 19.0.0.0.0 -  
Production  
Start Date           08-DEC-2023 16:26:09  
Uptime               0 days 1 hr. 54 min. 14 sec  
Trace Level          off  
Security             ON: Local OS Authentication  
SNMP                 OFF  
Listener Parameter File
```

```
/u01/app/oracle/product/19.0.0/grid/network/admin/listener.ora
Listener Log File          /u01/app/oracle/diag/tnslsnr/ip-172-30-15-
40/listener/alert/log.xml
Listening Endpoints Summary...
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp) (HOST=ip-172-30-15-
40.ec2.internal) (PORT=1521)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=ipc) (KEY=EXTPROC1521)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcps) (HOST=ip-172-30-15-
40.ec2.internal) (PORT=5500)) (Security=(my_wallet_directory=/u01/app/
oracle/product/19.0.0/NTAP1/admin/NTAP1/xdb_wallet)) (Presentation=HT
TP) (Session=RAW))
Services Summary...
Service "+ASM" has 1 instance(s).
  Instance "+ASM", status READY, has 1 handler(s) for this
service...
Service "+ASM_DATA" has 1 instance(s).
  Instance "+ASM", status READY, has 1 handler(s) for this
service...
Service "+ASM_LOGS" has 1 instance(s).
  Instance "+ASM", status READY, has 1 handler(s) for this
service...
Service "0c03aafa7c6fd2e5e063280f1eacfb0.solutions.netapp.com" has
1 instance(s).
  Instance "NTAP1", status READY, has 1 handler(s) for this
service...
Service "0c03ac0089acd352e063280f1eac12bd.solutions.netapp.com" has
1 instance(s).
  Instance "NTAP1", status READY, has 1 handler(s) for this
service...
Service "0c03aceaba54d386e063280f1eace573.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 the grid infrastructure and resources created.

```
[oracle@ip-172-30-15-40 ~]$ asm
[oracle@ip-172-30-15-40 ~]$ crsctl check has
CRS-4638: Oracle High Availability Services is online
[oracle@ip-172-30-15-40 ~]$ crsctl stat res -t
-----
-----
Name          Target  State         Server          State
details
-----
-----
Local Resources
-----
-----
ora.DATA.dg
          ONLINE  ONLINE        ip-172-30-15-40  STABLE
ora.LISTENER.lsnr
          ONLINE  ONLINE        ip-172-30-15-40  STABLE
ora.LOGS.dg
          ONLINE  ONLINE        ip-172-30-15-40  STABLE
ora.asm
          ONLINE  ONLINE        ip-172-30-15-40  Started,STABLE
ora.ons
          OFFLINE OFFLINE        ip-172-30-15-40  STABLE
-----
-----
Cluster Resources
-----
-----
ora.cssd
   1      ONLINE  ONLINE        ip-172-30-15-40  STABLE
ora.diskmon
   1      OFFLINE OFFLINE        ip-172-30-15-40  STABLE
ora.driver.afd
   1      ONLINE  ONLINE        ip-172-30-15-40  STABLE
ora.evmd
   1      ONLINE  ONLINE        ip-172-30-15-40  STABLE
ora.ntap1.db
   1      ONLINE  ONLINE        ip-172-30-15-40
```



```
Open,HOME=/u01/app/o
```

```
racle/product/19.0.0
```

```
/NTAP1, STABLE
```

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

4. Validate Oracle ASM.

```
[oracle@ip-172-30-15-40 ~]$ asmcmd  
ASMCMD> lsdg  
State      Type      Rebal  Sector  Logical_Sector  Block      AU  
Total_MB  Free_MB  Req_mir_free_MB  Usable_file_MB  Offline_disks  
Voting_files  Name  
MOUNTED  EXTERN  N      512     512     4096    4194304  
163840   155376      0      155376      0  
N  DATA/  
MOUNTED  EXTERN  N      512     512     4096    4194304  
81920   80972      0      80972      0  
N  LOGS/  
ASMCMDB> lsdsk  
Path  
AFD:ORA_01_DAT1_01  
AFD:ORA_01_DAT1_03  
AFD:ORA_01_DAT2_02  
AFD:ORA_01_DAT2_04  
AFD:ORA_01_LOGS_01  
AFD:ORA_01_LOGS_02  
ASMCMDB> afd_state  
ASMCMDB-9526: The AFD state is 'LOADED' and filtering is 'ENABLED' on  
host 'ip-172-30-15-40.ec2.internal'  
ASMCMDB> exit
```

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

Not secure | <https://172.30.15.40:5500/em/login>

ORACLE ENTERPRISE MANAGER DATABASE EXPRESS

Username

Password

Container Name

[Log in](#)

ORACLE

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

Not secure | <https://172.30.15.40:5500/em/shell>

ORACLE Enterprise Manager Database Express

NTAP1 (19.18.0.0.0) Performance Storage

Database Home

Time Zone: Browser (GMT-00:00) 1 min Auto-Refresh Refresh

Status

Up Time 1 hours, 21 minutes, 12 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) 5

Performance


Activity Services Containers

Resources

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
--------	----------	--------	---------------	-----------	----------	---------------	--------------	----------



Oracle backup, restore, and clone with SnapCenter

Refer to TR-4979 [Simplified, self-managed Oracle in VMware Cloud on AWS with guest-mounted FSx ONTAP](#) section Oracle backup, restore, and clone with SnapCenter for details on setting up SnapCenter and executing the database backup, restore, and clone workflows.

Where to find additional information

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

- Amazon FSx for NetApp ONTAP

<https://aws.amazon.com/fsx/netapp-ontap/>

- Amazon EC2

https://aws.amazon.com/pm/ec2/?trk=36c6da98-7b20-48fa-8225-4784bced9843&sc_channel=ps&s_kwcid=AL14422!3!467723097970!e!lg!!aws%20ec2&ef_id=Cj0KCQiA54KfBhCKARIsAJzSrdqwQrghn6!7!1jiWzSeaT9Uh1-vY-VfhJixF-xnv5rWwn2S7RqZOTQ0aAh7eEALw_wcB:G:s&s_kwcid=AL14422!3!467723097970!e!lg!!aws%20ec2

- Installing Oracle Grid Infrastructure for a Standalone Server with a New Database Installation

<https://docs.oracle.com/en/database/oracle/oracle-database/19/ladbi/installing-oracle-grid-infrastructure-for-a-standalone-server-with-a-new-database-installation.html#GUID-0B1CEE8C-C893-46AA-8A6A-7B5FAAEC72B3>

- Installing and Configuring Oracle Database Using Response Files

<https://docs.oracle.com/en/database/oracle/oracle-database/19/ladbi/installing-and-configuring-oracle->

- Use Red Hat Enterprise Linux 8.2 with ONTAP

https://docs.netapp.com/us-en/ontap-sanhost/hu_rhel_82.html#all-san-array-configurations

TR-4979: Simplified, Self-managed Oracle in VMware Cloud on AWS with guest-mounted FSx ONTAP

Allen Cao, Niyaz Mohamed, NetApp

This solution provides overview and details for Oracle deployment and protection in VMware Cloud in AWS with FSx ONTAP as primary database storage and Oracle database configured in standalone ReStart using asm as volume manager.

Purpose

Enterprises have been running Oracle on VMware in private data centers for decades. VMware Cloud (VMC) on AWS provides a push-button solution to bring VMware's enterprise-class Software-Defined Data Center (SDDC) software to the AWS Cloud's dedicated, elastic, bare-metal infrastructure. AWS FSx ONTAP offers premium storage to VMC SDDC and a data fabric that enables customers to run business-critical applications such as Oracle across vSphere®-based private, public, and hybrid cloud environments, with optimized access to AWS services. Whether it is an existing or new Oracle workload, VMC on AWS provides a familiar, simplified, and self-managed Oracle environment on VMware with all the benefits of AWS cloud while deferring all platform management and optimization to VMware.

This documentation demonstrates the deployment and protection of an Oracle database in a VMC environment with Amazon FSx ONTAP as primary database storage. Oracle database can be deployed to VMC on FSx storage as direct VM guest-mounted LUNs or NFS-mounted VMware VMDK datastore disks. This technical report focuses on Oracle database deployment as direct guest-mounted FSx storage to VMs in the VMC cluster with the iSCSI protocol and Oracle ASM. We also demonstrate how to use the NetApp SnapCenter UI tool to backup, restore, and clone an Oracle database for dev/test or other use cases for storage-efficient database operation in the VMC on AWS.

This solution addresses the following use cases:

- Oracle database deployment in VMC on AWS with Amazon FSx ONTAP as primary database storage
- Oracle database backup and restore in VMC on AWS using NetApp SnapCenter tool
- Oracle database clone for dev/test or other use cases in VMC on AWS using NetApp SnapCenter tool

Audience

This solution is intended for the following people:

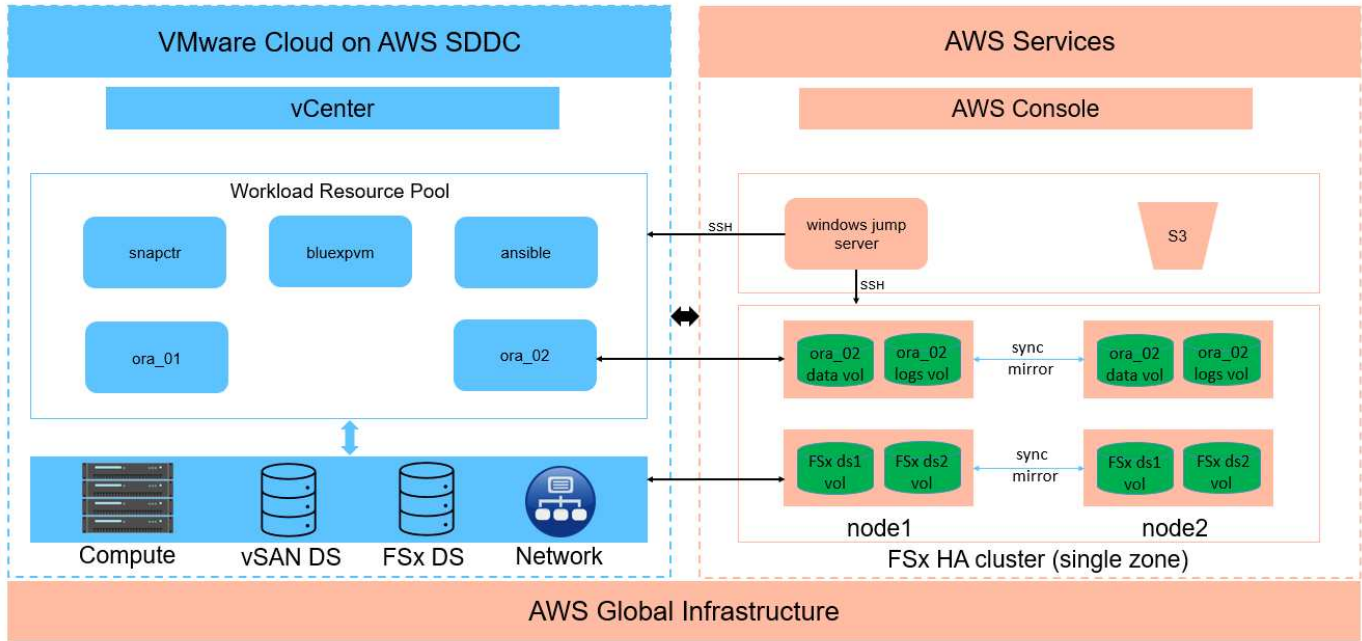
- A DBA who would like to deploy Oracle in VMC on AWS with Amazon FSx ONTAP
- A database solution architect who would like to test Oracle workloads in VMC on the AWS cloud
- A storage administrator who would like to deploy and manage an Oracle database deployed to VMC on AWS with Amazon FSx ONTAP
- An application owner who would like to stand up an Oracle database in VMC on the AWS cloud

Solution test and validation environment

The testing and validation of this solution was performed in a lab environment with VMC on AWS that might not match the final deployment environment. For more information, see the section [Key factors for deployment consideration](#).

Architecture

Oracle Database Deployment in VMware Cloud on AWS with Amazon FSx ONTAP



NetApp

Hardware and software components

Hardware

FSx ONTAP storage	Current version offered by AWS	One FSx ONTAP HA cluster in the same VPC and availability zone as VMC
VMC SDDC cluster	Amazon EC2 i3.metal single node/Intel Xeon E5-2686 CPU,36 cores/512G RAM	10.37 TB vSAN storage

Software

RedHat Linux	RHEL-8.6, 4.18.0-372.9.1.el8.x86_64 kernel	Deployed RedHat subscription for testing
Windows Server	2022 Standard, 10.0.20348 Build 20348	Hosting SnapCenter server
Oracle Grid Infrastructure	Version 19.18	Applied RU patch p34762026_190000_Linux-x86-64.zip

Oracle Database	Version 19.18	Applied RU patch p34765931_190000_Linux-x86-64.zip
Oracle OPatch	Version 12.2.0.1.36	Latest patch p6880880_190000_Linux-x86-64.zip
SnapCenter Server	Version 4.9P1	Workgroup deployment
BlueXP backup and recovery for VMs	Release 1.0	Deployed as an ova vSphere plugin VM
VMware vSphere	Version 8.0.1.00300	VMware Tools, Version: 11365 - Linux, 12352 - Windows
Open JDK	Version java-1.8.0-openjdk.x86_64	SnapCenter plugin requirement on DB VMs

Oracle database configuration in VMC on AWS

Server	Database	DB Storage
ora_01	cdb1(cdb1_pdb1,cdb1_pdb2,cdb1_pdb3)	VMDK datastore on FSx ONTAP
ora_01	cdb2(cdb2_pdb)	VMDK datastore on FSx ONTAP
ora_02	cdb3(cdb3_pdb1,cdb3_pdb2,cdb3_pdb3)	Direct guest mounted FSx ONTAP
ora_02	cdb4(cdb4_pdb)	Direct guest mounted FSx ONTAP

Key factors for deployment consideration

- **FSx to VMC connectivity.** When you deploy your SDDC on VMware Cloud on AWS, it is created within an AWS account and a VPC dedicated to your organization and managed by VMware. You must also connect the SDDC to an AWS account belonging to you, called the customer AWS account. This connection allows your SDDC to access AWS services belonging to your customer account. FSx for ONTAP is an AWS service deployed in your customer account. Once the VMC SDDC is connected to your customer account, FSx storage is available to VMs in VMC SDDC for direct guest mount.
- **FSx storage HA clusters single- or multi-zone deployment.** In these tests and validations, we deployed an FSx HA cluster in a single AWS availability zone. NetApp also recommends deploying FSx for NetApp ONTAP and VMware Cloud on AWS in the same availability zone to achieve better performance and avoid data transfer charges between availability zones.
- **FSx storage cluster sizing.** An Amazon FSx for ONTAP storage file system provides up to 160,000 raw SSD IOPS, up to 4GBps throughput, and a maximum of 192TiB capacity. However, you can size the cluster in terms of provisioned IOPS, throughput, and storage limit (minimum 1,024 GiB) based on your actual requirements at the time of deployment. The capacity can be adjusted dynamically on the fly without affecting application availability.
- **Oracle data and logs layout.** In our tests and validations, we deployed two ASM disk groups for data and logs respectively. Within the +DATA asm disk group, we provisioned four LUNs in a data volume. Within the +LOGS asm disk group, we provisioned two LUNs in a log volume. In general, multiple LUNs laid out within an Amazon FSx for ONTAP volume provide better performance.

- **iSCSI configuration.** The database VMs in VMC SDDC connect to FSx storage with the iSCSI protocol. It is important to gauge the Oracle database peak I/O throughput requirement by carefully analyzing the Oracle AWR report to determine the application and iSCSI traffic-throughput requirements. NetApp also recommends allocating four iSCSI connections to both FSx iSCSI endpoints with multipath properly configured.
- **Oracle ASM redundancy level to use for each Oracle ASM disk group that you create.** Because FSx ONTAP already mirrors the storage on the FSx cluster level, you should use External Redundancy, which means that the option does not allow Oracle ASM to mirror the contents of the disk group.
- **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 Oracle 19c deployment in VMC on AWS with directly mounted FSx ONTAP storage to DB VM in a single node Restart configuration with Oracle ASM as database volume manager.

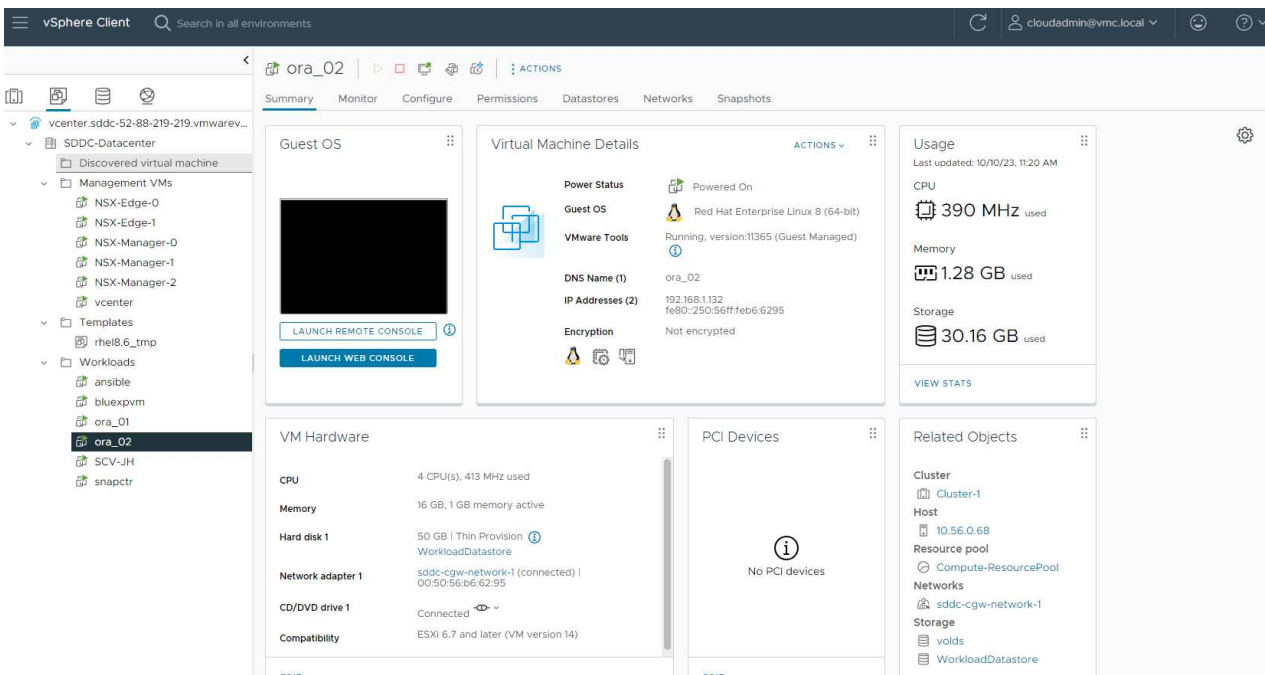
Prerequisites for deployment

Deployment requires the following prerequisites.

1. A software-defined data center (SDDC) using VMware Cloud on AWS has been created. For detailed instruction on how to create an SDDC in VMC, please refer to VMware documentation [Getting Started With VMware Cloud on AWS](#)
2. An AWS account has been set up, and the necessary VPC and network segments have been created within your AWS account. The AWS account is linked to your VMC SDDC.
3. From the AWS EC2 console, deploying an Amazon FSx for ONTAP storage HA clusters to host the Oracle database volumes. If you are not familiar with the deployment of FSx storage, see the documentation [Creating FSx for ONTAP file systems](#) for step-by-step instructions.
4. The above step can be performed using the following Terraform automation toolkit, which creates an EC2 instance as a jump host for SDDC in VMC access via SSH and an FSx file system. Review instructions carefully and change the variables to suit your environment before execution.

```
git clone https://github.com/NetApp-Automation/na_aws_fsx_ec2_deploy.git
```

5. Build VMs in VMware SDDC on AWS for hosting your Oracle environment to be deployed in VMC. In our demonstration, we have built two Linux VMs as Oracle DB servers, one Windows server for the SnapCenter server, and one optional Linux server as an Ansible controller for automated Oracle installation or configuration if desired. Following is a snapshot of the lab environment for the solution validation.



6. Optionally, NetApp also provides several automation toolkits to run Oracle deployment and configuration when applicable. Refer to [DB Automation Toolkits](#) for more information.



Ensure that you have allocated at least 50G in Oracle VM root volume in order to have sufficient space to stage Oracle installation files.

DB VM kernel configuration

With the prerequisites provisioned, login to the Oracle VM as an admin user via SSH and sudo to the root user to configure the Linux kernel for Oracle installation. Oracle install files can be staged in an AWS S3 bucket and transferred into the VM.

1. Create a staging directory `/tmp/archive` folder and set the `777` permission.

```
mkdir /tmp/archive
```

```
chmod 777 /tmp/archive
```

2. Download and stage the Oracle binary installation files and other required rpm files to the `/tmp/archive` directory.

See the following list of installation files to be staged in `/tmp/archive` on the DB VM.

```
[admin@ora_02 ~]$ ls -l /tmp/archive/
total 10539364
-rw-rw-r--. 1 admin admin      19112 Oct  4 17:04 compat-
libcap1-1.10-7.el7.x86_64.rpm
-rw-rw-r--. 1 admin admin  3059705302 Oct  4 17:10
LINUX.X64_193000_db_home.zip
-rw-rw-r--. 1 admin admin  2889184573 Oct  4 17:11
LINUX.X64_193000_grid_home.zip
-rw-rw-r--. 1 admin admin    589145 Oct  4 17:04
netapp_linux_unified_host_utilities-7-1.x86_64.rpm
-rw-rw-r--. 1 admin admin    31828 Oct  4 17:04 oracle-
database-preinstall-19c-1.0-2.el8.x86_64.rpm
-rw-rw-r--. 1 admin admin  2872741741 Oct  4 17:12
p34762026_190000_Linux-x86-64.zip
-rw-rw-r--. 1 admin admin  1843577895 Oct  4 17:13
p34765931_190000_Linux-x86-64.zip
-rw-rw-r--. 1 admin admin  124347218 Oct  4 17:13
p6880880_190000_Linux-x86-64.zip
-rw-rw-r--. 1 admin admin    257136 Oct  4 17:04
policycoreutils-python-utils-2.9-9.el8.noarch.rpm
[admin@ora_02 ~]$
```

3. Install Oracle 19c preinstall RPM, which satisfies most kernel configuration requirements.

```
yum install /tmp/archive/oracle-database-preinstall-19c-1.0-
2.el8.x86_64.rpm
```

4. Download and install the missing `compat-libcap1` in Linux 8.

```
yum install /tmp/archive/compat-libcap1-1.10-7.el7.x86_64.rpm
```

5. From NetApp, download and install NetApp host utilities.

```
yum install /tmp/archive/netapp_linux_unified_host_utilities-7-1.x86_64.rpm
```

6. Install `policycoreutils-python-utils`.

```
yum install /tmp/archive/policycoreutils-python-utils-2.9-9.el8.noarch.rpm
```

7. Install open JDK version 1.8.

```
yum install java-1.8.0-openjdk.x86_64
```

8. Install iSCSI initiator utils.

```
yum install iscsi-initiator-utils
```

9. Install `sg3_utils`.

```
yum install sg3_utils
```

10. Install `device-mapper-multipath`.

```
yum install device-mapper-multipath
```

11. Disable transparent hugepages in the current system.

```
echo never > /sys/kernel/mm/transparent_hugepage/enabled
```

```
echo never > /sys/kernel/mm/transparent_hugepage/defrag
```

12. Add the following lines in `/etc/rc.local` to disable `transparent_hugepage` after reboot.

```
vi /etc/rc.local
```

```
# Disable transparent hugepages
    if test -f /sys/kernel/mm/transparent_hugepage/enabled;
then
    echo never > /sys/kernel/mm/transparent_hugepage/enabled
fi
    if test -f /sys/kernel/mm/transparent_hugepage/defrag;
then
    echo never > /sys/kernel/mm/transparent_hugepage/defrag
fi
```

13. Disable selinux by changing `SELINUX=enforcing` to `SELINUX=disabled`. You must reboot the host to make the change effective.

```
vi /etc/sysconfig/selinux
```

14. Add the following lines to `limit.conf` to set the file descriptor limit and stack size.

```
vi /etc/security/limits.conf
```

```
*          hard    nofile    65536
*          soft    stack     10240
```

15. Add swap space to DB VM if there is no swap space configured with this instruction: [How do I allocate memory to work as swap space in an Amazon EC2 instance by using a swap file?](#) The exact amount of space to add depends on the size of RAM up to 16G.

16. Change `node.session.timeo.replacement_timeout` in the `iscsi.conf` configuration file from 120 to 5 seconds.

```
vi /etc/iscsi/iscsid.conf
```

17. Enable and start the iSCSI service on the EC2 instance.

```
systemctl enable iscsid
```

```
systemctl start iscsid
```

18. Retrieve the iSCSI initiator address to be used for database LUN mapping.

```
cat /etc/iscsi/initiatorname.iscsi
```

19. Add the asm groups for asm management user (oracle).

```
groupadd asmadmin
```

```
groupadd asmdba
```

```
groupadd asmoper
```

20. Modify the oracle user to add asm groups as secondary groups (the oracle user should have been created after Oracle preinstall RPM installation).

```
usermod -a -G asmadmin oracle
```

```
usermod -a -G asmdba oracle
```

```
usermod -a -G asmoper oracle
```

21. Stop and disable the Linux firewall if it is active.

```
systemctl stop firewalld
```

```
systemctl disable firewalld
```

22. Enable password-less sudo for admin user by uncommenting # %wheel ALL=(ALL) NOPASSWD: ALL line in /etc/sudoers file. Change the file permission to make the edit.

```
chmod 640 /etc/sudoers
```

```
vi /etc/sudoers
```

```
chmod 440 /etc/sudoers
```

23. Reboot the EC2 instance.

Provision and map FSx ONTAP LUNs to the DB VM

Provision three volumes from the command line by login to FSx cluster as fsxadmin user via ssh and FSx cluster management IP. Create LUNs within the volumes to host the Oracle database binary, data, and logs files.

1. Log into the FSx cluster through SSH as the fsxadmin user.

```
ssh fsxadmin@10.49.0.74
```

2. Execute the following command to create a volume for the Oracle binary.

```
vol create -volume ora_02_biny -aggregate aggr1 -size 50G -state  
online -type RW -snapshot-policy none -tiering-policy snapshot-only
```

3. Execute the following command to create a volume for Oracle data.

```
vol create -volume ora_02_data -aggregate aggr1 -size 100G -state  
online -type RW -snapshot-policy none -tiering-policy snapshot-only
```

4. Execute the following command to create a volume for Oracle logs.

```
vol create -volume ora_02_logs -aggregate aggr1 -size 100G -state  
online -type RW -snapshot-policy none -tiering-policy snapshot-only
```

5. Validate the volumes created.

```
vol show ora*
```

Output from the command:

```
FsxId0c00cec8dad373fd1::> vol show ora*  
Vserver   Volume           Aggregate      State      Type      Size  
Available Used%  
-----  
nim       ora_02_biny     aggr1         online    RW        50GB  
22.98GB  51%  
nim       ora_02_data     aggr1         online    RW        100GB  
18.53GB  80%  
nim       ora_02_logs     aggr1         online    RW        50GB  
7.98GB   83%
```

6. Create a binary LUN within the database binary volume.

```
lun create -path /vol/ora_02_biny/ora_02_biny_01 -size 40G -ostype linux
```

7. Create data LUNs within the database data volume.

```
lun create -path /vol/ora_02_data/ora_02_data_01 -size 20G -ostype linux
```

```
lun create -path /vol/ora_02_data/ora_02_data_02 -size 20G -ostype linux
```

```
lun create -path /vol/ora_02_data/ora_02_data_03 -size 20G -ostype linux
```

```
lun create -path /vol/ora_02_data/ora_02_data_04 -size 20G -ostype linux
```

8. Create log LUNs within the database logs volume.

```
lun create -path /vol/ora_02_logs/ora_02_logs_01 -size 40G -ostype linux
```

```
lun create -path /vol/ora_02_logs/ora_02_logs_02 -size 40G -ostype linux
```

9. Create an igroup for the EC2 instance with the initiator retrieved from step 14 of the EC2 kernel configuration above.

```
igroup create -igroup ora_02 -protocol iscsi -ostype linux  
-initiator iqn.1994-05.com.redhat:f65fed7641c2
```

10. Map the LUNs to the igroup created above. Increment the LUN ID sequentially for each additional LUN.


```

lun map -path /vol/ora_02_biny/ora_02_biny_01 -igroup ora_02
-vserver svm_ora -lun-id 0
lun map -path /vol/ora_02_data/ora_02_data_01 -igroup ora_02
-vserver svm_ora -lun-id 1
lun map -path /vol/ora_02_data/ora_02_data_02 -igroup ora_02
-vserver svm_ora -lun-id 2
lun map -path /vol/ora_02_data/ora_02_data_03 -igroup ora_02
-vserver svm_ora -lun-id 3
lun map -path /vol/ora_02_data/ora_02_data_04 -igroup ora_02
-vserver svm_ora -lun-id 4
lun map -path /vol/ora_02_logs/ora_02_logs_01 -igroup ora_02
-vserver svm_ora -lun-id 5
lun map -path /vol/ora_02_logs/ora_02_logs_02 -igroup ora_02
-vserver svm_ora -lun-id 6

```

11. Validate the LUN mapping.

```
mapping show
```

This is expected to return:

```

FsxId0c00cec8dad373fd1::> mapping show
(lun mapping show)
Vserver      Path                                          Igroup    LUN ID
Protocol
-----
nim          /vol/ora_02_biny/ora_02_u01_01            ora_02     0
iscsi
nim          /vol/ora_02_data/ora_02_u02_01            ora_02     1
iscsi
nim          /vol/ora_02_data/ora_02_u02_02            ora_02     2
iscsi
nim          /vol/ora_02_data/ora_02_u02_03            ora_02     3
iscsi
nim          /vol/ora_02_data/ora_02_u02_04            ora_02     4
iscsi
nim          /vol/ora_02_logs/ora_02_u03_01            ora_02     5
iscsi
nim          /vol/ora_02_logs/ora_02_u03_02            ora_02     6
iscsi

```

DB VM storage configuration

Now, import and set up the FSx ONTAP storage for the Oracle grid infrastructure and database installation on the VMC database VM.

1. Login to the DB VM via SSH as the admin user using Putty from Windows jump server.
2. Discover the FSx iSCSI endpoints using either SVM iSCSI IP address. Change to your environment-specific portal address.

```
sudo iscsiadm iscsiadm --mode discovery --op update --type  
sendtargets --portal 10.49.0.12
```

3. Establish iSCSI sessions by logging into each target.

```
sudo iscsiadm --mode node -l all
```

The expected output from the command is:

```
[ec2-user@ip-172-30-15-58 ~]$ sudo iscsiadm --mode node -l all  
Logging in to [iface: default, target: iqn.1992-  
08.com.netapp:sn.1f795e65c74911edb785affbf0a2b26e:vs.3, portal:  
10.49.0.12,3260]  
Logging in to [iface: default, target: iqn.1992-  
08.com.netapp:sn.1f795e65c74911edb785affbf0a2b26e:vs.3, portal:  
10.49.0.186,3260]  
Login to [iface: default, target: iqn.1992-  
08.com.netapp:sn.1f795e65c74911edb785affbf0a2b26e:vs.3, portal:  
10.49.0.12,3260] successful.  
Login to [iface: default, target: iqn.1992-  
08.com.netapp:sn.1f795e65c74911edb785affbf0a2b26e:vs.3, portal:  
10.49.0.186,3260] successful.
```

4. View and validate a list of active iSCSI sessions.

```
sudo iscsiadm --mode session
```

Return the iSCSI sessions.

```
[ec2-user@ip-172-30-15-58 ~]$ sudo iscsiadm --mode session  
tcp: [1] 10.49.0.186:3260,1028 iqn.1992-  
08.com.netapp:sn.545a38bf06ac11ee8503e395ab90d704:vs.3 (non-flash)  
tcp: [2] 10.49.0.12:3260,1029 iqn.1992-  
08.com.netapp:sn.545a38bf06ac11ee8503e395ab90d704:vs.3 (non-flash)
```

5. Verify that the LUNs were imported into the host.

```
sudo sanlun lun show
```

This will return a list of Oracle LUNs from FSx.

```
[admin@ora_02 ~]$ sudo sanlun lun show
controller(7mode/E-Series)/
device          host          lun
vserver(cDOT/FlashRay)
filename        adapter      protocol    size    product
-----
nim             /vol/ora_02_logs/ora_02_u03_02
/dev/sdo       host34       iSCSI       20g    cDOT
nim             /vol/ora_02_logs/ora_02_u03_01
/dev/sdn       host34       iSCSI       20g    cDOT
nim             /vol/ora_02_data/ora_02_u02_04
/dev/sdm       host34       iSCSI       20g    cDOT
nim             /vol/ora_02_data/ora_02_u02_03
/dev/sdl       host34       iSCSI       20g    cDOT
nim             /vol/ora_02_data/ora_02_u02_02
/dev/sdk       host34       iSCSI       20g    cDOT
nim             /vol/ora_02_data/ora_02_u02_01
/dev/sdj       host34       iSCSI       20g    cDOT
nim             /vol/ora_02_biny/ora_02_u01_01
/dev/sdi       host34       iSCSI       40g    cDOT
nim             /vol/ora_02_logs/ora_02_u03_02
/dev/sdh       host33       iSCSI       20g    cDOT
nim             /vol/ora_02_logs/ora_02_u03_01
/dev/sdg       host33       iSCSI       20g    cDOT
nim             /vol/ora_02_data/ora_02_u02_04
/dev/sdf       host33       iSCSI       20g    cDOT
nim             /vol/ora_02_data/ora_02_u02_03
/dev/sde       host33       iSCSI       20g    cDOT
nim             /vol/ora_02_data/ora_02_u02_02
/dev/sdd       host33       iSCSI       20g    cDOT
nim             /vol/ora_02_data/ora_02_u02_01
/dev/sdc       host33       iSCSI       20g    cDOT
nim             /vol/ora_02_biny/ora_02_u01_01
/dev/sdb       host33       iSCSI       40g    cDOT
```

6. Configure the `multipath.conf` file with following default and blacklist entries.

```
sudo vi /etc/multipath.conf
```

Add following entries:

```
defaults {
    find_multipaths yes
    user_friendly_names yes
}

blacklist {
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9]*"
    devnode "^hd[a-z]"
    devnode "^cciss.*"
}
```

7. Start the multipath service.

```
sudo systemctl start multipathd
```

Now multipath devices appear in the `/dev/mapper` directory.

```
[ec2-user@ip-172-30-15-58 ~]$ ls -l /dev/mapper
total 0
lrwxrwxrwx 1 root root          7 Mar 21 20:13
3600a09806c574235472455534e68512d -> ../dm-0
lrwxrwxrwx 1 root root          7 Mar 21 20:13
3600a09806c574235472455534e685141 -> ../dm-1
lrwxrwxrwx 1 root root          7 Mar 21 20:13
3600a09806c574235472455534e685142 -> ../dm-2
lrwxrwxrwx 1 root root          7 Mar 21 20:13
3600a09806c574235472455534e685143 -> ../dm-3
lrwxrwxrwx 1 root root          7 Mar 21 20:13
3600a09806c574235472455534e685144 -> ../dm-4
lrwxrwxrwx 1 root root          7 Mar 21 20:13
3600a09806c574235472455534e685145 -> ../dm-5
lrwxrwxrwx 1 root root          7 Mar 21 20:13
3600a09806c574235472455534e685146 -> ../dm-6
crw----- 1 root root 10, 236 Mar 21 18:19 control
```

8. Log into the FSx ONTAP cluster as the `fsxadmin` user via SSH to retrieve the serial-hex number for each LUN starting with `6c574xxx...`, the HEX number starts with `3600a0980`, which is the AWS vendor ID.

```
lun show -fields serial-hex
```

and return as follow:

```
FsxId02ad7bf3476b741df::> lun show -fields serial-hex
vserver path                               serial-hex
-----
svm_ora /vol/ora_02_biny/ora_02_biny_01 6c574235472455534e68512d
svm_ora /vol/ora_02_data/ora_02_data_01 6c574235472455534e685141
svm_ora /vol/ora_02_data/ora_02_data_02 6c574235472455534e685142
svm_ora /vol/ora_02_data/ora_02_data_03 6c574235472455534e685143
svm_ora /vol/ora_02_data/ora_02_data_04 6c574235472455534e685144
svm_ora /vol/ora_02_logs/ora_02_logs_01 6c574235472455534e685145
svm_ora /vol/ora_02_logs/ora_02_logs_02 6c574235472455534e685146
7 entries were displayed.
```

9. Update the `/dev/multipath.conf` file to add a user-friendly name for the multipath device.

```
sudo vi /etc/multipath.conf
```

with following entries:

```

multipaths {
    multipath {
        wwid          3600a09806c574235472455534e68512d
        alias         ora_02_biny_01
    }
    multipath {
        wwid          3600a09806c574235472455534e685141
        alias         ora_02_data_01
    }
    multipath {
        wwid          3600a09806c574235472455534e685142
        alias         ora_02_data_02
    }
    multipath {
        wwid          3600a09806c574235472455534e685143
        alias         ora_02_data_03
    }
    multipath {
        wwid          3600a09806c574235472455534e685144
        alias         ora_02_data_04
    }
    multipath {
        wwid          3600a09806c574235472455534e685145
        alias         ora_02_logs_01
    }
    multipath {
        wwid          3600a09806c574235472455534e685146
        alias         ora_02_logs_02
    }
}

```

10. Reboot the multipath service to verify that the devices under `/dev/mapper` have changed to LUN names versus serial-hex IDs.

```
sudo systemctl restart multipathd
```

Check `/dev/mapper` to return as following:

```
[ec2-user@ip-172-30-15-58 ~]$ ls -l /dev/mapper
total 0
crw----- 1 root root 10, 236 Mar 21 18:19 control
lrwxrwxrwx 1 root root      7 Mar 21 20:41 ora_02_biny_01 -> ../dm-
0
lrwxrwxrwx 1 root root      7 Mar 21 20:41 ora_02_data_01 -> ../dm-
1
lrwxrwxrwx 1 root root      7 Mar 21 20:41 ora_02_data_02 -> ../dm-
2
lrwxrwxrwx 1 root root      7 Mar 21 20:41 ora_02_data_03 -> ../dm-
3
lrwxrwxrwx 1 root root      7 Mar 21 20:41 ora_02_data_04 -> ../dm-
4
lrwxrwxrwx 1 root root      7 Mar 21 20:41 ora_02_logs_01 -> ../dm-
5
lrwxrwxrwx 1 root root      7 Mar 21 20:41 ora_02_logs_02 -> ../dm-
6
```

11. Partition the binary LUN with a single primary partition.

```
sudo fdisk /dev/mapper/ora_02_biny_01
```

12. Format the partitioned binary LUN with an XFS file system.

```
sudo mkfs.xfs /dev/mapper/ora_02_biny_01p1
```

13. Mount the binary LUN to /u01.

```
sudo mkdir /u01
```

```
sudo mount -t xfs /dev/mapper/ora_02_biny_01p1 /u01
```

14. Change /u01 mount point ownership to the oracle user and it's associated primary group.

```
sudo chown oracle:oinstall /u01
```

15. Find the UUI of the binary LUN.

```
sudo blkid /dev/mapper/ora_02_biny_01p1
```


16. Add a mount point to `/etc/fstab`.

```
sudo vi /etc/fstab
```

Add the following line.

```
UUID=d89fb1c9-4f89-4de4-b4d9-17754036d11d    /u01    xfs
defaults,nofail 0                2
```

17. As the root user, add the udev rule for Oracle devices.

```
vi /etc/udev/rules.d/99-oracle-asmdevices.rules
```

Include following entries:

```
ENV{DM_NAME}=="ora*", GROUP=="oinstall", OWNER=="oracle",
MODE=="660"
```

18. As the root user, reload the udev rules.

```
udevadm control --reload-rules
```

19. As the root user, trigger the udev rules.

```
udevadm trigger
```

20. As the root user, reload multipathd.

```
systemctl restart multipathd
```

21. Reboot the EC2 instance host.

Oracle grid infrastructure installation

1. Log into the DB VM as the admin user via SSH and enable password authentication by uncommenting `PasswordAuthentication yes` and then commenting out `PasswordAuthentication no`.

```
sudo vi /etc/ssh/sshd_config
```

2. Restart the sshd service.

```
sudo systemctl restart sshd
```

3. Reset the Oracle user password.

```
sudo passwd oracle
```

4. Log in as the Oracle Restart software owner user (oracle). Create an Oracle directory as follows:

```
mkdir -p /u01/app/oracle
```

```
mkdir -p /u01/app/oraInventory
```

5. Change the directory permission setting.

```
chmod -R 775 /u01/app
```

6. Create a grid home directory and change to it.

```
mkdir -p /u01/app/oracle/product/19.0.0/grid
```

```
cd /u01/app/oracle/product/19.0.0/grid
```

7. Unzip the grid installation files.

```
unzip -q /tmp/archive/LINUX.X64_193000_grid_home.zip
```

8. From grid home, delete the `OPatch` directory.

```
rm -rf OPatch
```

9. From grid home, unzip p6880880_190000_Linux-x86-64.zip.

```
unzip -q /tmp/archive/p6880880_190000_Linux-x86-64.zip
```

10. From grid home, revise cv/admin/cvu_config, uncomment and replace CV_ASSUME_DISTID=OEL5 with CV_ASSUME_DISTID=OL7.

```
vi cv/admin/cvu_config
```

11. Prepare a gridsetup.rsp file for silent installation and place the rsp file in the /tmp/archive directory. The rsp file should cover sections A, B, and G with the following information:

```
INVENTORY_LOCATION=/u01/app/oraInventory
oracle.install.option=HA_CONFIG
ORACLE_BASE=/u01/app/oracle
oracle.install.asm.OSDBA=asmdba
oracle.install.asm.OSOPER=asmoper
oracle.install.asm.OSASM=asmadmin
oracle.install.asm.SYSASMPassword="SetPWD"
oracle.install.asm.diskGroup.name=DATA
oracle.install.asm.diskGroup.redundancy=EXTERNAL
oracle.install.asm.diskGroup.AUSize=4
oracle.install.asm.diskGroup.disks=/dev/mapper/ora_02_data_01,/dev/mapper/ora_02_data_02,/dev/mapper/ora_02_data_03,/dev/mapper/ora_02_data_04
oracle.install.asm.diskGroup.diskDiscoveryString=/dev/mapper/*
oracle.install.asm.monitorPassword="SetPWD"
oracle.install.asm.configureAFD=true
```

12. Log into the EC2 instance as the root user and set ORACLE_HOME and ORACLE_BASE.

```
export ORACLE_HOME=/u01/app/oracle/product/19.0.0/
```

```
export ORACLE_BASE=/tmp
```

```
cd /u01/app/oracle/product/19.0.0/grid/bin
```

13. Initialize disk devices for use with the Oracle ASM filter driver.

```
./asmcmd afd_label DATA01 /dev/mapper/ora_02_data_01 --init
```

```
./asmcmd afd_label DATA02 /dev/mapper/ora_02_data_02 --init
```

```
./asmcmd afd_label DATA03 /dev/mapper/ora_02_data_03 --init
```

```
./asmcmd afd_label DATA04 /dev/mapper/ora_02_data_04 --init
```

```
./asmcmd afd_label LOGS01 /dev/mapper/ora_02_logs_01 --init
```

```
./asmcmd afd_label LOGS02 /dev/mapper/ora_02_logs_02 --init
```

14. Install cvuqdisk-1.0.10-1.rpm.

```
rpm -ivh /u01/app/oracle/product/19.0.0/grid/cv/rpm/cvuqdisk-1.0.10-1.rpm
```

15. Unset \$ORACLE_BASE.

```
unset ORACLE_BASE
```

16. Log into the EC2 instance as the Oracle user and extract the patch in the /tmp/archive folder.

```
unzip -q /tmp/archive/p34762026_190000_Linux-x86-64.zip -d /tmp/archive
```

17. From grid home /u01/app/oracle/product/19.0.0/grid and as the oracle user, launch gridSetup.sh for grid infrastructure installation.

```
./gridSetup.sh -applyRU /tmp/archive/34762026/ -silent -responseFile /tmp/archive/gridsetup.rsp
```

18. As root user, execute the following script(s):

```
/u01/app/oraInventory/orainstRoot.sh
```

```
/u01/app/oracle/product/19.0.0/grid/root.sh
```

19. As root user, reload the multipathd.

```
systemctl restart multipathd
```

20. As the Oracle user, execute the following command to complete the configuration:

```
/u01/app/oracle/product/19.0.0/grid/gridSetup.sh -executeConfigTools  
-responseFile /tmp/archive/gridsetup.rsp -silent
```

21. As the Oracle user, create the LOGS disk group.

```
bin/asmca -silent -sysAsmPassword 'yourPWD' -asmsnmpPassword  
'yourPWD' -createDiskGroup -diskGroupName LOGS -disk 'AFD:LOGS*'  
-redundancy EXTERNAL -au_size 4
```

22. As the Oracle user, validate grid services after installation configuration.

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

```
[oracle@ora_02 grid]$ bin/crsctl stat res -t
```

```
-----  
-----  
Name          Target  State          Server          State  
details  
-----  
-----  
Local Resources  
-----  
-----  
ora.DATA.dg  
          ONLINE  ONLINE          ora_02          STABLE  
ora.LISTENER.lsnr  
          ONLINE  INTERMEDIATE   ora_02          Not All  
Endpoints Re  
gistered, STABLE  
ora.LOGS.dg  
          ONLINE  ONLINE          ora_02          STABLE  
ora.asm  
          ONLINE  ONLINE          ora_02  
Started, STABLE  
ora.ons  
          OFFLINE OFFLINE          ora_02          STABLE  
-----  
-----  
Cluster Resources  
-----  
-----  
ora.cssd  
   1      ONLINE  ONLINE          ora_02          STABLE  
ora.diskmon  
   1      OFFLINE OFFLINE          STABLE  
ora.driver.afd  
   1      ONLINE  ONLINE          ora_02          STABLE  
ora.evmd  
   1      ONLINE  ONLINE          ora_02          STABLE  
-----  
-----
```

23. Valiate ASM filter driver status.

```

[oracle@ora_02 grid]$ export
ORACLE_HOME=/u01/app/oracle/product/19.0.0/grid
[oracle@ora_02 grid]$ export ORACLE_SID=+ASM
[oracle@ora_02 grid]$ export PATH=$PATH:$ORACLE_HOME/bin
[oracle@ora_02 grid]$ asmcmd
ASMCMDB> lsdg
State      Type      Rebal  Sector  Logical_Sector  Block      AU
Total_MB  Free_MB  Req_mir_free_MB  Usable_file_MB  Offline_disks
Voting_files  Name
MOUNTED  EXTERN  N      512     512     4096    4194304
81920    81780      0      81780
N  DATA/
MOUNTED  EXTERN  N      512     512     4096    4194304
40960    40852      0      40852
N  LOGS/
ASMCMDB> afd_state
ASMCMDB-9526: The AFD state is 'LOADED' and filtering is 'ENABLED' on
host 'ora_02'
ASMCMDB> exit
[oracle@ora_02 grid]$

```

24. Validate HA service status.

```

[oracle@ora_02 bin]$ ./crsctl check has
CRS-4638: Oracle High Availability Services is online

```

Oracle database installation

1. Log in as the Oracle user and unset `$ORACLE_HOME` and `$ORACLE_SID` if it is set.

```
unset ORACLE_HOME
```

```
unset ORACLE_SID
```

2. Create the Oracle DB home directory and change the directory to it.

```
mkdir /u01/app/oracle/product/19.0.0/cdb3
```

```
cd /u01/app/oracle/product/19.0.0/cdb3
```

3. Unzip the Oracle DB installation files.

```
unzip -q /tmp/archive/LINUX.X64_193000_db_home.zip
```

4. From the DB home, delete the OPatch directory.

```
rm -rf OPatch
```

5. From DB home, unzip `p6880880_190000_Linux-x86-64.zip`.

```
unzip -q /tmp/archive/p6880880_190000_Linux-x86-64.zip
```

6. From DB home, revise `cv/admin/cvu_config` and uncomment and replace `CV_ASSUME_DISTID=OEL5` with `CV_ASSUME_DISTID=OL7`.

```
vi cv/admin/cvu_config
```

7. From the `/tmp/archive` directory, unpack the DB 19.18 RU patch.

```
unzip -q /tmp/archive/p34765931_190000_Linux-x86-64.zip -d  
/tmp/archive
```

8. Prepare the DB silent install `rsp` file in `/tmp/archive/dbinstall.rsp` directory with the following values:


```
oracle.install.option=INSTALL_DB_SWONLY
UNIX_GROUP_NAME=oinstall
INVENTORY_LOCATION=/u01/app/oraInventory
ORACLE_HOME=/u01/app/oracle/product/19.0.0/cdb3
ORACLE_BASE=/u01/app/oracle
oracle.install.db.InstallEdition=EE
oracle.install.db.OSDBA_GROUP=dba
oracle.install.db.OSOPER_GROUP=oper
oracle.install.db.OSBACKUPDBA_GROUP=oper
oracle.install.db.OSDGDBA_GROUP=dba
oracle.install.db.OSKMDBA_GROUP=dba
oracle.install.db.OSRACDBA_GROUP=dba
oracle.install.db.rootconfig.executeRootScript=false
```

9. From cdb3 home /u01/app/oracle/product/19.0.0/cdb3, execute silent software-only DB installation.

```
./runInstaller -applyRU /tmp/archive/34765931/ -silent
-ignorePrereqFailure -responseFile /tmp/archive/dbinstall.rsp
```

10. As root user, run the `root.sh` script after software-only installation.

```
/u01/app/oracle/product/19.0.0/db1/root.sh
```

11. As oracle user, create the `dbca.rsp` file with the following entries:

```
gdbName=cdb3.demo.netapp.com
sid=cdb3
createAsContainerDatabase=true
numberOfPDBs=3
pdbName=cdb3_pdb
useLocalUndoForPDBs=true
pdbAdminPassword="yourPWD"
templateName=General_Purpose.dbc
sysPassword="yourPWD"
systemPassword="yourPWD"
dbsnmpPassword="yourPWD"
datafileDestination=+DATA
recoveryAreaDestination=+LOGS
storageType=ASM
diskGroupName=DATA
characterSet=AL32UTF8
nationalCharacterSet=AL16UTF16
listeners=LISTENER
databaseType=MULTIPURPOSE
automaticMemoryManagement=false
totalMemory=8192
```

12. As oracle user, launch DB creation with dbca.

```
bin/dbca -silent -createDatabase -responseFile /tmp/archive/dbca.rsp
```

output:

```

Prepare for db operation
7% complete
Registering database with Oracle Restart
11% complete
Copying database files
33% complete
Creating and starting Oracle instance
35% complete
38% complete
42% complete
45% complete
48% complete
Completing Database Creation
53% complete
55% complete
56% complete
Creating Pluggable Databases
60% complete
64% complete
69% complete
78% complete
Executing Post Configuration Actions
100% complete
Database creation complete. For details check the logfiles at:
  /u01/app/oracle/cfgtoollogs/dbca/cdb3.
Database Information:
Global Database Name:cdb3.vmc.netapp.com
System Identifier(SID):cdb3
Look at the log file "/u01/app/oracle/cfgtoollogs/dbca/cdb3/cdb3.log"
for further details.

```

1. Repeat the same procedures from step 2 to create a container database cdb4 in a separate ORACLE_HOME /u01/app/oracle/product/19.0.0/cdb4 with a single PDB.
2. As Oracle user, validate Oracle Restart HA services after DB creation that all databases (cdb3, cdb4) are registered with HA services.

```
/u01/app/oracle/product/19.0.0/grid/crsctl stat res -t
```

output:

```

[oracle@ora_02 bin]$ ./crsctl stat res -t
-----
-----
Name                Target  State          Server                State

```

details

Local Resources

ora.DATA.dg
 ONLINE ONLINE ora_02 STABLE
ora.LISTENER.lsnr
 ONLINE INTERMEDIATE ora_02 Not All

Endpoints Re

gistered, STABLE

ora.LOGS.dg
 ONLINE ONLINE ora_02 STABLE
ora.asm
 ONLINE ONLINE ora_02

Started, STABLE

ora.ons
 OFFLINE OFFLINE ora_02 STABLE

Cluster Resources

ora.cdb3.db
 1 ONLINE ONLINE ora_02

Open,HOME=/u01/app/o

racle/product/19.0.0

/cdb3, STABLE

ora.cdb4.db
 1 ONLINE ONLINE ora_02

Open,HOME=/u01/app/o

racle/product/19.0.0

/cdb4, STABLE

ora.cssd
 1 ONLINE ONLINE ora_02 STABLE

ora.diskmon
 1 OFFLINE OFFLINE STABLE

ora.driver.afd
 1 ONLINE ONLINE ora_02 STABLE

ora.evmd

```
1          ONLINE  ONLINE          ora_02          STABLE
-----
-----
```

3. Set the Oracle user .bash_profile.

```
vi ~/.bash_profile
```

Add following entries:

```
export ORACLE_HOME=/u01/app/oracle/product/19.0.0/db3
export ORACLE_SID=db3
export PATH=$PATH:$ORACLE_HOME/bin
alias asm='export
ORACLE_HOME=/u01/app/oracle/product/19.0.0/grid;export
ORACLE_SID=+ASM;export PATH=$PATH:$ORACLE_HOME/bin'
alias cdb3='export
ORACLE_HOME=/u01/app/oracle/product/19.0.0/cdb3;export
ORACLE_SID=cdb3;export PATH=$PATH:$ORACLE_HOME/bin'
alias cdb4='export
ORACLE_HOME=/u01/app/oracle/product/19.0.0/cdb4;export
ORACLE_SID=cdb4;export PATH=$PATH:$ORACLE_HOME/bin'
```

4. Validate the CDB/PDB created for cdb3.

```
cdb3
```

```
[oracle@ora_02 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Mon Oct 9 08:19:20 2023
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 from v$database;
```

```
NAME          OPEN_MODE
-----
```

```
CDB3          READ WRITE
```

```
SQL> show pdbs
```

```
CON_ID CON_NAME                                OPEN MODE  RESTRICTED
-----
```

2	PDB\$SEED	READ ONLY	NO
3	CDB3_PDB1	READ WRITE	NO
4	CDB3_PDB2	READ WRITE	NO
5	CDB3_PDB3	READ WRITE	NO

```
SQL>
```

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

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

```
+DATA/CDB3/DATAFILE/system.257.1149420273
+DATA/CDB3/DATAFILE/sysaux.258.1149420317
+DATA/CDB3/DATAFILE/undotbs1.259.1149420343
+DATA/CDB3/86B637B62FE07A65E053F706E80A27CA/DATAFILE/system.266.1149
421085
+DATA/CDB3/86B637B62FE07A65E053F706E80A27CA/DATAFILE/sysaux.267.1149
421085
+DATA/CDB3/DATAFILE/users.260.1149420343
+DATA/CDB3/86B637B62FE07A65E053F706E80A27CA/DATAFILE/undotbs1.268.11
49421085
+DATA/CDB3/06FB206DF15ADEE8E065025056B66295/DATAFILE/system.272.1149
422017
+DATA/CDB3/06FB206DF15ADEE8E065025056B66295/DATAFILE/sysaux.273.1149
422017
+DATA/CDB3/06FB206DF15ADEE8E065025056B66295/DATAFILE/undotbs1.271.11
49422017
+DATA/CDB3/06FB206DF15ADEE8E065025056B66295/DATAFILE/users.275.11494
22033
```

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

```
+DATA/CDB3/06FB21766256DF9AE065025056B66295/DATAFILE/system.277.1149
422033
+DATA/CDB3/06FB21766256DF9AE065025056B66295/DATAFILE/sysaux.278.1149
422033
+DATA/CDB3/06FB21766256DF9AE065025056B66295/DATAFILE/undotbs1.276.11
```

```

49422033
+DATA/CDB3/06FB21766256DF9AE065025056B66295/DATAFILE/users.280.11494
22049
+DATA/CDB3/06FB22629AC1DFD7E065025056B66295/DATAFILE/system.282.1149
422049
+DATA/CDB3/06FB22629AC1DFD7E065025056B66295/DATAFILE/sysaux.283.1149
422049
+DATA/CDB3/06FB22629AC1DFD7E065025056B66295/DATAFILE/undotbs1.281.11
49422049
+DATA/CDB3/06FB22629AC1DFD7E065025056B66295/DATAFILE/users.285.11494
22063

```

19 rows selected.

SQL>

5. Validate the CDB/PDB created for cdb4.

```

cdb4

```

```

[oracle@ora_02 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Mon Oct 9 08:20:26 2023
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 from v$database;

NAME          OPEN_MODE
-----
CDB4          READ WRITE

SQL> show pdbs

          CON_ID CON_NAME                                OPEN MODE  RESTRICTED
          -----
          2 PDB$SEED                                READ ONLY  NO

```

3 CDB4_PDB

READ WRITE NO

SQL>

SQL> select name from v\$datafile;

NAME

```
-----  
-----  
+DATA/CDB4/DATAFILE/system.286.1149424943  
+DATA/CDB4/DATAFILE/sysaux.287.1149424989  
+DATA/CDB4/DATAFILE/undotbs1.288.1149425015  
+DATA/CDB4/86B637B62FE07A65E053F706E80A27CA/DATAFILE/system.295.1149  
425765  
+DATA/CDB4/86B637B62FE07A65E053F706E80A27CA/DATAFILE/sysaux.296.1149  
425765  
+DATA/CDB4/DATAFILE/users.289.1149425015  
+DATA/CDB4/86B637B62FE07A65E053F706E80A27CA/DATAFILE/undotbs1.297.11  
49425765  
+DATA/CDB4/06FC3070D5E12C23E065025056B66295/DATAFILE/system.301.1149  
426581  
+DATA/CDB4/06FC3070D5E12C23E065025056B66295/DATAFILE/sysaux.302.1149  
426581  
+DATA/CDB4/06FC3070D5E12C23E065025056B66295/DATAFILE/undotbs1.300.11  
49426581  
+DATA/CDB4/06FC3070D5E12C23E065025056B66295/DATAFILE/users.304.11494  
26597
```

11 rows selected.

6. Login to each cdb as sysdba with sqlplus and set the DB recovery destination size to the +LOGS disk group size for both cdb's.

```
alter system set db_recovery_file_dest_size = 40G scope=both;
```

7. Login to each cdb as sysdba with sqlplus and enable archive log mode with following command sets in sequence.

```
sqlplus /as sysdba
```

```
shutdown immediate;
```

```
startup mount;
```



```
alter database archivelog;
```

```
alter database open;
```

This completes Oracle 19c version 19.18 Restart deployment on an Amazon FSx for ONTAP storage and a VMC DB VM. If desired, NetApp recommends relocating the Oracle control file and online log files to the +LOGS disk group.

Oracle backup, restore, and clone with SnapCenter

SnapCenter Setup

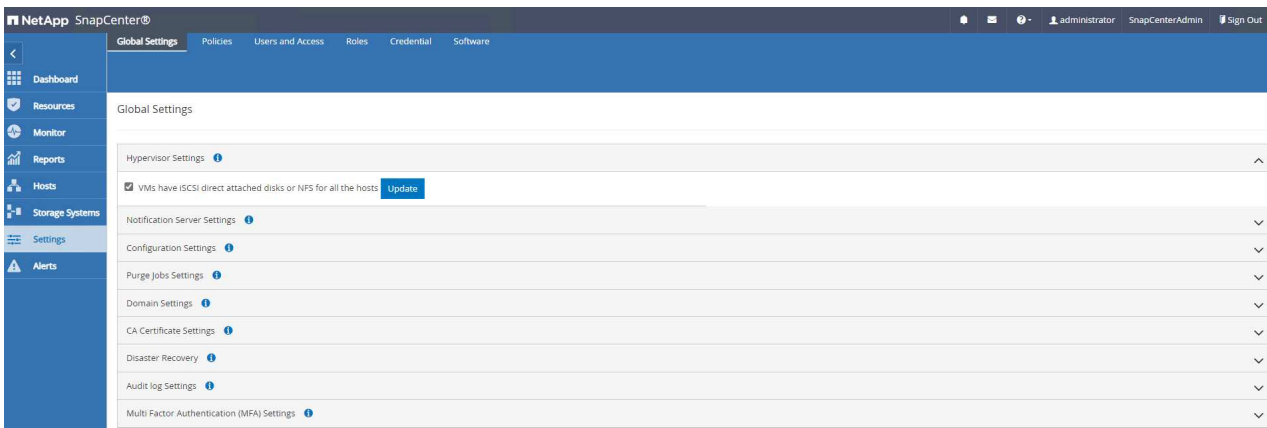
SnapCenter relies on a host-side plug-in on database VM to perform application-aware data protection management activities. For detailed information on NetApp SnapCenter plugin for Oracle, refer to this documentation [What can you do with the Plug-in for Oracle Database](#). The following provides high level steps to setup SnapCenter for Oracle database backup, recovery, and clone.

1. Download the latest version of SnapCenter software from NetApp support site: [NetApp Support Downloads](#).
2. As administrator, install latest java JDK from [Get Java for desktop applications](#) on SnapCenter server Windows host.




If Windows server is deployed in a domain environment, add a domain user to SnapCenter server local administrators group and run SnapCenter installation with the domain user.


3. Login to SnapCenter UI via HTTPS port 8846 as installation user to configure SnapCenter for Oracle.
4. Update Hypervisor Settings in global settings.



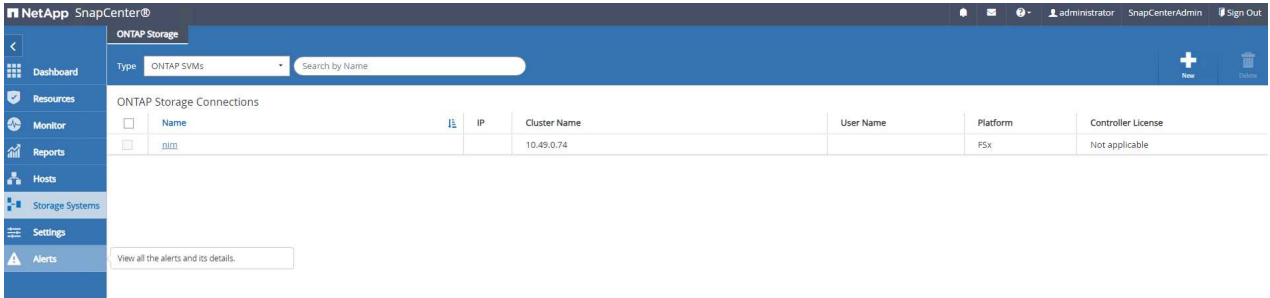
5. Create Oracle database backup policies. Ideally, create a separate archive log backup policy to allow more frequent backup interval to minimize data loss in the event of a failure.




6. Add database server Credential for SnapCenter access to DB VM. The credential should have sudo privilege on a Linux VM or administrator privilege on a Windows VM.



7. Add FSx ONTAP storage cluster to `Storage Systems` with cluster management IP and authenticated via fsxadmin user ID.



8. Add Oracle database VM in VMC to `Hosts` with server credential created in previous step 6.

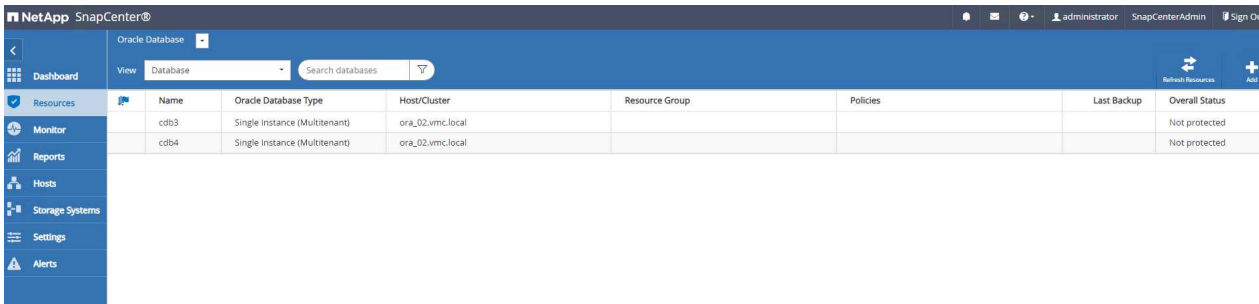


Ensure that the SnapCenter server name can be resolved to the IP address from the DB VM and DB VM name can be resolved to the IP address from the SnapCenter server.

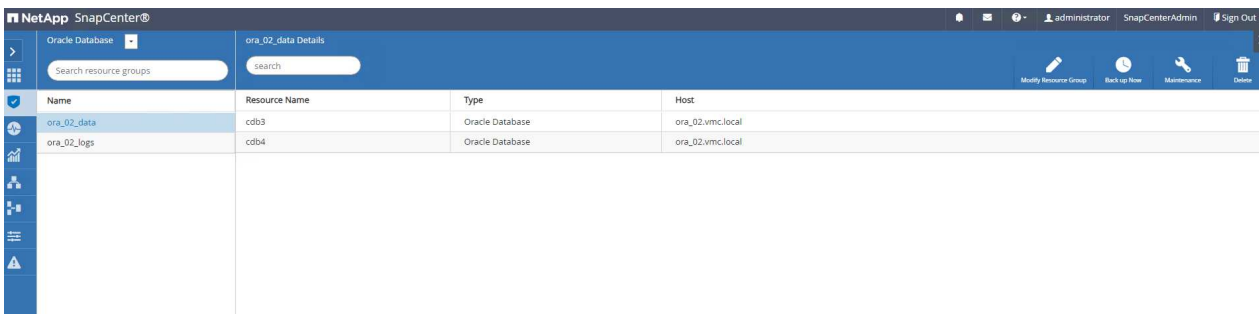
Database backup

SnapCenter leverages FSx ONTAP volume snapshot for much quicker database backup, restore, or clone compared with traditional RMAN based methodology. The snapshots are application-consistent as the database is put in Oracle backup mode before a snapshot.

1. From the Resources tab, any databases on the VM are auto-discovered after the VM is added to SnapCenter. Initially, the database status shows as Not protected.




2. Create a resources group to backup the database in a logical grouping such as by DB VM etc. In this example, we created an ora_02_data group to do a full online database backup for all databases on VM ora_02. Resources group ora_02_log performs the backup of archived logs only on the VM. Creating a resources group also defines a schedule to execute the backup.




3. Resources group backup can also be triggered manually by clicking on Back up Now and executing the backup with the policy defined in the resources group.

Add schedules for policy Oracle Online Full Backup ✕

Hourly

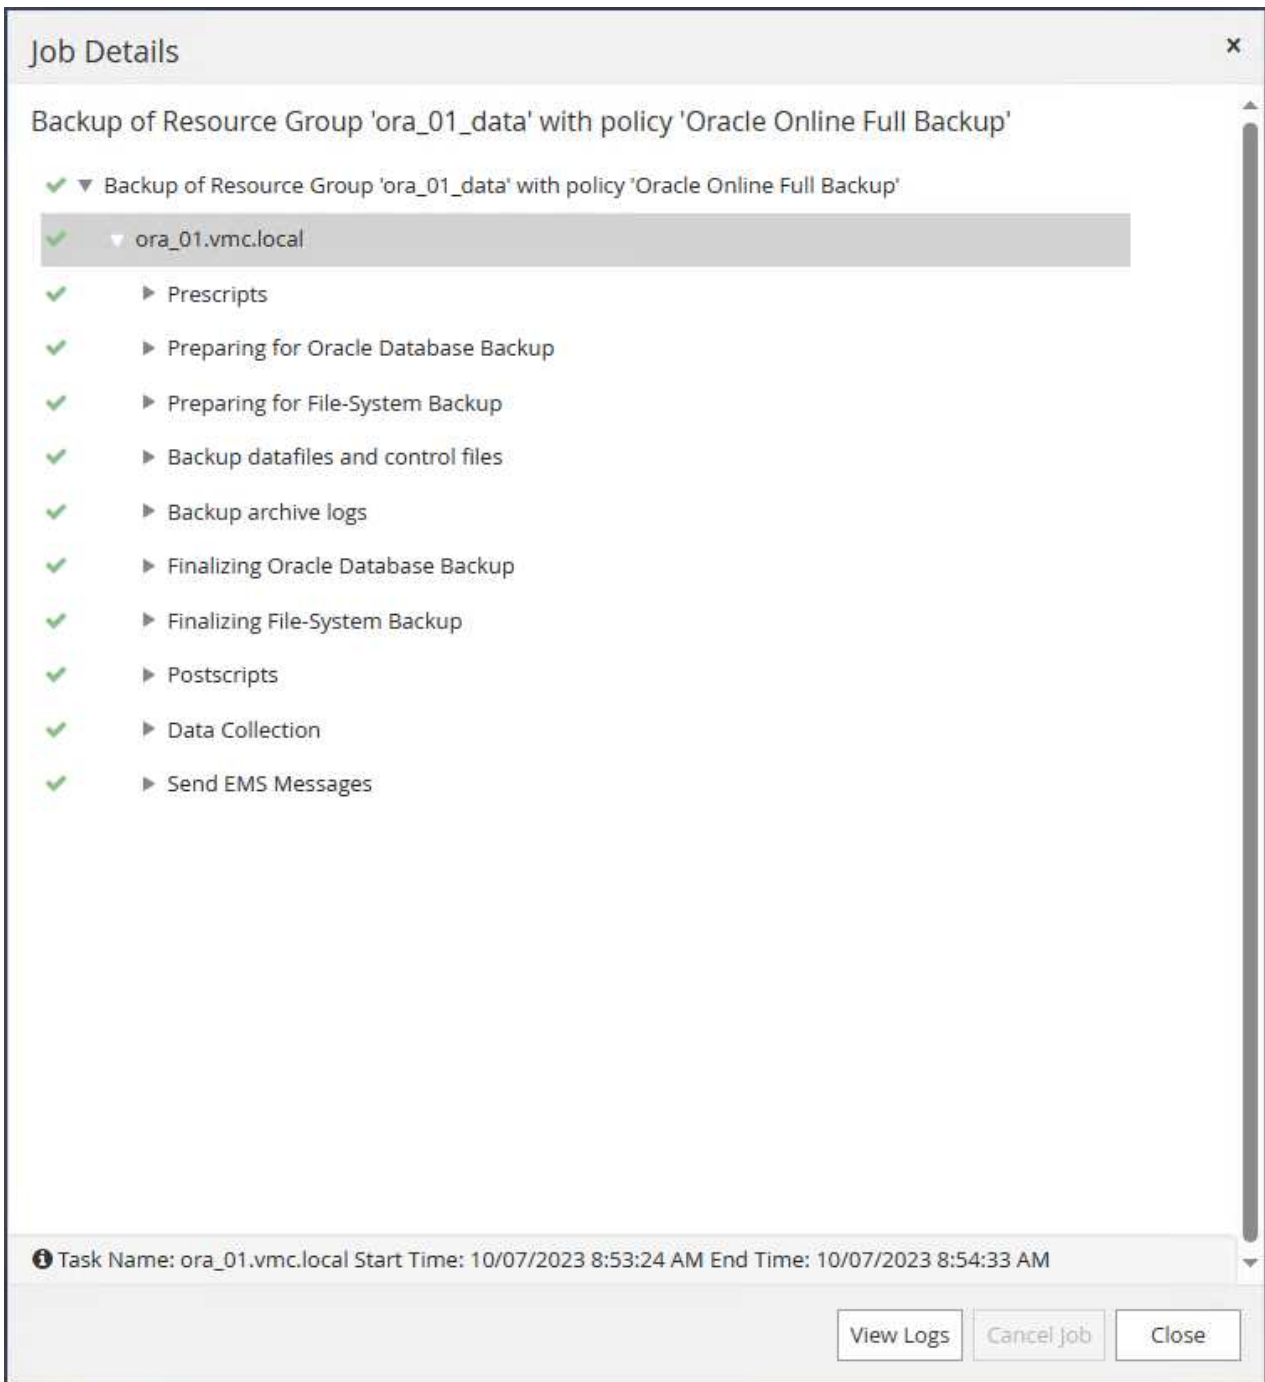
Start date 

Expires on 

Repeat every hours mins

i The schedules are triggered in the SnapCenter Server time zone. ✕

4. The backup job can be monitored at the `Monitor` tab by clicking on the running job.



5. After a successful backup, the database status shows the job status and the most recent backup time.

Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Overall Status
cdb1	Single Instance (Multitenant)	ora_01.vmc.local	ora_01_data ora_01_logs	Oracle Archive Logs Backup Oracle Online Full Backup	10/07/2023 12:00:25 PM	Backup succeeded
cdb2	Single Instance (Multitenant)	ora_01.vmc.local	ora_01_data ora_01_logs	Oracle Archive Logs Backup Oracle Online Full Backup	10/07/2023 12:00:25 PM	Backup succeeded
cdb3	Single Instance (Multitenant)	ora_02.vmc.local	ora_02_data ora_02_logs	Oracle Archive Logs Backup Oracle Online Full Backup	10/07/2023 8:05:25 AM	Backup succeeded
cdb4	Single Instance (Multitenant)	ora_02.vmc.local	ora_02_data ora_02_logs	Oracle Archive Logs Backup Oracle Online Full Backup	10/07/2023 8:05:25 AM	Backup succeeded

6. Click on database to review the backup sets for each database.

NetApp SnapCenter®

Oracle Database | cdb3 Topology

Search databases

Administrator | SnapCenterAdmin | Sign Out

Backup to Object Store | Database Settings | Protect | Refresh

Manage Copies

22 Backups
0 Clones
Local copies

Summary Card

22 Backups
8 Data Backups
14 Log Backups
0 Clones

Primary Backup(s)

Backup Name	Count	Type	IF	End Date	Verified	Mounted	RMAN Cataloged	SCN
ora_02_10-07-2023_08.05.02.4105_1	1	Log		10/07/2023 8:05:26 AM	Not Applicable	False	Not Cataloged	2928738
ora_02_10-07-2023_07.50.02.4250_1	1	Log		10/07/2023 7:50:27 AM	Not Applicable	False	Not Cataloged	2927731
ora_02_10-07-2023_07.45.02.4192_1	1	Log		10/07/2023 7:45:49 AM	Not Applicable	False	Not Cataloged	2927497
ora_02_10-07-2023_07.45.02.4192_0	1	Data		10/07/2023 7:45:31 AM	Unverified	False	Not Cataloged	2927446
ora_02_10-07-2023_07.35.02.3846_1	1	Log		10/07/2023 7:35:25 AM	Not Applicable	False	Not Cataloged	2926747
ora_02_10-07-2023_07.20.02.3803_1	1	Log		10/07/2023 7:20:25 AM	Not Applicable	False	Not Cataloged	2925995
ora_02_10-07-2023_07.05.02.3948_1	1	Log		10/07/2023 7:05:26 AM	Not Applicable	False	Not Cataloged	2924987
ora_02_10-07-2023_06.50.02.3786_1	1	Log		10/07/2023 6:50:26 AM	Not Applicable	False	Not Cataloged	2923925

Database recovery

SnapCenter provides a number of restore and recovery options for Oracle databases from snapshot backup. In this example, we demonstrate a point in time restoration to recover a dropped table by mistake. On VM ora_02, two databases cdb3, cdb4 share the same +DATA and +LOGS disk groups. Database restoration for one database does not impact the availability of the other database.

1. First, create a test table and insert a row into table to validate a point in time recovery.

```
[oracle@ora_02 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Fri Oct 6 14:15:21 2023
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 from v$database;

NAME          OPEN_MODE
-----
CDB3          READ WRITE

SQL> show pdbs

          CON_ID CON_NAME                                OPEN MODE  RESTRICTED
-----
          2 PDB$SEED                                READ ONLY  NO
          3 CDB3_PDB1                            READ WRITE NO
          4 CDB3_PDB2                            READ WRITE NO
          5 CDB3_PDB3                            READ WRITE NO

SQL>

SQL> alter session set container=cdb3_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 recovery on
guest mounted fsx storage to VMC guest vm ora_02');
```

```
1 row created.
```

```
SQL> commit;
```

```
Commit complete.
```

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

```
          ID
-----
DT
-----
EVENT
-----
          1
06-OCT-23 03.18.24.000000 PM
test oracle recovery on guest mounted fsx storage to VMC guest vm
ora_02
```

```
SQL> select current_timestamp from dual;
```

```
CURRENT_TIMESTAMP
-----
06-OCT-23 03.18.53.996678 PM -07:00
```

2. We run a manual snapshot backup from SnapCenter. Then drop the table.

```

SQL> drop table test;

Table dropped.

SQL> commit;

Commit complete.

SQL> select current_timestamp from dual;

CURRENT_TIMESTAMP
-----
06-OCT-23 03.26.30.169456 PM -07:00

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

```

- From backup set created from last step, take a note of the SCN number of log backup. Click on Restore to launch restore-recover workflow.

The screenshot shows the NetApp SnapCenter interface for an Oracle Database. The main area displays a table of Primary Backup(s) with the following data:

Backup Name	Count	Type	End Date	Verified	Mounted	RMAN Cataloged	SCN
ora_02_10-06-2023_14.22.59.0383_1	1	Log	10/06/2023 2:23:43 PM	Not Applicable	False	Not Cataloged	2795205
ora_02_10-06-2023_14.22.59.0383_0	1	Data	10/06/2023 2:23:27 PM	Unverified	False	Not Cataloged	2795113
ora_02_10-06-2023_14.20.01.8472_1	1	Log	10/06/2023 2:20:24 PM	Not Applicable	False	Not Cataloged	2794928
ora_02_10-06-2023_14.05.01.8346_1	1	Log	10/06/2023 2:05:24 PM	Not Applicable	False	Not Cataloged	2793950
ora_02_10-06-2023_13.52.09.1111_1	1	Log	10/06/2023 1:52:59 PM	Not Applicable	False	Not Cataloged	2792888
ora_02_10-06-2023_13.52.09.1111_0	1	Data	10/06/2023 1:52:43 PM	Unverified	False	Not Cataloged	2792838

A Summary Card on the right side of the interface shows the following statistics:

- 6 Backups
- 2 Data Backups
- 4 Log Backups
- 0 Clones

- Choose restore scope.

Restore cdb3

- 1 Restore Scope
- 2 Recovery Scope
- 3 PreOps
- 4 PostOps
- 5 Notification
- 6 Summary

Restore Scope ?

All Datafiles

Pluggable databases (PDBs)

Pluggable database (PDB) tablespaces

Control files

Database State

Change database state if needed for restore and recovery

Restore Mode ?

Force in place restore

If this check box is not selected and if any of the in place restore criteria is not met, restore will be performed using the connect and copy method. The connect and copy restore method might take time based on the files being restored.

Previous Next

5. Choose recovery scope up to the log SCN from last full database backup.

Restore cdb3

- 1 Restore Scope
- 2 Recovery Scope**
- 3 PreOps
- 4 PostOps
- 5 Notification
- 6 Summary

Choose Recovery Scope

All Logs ?

Until SCN (System Change Number)

SCN ?

Date and Time

No recovery

Specify external archive log files locations + - ?

i After the operation is complete, it is recommended to create a full backup of the Oracle database. x

Previous Next

6. Specify any optional pre-scripts to run.

Restore cdb3 x

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Specify optional scripts to run before performing a restore job ⓘ

Prescript full path

Arguments

Script timeout

7. Specify any optional after-script to run.

Restore cdb3 x

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Specify optional scripts to run after performing a restore job ⓘ

Postscript full path

Arguments

Open the database or container database in READ-WRITE mode after recovery

8. Send a job report if desired.

Restore cdb3 ×

- 1 Restore Scope
- 2 Recovery Scope
- 3 PreOps
- 4 PostOps
- 5 Notification**
- 6 Summary

Provide email settings ⓘ

Email preference:

From:

To:

Subject:

Attach job report

9. Review the summary and click on `Finish` to launch the restoration and recovery.

Restore cdb3
✕

- 1 Restore Scope
- 2 Recovery Scope
- 3 PreOps
- 4 PostOps
- 5 Notification
- 6 Summary

Summary

Backup name	ora_02_10-06-2023_14.22.59.0383_0
Backup date	10/06/2023 2:23:27 PM
Restore scope	All DataFiles
Recovery scope	Until SCN 2795205
Auxiliary destination	
Options	Change database state if necessary , Open the database or container database in READ-WRITE mode after recovery
Prescript full path	None
Prescript arguments	
Postscript full path	None
Postscript arguments	
Send email	No

Previous
Finish

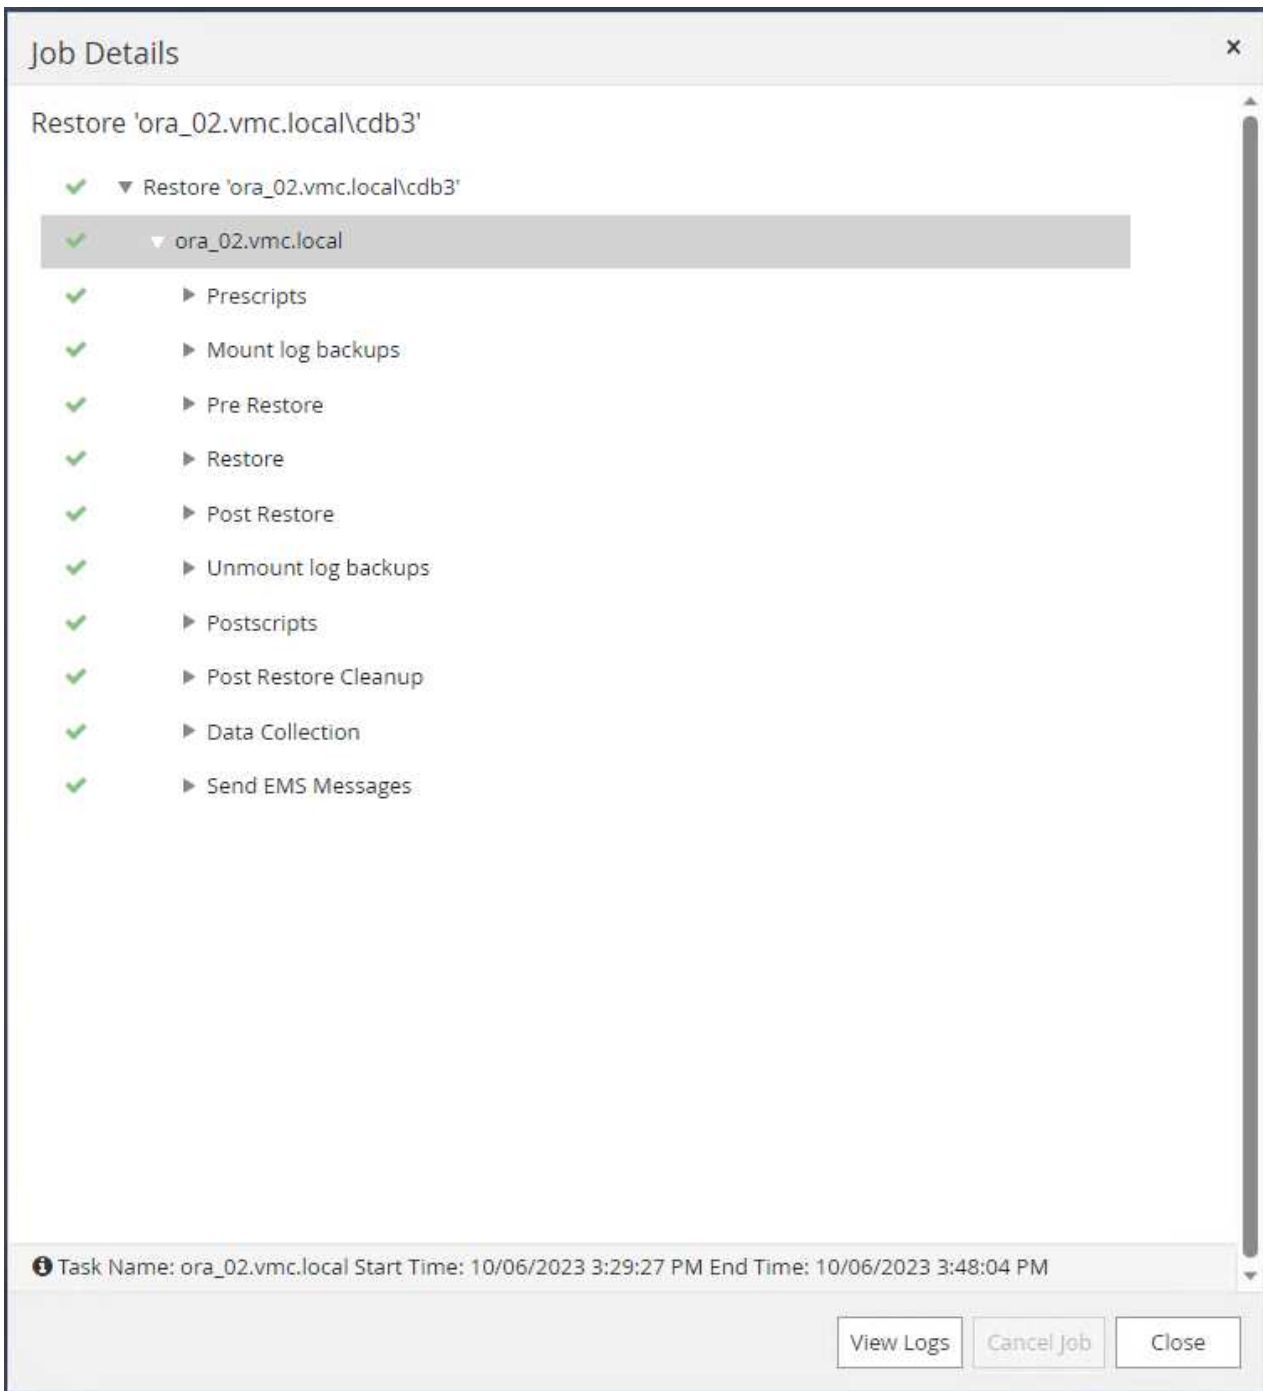
10. From Oracle Restart grid control, we observe that while cdb3 is under restoration and recovery cdb4 is online and available.


```

[oracle@ora_02 bin]$ ./crsctl stat res -t
-----
Name                Target  State        Server        State details
-----
Local Resources
-----
ora.DATA.dg         ONLINE  ONLINE       ora_02        STABLE
ora.LISTENER.lsnr   ONLINE  INTERMEDIATE ora_02        Not All Endpoints Re
                    registered, STABLE
ora.LOGS.dg         ONLINE  ONLINE       ora_02        STABLE
ora.LOGS_CDB3_22.dg ONLINE  ONLINE       ora_02        STABLE
ora.asm              ONLINE  ONLINE       ora_02        Started, STABLE
ora.ons              OFFLINE OFFLINE       ora_02        STABLE
-----
Cluster Resources
-----
ora.cdb3.db         1       ONLINE  INTERMEDIATE ora_02        Dismounted, Mount Ini
                    tiated, HOME=/u01/app
                    /oracle/product/19.0
                    .0/cdb3, STABLE
ora.cdb4.db         1       ONLINE  ONLINE       ora_02        Open, HOME=/u01/app/o
                    racle/product/19.0.0
                    /cdb4, STABLE
ora.cssd            1       ONLINE  ONLINE       ora_02        STABLE
ora.diskmon         1       OFFLINE OFFLINE       STABLE
ora.driver.afd      1       ONLINE  ONLINE       ora_02        STABLE
ora.evmd            1       ONLINE  ONLINE       ora_02        STABLE
-----
[oracle@ora_02 bin]$ █

```

11. From Monitor tab, open the job to review the details.



12. From DB VM ora_02, validate the dropped table is recovered after a successful recovery.

```
[oracle@ora_02 bin]$ sqlplus / as sysdba
```

```
SQL*Plus: Release 19.0.0.0.0 - Production on Fri Oct 6 17:01:28 2023  
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 from v\$database;

NAME	OPEN_MODE
CDB3	READ WRITE

SQL> show pdbs

CON_ID	CON_NAME	OPEN MODE	RESTRICTED
2	PDB\$SEED	READ ONLY	NO
3	CDB3_PDB1	READ WRITE	NO
4	CDB3_PDB2	READ WRITE	NO
5	CDB3_PDB3	READ WRITE	NO

SQL> alter session set container=CDB3_PDB1;

Session altered.

SQL> select * from test;

ID
DT
EVENT
1
06-OCT-23 03.18.24.000000 PM
test oracle recovery on guest mounted fsx storage to VMC guest vm
ora_02

SQL> select current_timestamp from dual;

CURRENT_TIMESTAMP
06-OCT-23 05.02.20.382702 PM -07:00

SQL>

Database clone

In this example, the same backup sets is used to clone a database on the same VM in a different ORACLE_HOME. The procedures are equally applicable to clone a database from the backup to separate VM in VMC if needed.

1. Open the database cdb3 backup list. From a data backup of choice, click on Clone button to launch database clone workflow.

The screenshot displays the NetApp SnapCenter interface for an Oracle Database. The main area shows the 'Manage Copies' section for 'cdb3 Topology'. A 'Summary Card' indicates 19 Backups, 6 Data Backups, 13 Log Backups, and 0 Clones. Below this is a table of 'Primary Backup(s)' with columns for Backup Name, Count, Type, IF, End Date, Verified, Mounted, RMAN Cataloged, and SCN. The table lists several log backups and one data backup.

Backup Name	Count	Type	IF	End Date	Verified	Mounted	RMAN Cataloged	SCN
ora_02_10-06-2023_17.20.01.9983_1	1	Log		10/06/2023 5:20:23 PM	Not Applicable	False	Not Cataloged	2814539
ora_02_10-06-2023_17.05.01.9656_1	1	Log		10/06/2023 5:05:24 PM	Not Applicable	False	Not Cataloged	2813819
ora_02_10-06-2023_16.50.01.9670_1	1	Log		10/06/2023 4:50:25 PM	Not Applicable	False	Not Cataloged	2812382
ora_02_10-06-2023_16.45.02.2685_1	1	Log		10/06/2023 4:45:45 PM	Not Applicable	False	Not Cataloged	2812040
ora_02_10-06-2023_16.45.02.2685_0	1	Data		10/06/2023 4:45:30 PM	Unverified	False	Not Cataloged	2811991
ora_02_10-06-2023_16.35.01.9959_1	1	Log		10/06/2023 4:35:22 PM	Not Applicable	False	Not Cataloged	2811534

2. Name the clone database SID.

Clone from cdb3 x

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps
- 6 Notification
- 7 Summary

Complete Database Clone

Clone SID:

Exclude PDBs:

PDB Clone

3. Select a VM in VMC as the target database host. Identical Oracle version should have been installed and configured on the host.

×
Clone from cdb3

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Select the host to create a clone

Clone host

Datafile locations ⓘ

Reset

Control files ⓘ

✕

✕ Reset

Redo logs ⓘ

Group	Size	Unit	Number of files
▶ RedoGroup 1	<input type="text" value="200"/> ✕	MB	2
▶ RedoGroup 2	<input type="text" value="200"/> ✕	MB	2
▶ RedoGroup 3	<input type="text" value="200"/> ✕	MB	2

Previous
Next

4. Select the proper ORACLE_HOME, user and group on the target host. Keep credential at default.

Clone from cdb3

1 Name

2 Locations

3 **Credentials**

4 PreOps

5 PostOps

6 Notification

7 Summary

Database Credentials for the clone

Credential name for sys user: None + ⓘ

ASM instance Credential name: None + ⓘ

Database port: 1521

ASM Port: 1521

Oracle Home Settings ⓘ

Oracle Home: /u01/app/oracle/product/19.0.0/cdb4

Oracle OS User: oracle

Oracle OS Group: oinstall

Previous Next

5. Change clone database parameters to meet configuration or resources requirements for the clone database.

x
Clone from cdb3

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps
- 6 Notification
- 7 Summary

Specify scripts to run before clone operation ?

Prescript full path

Arguments

Script timeout secs

Database Parameter settings

processes	320	x	▲
remote_login_passwordfile	EXCLUSIVE	x	+
sga_target	2048M	x	▼
undo_tablespace	UNDOTBS1	x	▼

6. Choose recovery scope. `Until Cancel` recovers the clone up to last available log file in the backup set.

Clone from cdb3

1 Name
2 Locations
3 Credentials
4 PreOps
5 PostOps
6 Notification
7 Summary

Recover Database

Until Cancel ⓘ
 Date and Time ⓘ
Date-time format: MM/DD/YYYY hh:mm:ss
 Until SCN (System Change Number) ⓘ

Specify external archive log locations ⓘ

Create new DBID ⓘ
 Create tempfile for temporary tablespace ⓘ
 Enter SQL queries to apply when clone is created
 Enter scripts to run after clone operation ⓘ

Previous Next

7. Review the summary and launch the clone job.

x
Clone from cdb3

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps
- 6 Notification
- 7 Summary

Summary

Clone from backup	ora_02_10-06-2023_16.45.02.2685_0
Clone SID	cdb3tst
Clone server	ora_01.vmc.local
Exclude PDBs	none
Oracle home	/u01/app/oracle/product/19.0.0/cdb2
Oracle OS user	oracle
Oracle OS group	oinstall
Datafile mountpaths	+SC_2090922_cdb3tst
Control files	+SC_2090922_cdb3tst/cdb3tst/control/control01.ctl +SC_2090922_cdb3tst/cdb3tst/control/control02.ctl
Redo groups	RedoGroup =1 TotalSize =200 Path =+SC_2090922_cdb3tst/cdb3tst/redo01_01.log RedoGroup =1 TotalSize =200 Path =+SC_2090922_cdb3tst/cdb3tst/redo01_02.log RedoGroup =2 TotalSize =200 Path =+SC_2090922_cdb3tst/cdb3tst/redo02_01.log RedoGroup =2 TotalSize =200 Path =+SC_2090922_cdb3tst/cdb3tst/redo02_02.log RedoGroup =3 TotalSize =200 Path =+SC_2090922_cdb3tst/cdb3tst/redo03_01.log RedoGroup =3 TotalSize =200 Path =+SC_2090922_cdb3tst/cdb3tst/redo03_02.log
Recovery scope	Until Cancel
Prescript full path	none
Prescript arguments	
Postscript full path	none
Postscript arguments	
Send email	No

Previous
Finish

8. Monitor the clone job execution from Monitor tab.

Job Details

Clone from backup 'ora_02_10-06-2023_16.45.02.2685_0'

- ✓ ▾ Clone from backup 'ora_02_10-06-2023_16.45.02.2685_0'
- ✓ ▾ ora_02.vmc.local
- ✓ ▶ Prescripts
- ✓ ▶ Query Host Information
- ✓ ▶ Prepare for Cloning
- ✓ ▶ Cloning Resources
- ✓ ▶ FileSystem Clone
- ✓ ▶ Application Clone
- ✓ ▶ Postscripts
- ✓ ▶ Register Clone
- ✓ ▶ Unmount Clone
- ✓ ▶ Data Collection
- ✓ ▶ Send EMS Messages

Task Name: ora_02.vmc.local Start Time: 10/06/2023 5:48:15 PM End Time: 10/06/2023 6:05:41 PM

View Logs Cancel Job Close

9. Cloned database is immediately registered in SnapCenter.

NetApp SnapCenter®

Oracle Database

View Database Search databases

Resources	Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Overall Status
	cdb1	Single Instance (Multitenant)	ora_01.vmc.local				Not protected
	cdb2	Single Instance (Multitenant)	ora_01.vmc.local				Not protected
	cdb3	Single Instance (Multitenant)	ora_02.vmc.local	ora_02_data ora_02_logs	Oracle Archive Logs Backup Oracle Online Full Backup	10/06/2023 6:20:23 PM	Backup succeeded
	cdb3test	Single Instance (Multitenant)	ora_02.vmc.local				Not protected
	cdb4	Single Instance (Multitenant)	ora_02.vmc.local	ora_02_data ora_02_logs	Oracle Archive Logs Backup Oracle Online Full Backup	10/06/2023 6:20:23 PM	Backup succeeded

10. From DB VM ora_02, the cloned database is also registered in Oracle Restart grid control and the

dropped test table is recovered in the cloned database cdb3tst as shown below.

```
[oracle@ora_02 ~]$ /u01/app/oracle/product/19.0.0/grid/bin/crsctl
stat res -t
-----
-----
Name          Target  State          Server          State
details
-----
-----
Local Resources
-----
-----
ora.DATA.dg
          ONLINE ONLINE          ora_02          STABLE
ora.LISTENER.lsnr
          ONLINE INTERMEDIATE ora_02          Not All
Endpoints Re
gistered, STABLE
ora.LOGS.dg
          ONLINE ONLINE          ora_02          STABLE
ora.SC_2090922_CDB3TST.dg
          ONLINE ONLINE          ora_02          STABLE
ora.asm
          ONLINE ONLINE          ora_02
Started, STABLE
ora.ons
          OFFLINE OFFLINE        ora_02          STABLE
-----
-----
Cluster Resources
-----
-----
ora.cdb3.db
      1      ONLINE ONLINE          ora_02
Open, HOME=/u01/app/o
racle/product/19.0.0
/cdb3, STABLE
ora.cdb3tst.db
      1      ONLINE ONLINE          ora_02
Open, HOME=/u01/app/o
racle/product/19.0.0
```

```

/cdb4,STABLE
ora.cdb4.db
      1      ONLINE  ONLINE      ora_02
Open,HOME=/u01/app/o

racle/product/19.0.0

/cdb4,STABLE
ora.cssd
      1      ONLINE  ONLINE      ora_02      STABLE
ora.diskmon
      1      OFFLINE OFFLINE      STABLE
ora.driver.afd
      1      ONLINE  ONLINE      ora_02      STABLE
ora.evmd
      1      ONLINE  ONLINE      ora_02      STABLE

```

```

-----
-----

```

```

[oracle@ora_02 ~]$ export
ORACLE_HOME=/u01/app/oracle/product/19.0.0/cdb4
[oracle@ora_02 ~]$ export ORACLE_SID=cdb3tst
[oracle@ora_02 ~]$ sqlplus / as sysdba

```

```

SQL*Plus: Release 19.0.0.0.0 - Production on Sat Oct 7 08:04:51 2023
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 from v$database;

```

```

NAME      OPEN_MODE
-----
CDB3TST   READ WRITE

```

```

SQL> show pdbs

```

```

      CON_ID CON_NAME      OPEN MODE  RESTRICTED
-----
      2 PDB$SEED      READ ONLY  NO

```

```

        3 CDB3_PDB1                READ WRITE NO
        4 CDB3_PDB2                READ WRITE NO
        5 CDB3_PDB3                READ WRITE NO
SQL> alter session set container=CDB3_PDB1;

Session altered.

SQL> select * from test;

          ID
-----
DT
-----
EVENT
-----
          1
06-OCT-23 03.18.24.000000 PM
test oracle recovery on guest mounted fsx storage to VMC guest vm
ora_02

```

```
SQL>
```

This completes the demonstration of SnapCenter backup, restore, and clone of Oracle database in VMC SDDC on AWS.

Where to find additional information

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

- VMware Cloud on AWS Documentation

<https://docs.vmware.com/en/VMware-Cloud-on-AWS/index.html>

- Installing Oracle Grid Infrastructure for a Standalone Server with a New Database Installation

<https://docs.oracle.com/en/database/oracle/oracle-database/19/ladbi/installing-oracle-grid-infrastructure-for-a-standalone-server-with-a-new-database-installation.html#GUID-0B1CEE8C-C893-46AA-8A6A-7B5FAAEC72B3>

- Installing and Configuring Oracle Database Using Response Files

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

- Amazon FSx for NetApp ONTAP

TR-4981: Oracle Active Data Guard Cost Reduction with Amazon FSx ONTAP

Allen Cao, Niyaz Mohamed, NetApp

This solution provides overview and details for configuring Oracle Data Guard using AWS FSx ONTAP as standby site Oracle database storage to reduce licensing and operational cost of Oracle Data Guard HA/DR solution in AWS.

Purpose

Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data in a primary database and standby database replication configuration. Oracle Active Data Guard empowers users to access standby databases while data replication is active from the primary database to standby databases. Data Guard is a feature of Oracle Database Enterprise Edition. It does not require separate licensing. On the other hand, Active Data Guard is an Oracle Database Enterprise Edition Option therefore requires separate licensing. Multiple standby databases can receive data replication from a primary database in the Active Data Guard setup. However, each additional standby database requires an Active Data Guard license and extra storage as the size of primary database. The operational costs add up quickly.

If you are keen on cutting back cost of your Oracle database operation and are planning to set up an Active Data Guard in AWS, you should consider an alternative. Instead of Active Data Guard, use Data Guard to replicate from primary database to a single physical standby database on Amazon FSx ONTAP storage. Subsequently, multiple copies of this standby database can be cloned and opened for read/write access to serve many other use cases such as reporting, development, test etc. The net results effectively deliver functionalities of Active Data Guard while eliminating Active Data Guard license and extra storage cost for each additional standby database. In this documentation, we demonstrate how to setup an Oracle Data Guard with your existing primary database in AWS and place physical standby database on Amazon FSx ONTAP storage. The standby database is backed up via snapshot and cloned for read/write access for use cases as desired.

This solution addresses the following use cases:

- Oracle Data Guard between a primary database on any storage in AWS to standby database on Amazon FSx ONTAP storage.
- Clone the standby database while closed for data replication to serve use cases such as reporting, dev, test, etc.

Audience

This solution is intended for the following people:

- A DBA who set up Oracle Active Data Guard in AWS for high availability, data protection, and disaster recovery.
- A database solution architect interested in Oracle Active Data Guard configuration in the AWS cloud.
- A storage administrator who manages AWS FSx ONTAP storage that supports Oracle Data Guard.
- An application owner who like to stand up Oracle Data Guard in AWS FSx/EC2 environment.

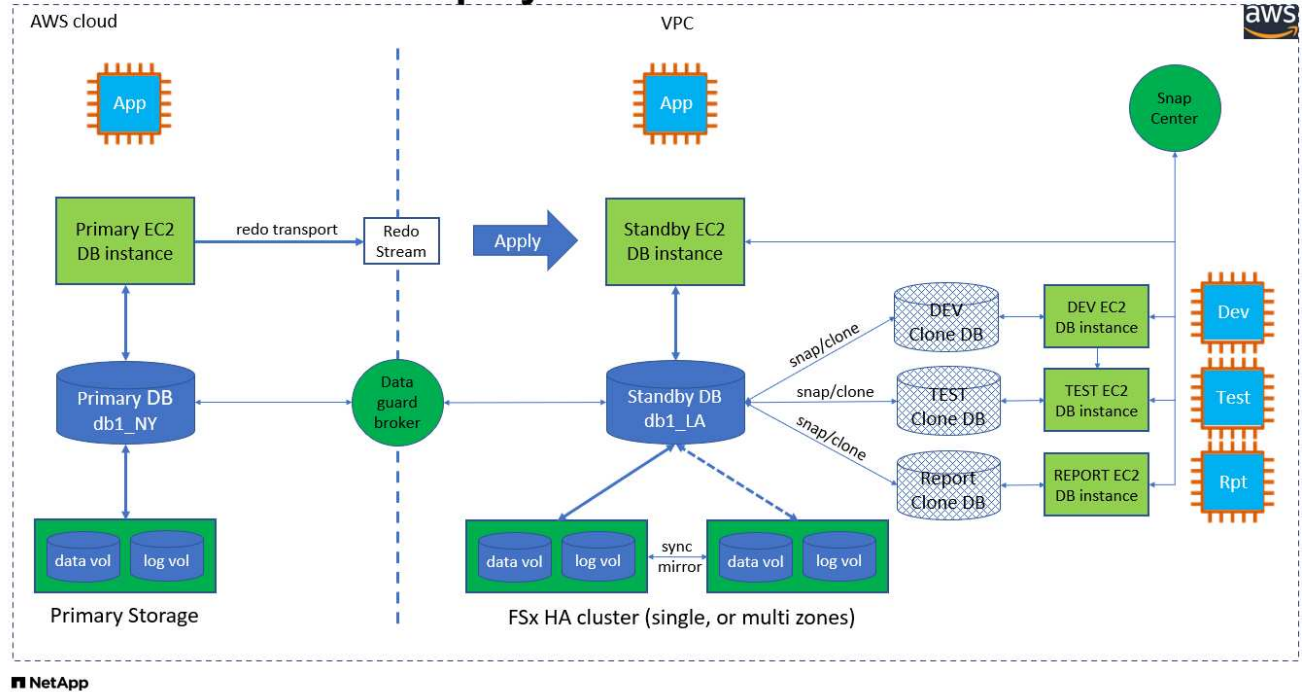
Solution test and validation environment

The testing and validation of this solution was performed in an AWS FSx ONTAP and EC2 lab environment that

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

Architecture

Oracle Data Guard Deployment with Amazon FSx for ONTAP



Hardware and software components

Hardware

FSx ONTAP storage	Current version offered by AWS	One FSx HA cluster in the same VPC and availability zone
EC2 instance for compute	t2.xlarge/4vCPU/16G	Three EC2 T2 xlarge EC2 instances, one as primary DB server, one as standby DB server, and the third as a clone DB server

Software

RedHat Linux	RHEL-8.6.0_HVM-20220503-x86_64-2-Hourly2-GP2	Deployed RedHat subscription for testing
Oracle Grid Infrastructure	Version 19.18	Applied RU patch p34762026_190000_Linux-x86-64.zip
Oracle Database	Version 19.18	Applied RU patch p34765931_190000_Linux-x86-64.zip
Oracle OPatch	Version 12.2.0.1.36	Latest patch p6880880_190000_Linux-x86-64.zip

Oracle Data Guard configuration with hypothetical NY to LA DR setup

Database	DB_UNIQUE_NAME	Oracle Net Service Name
Primary	db1_NY	db1_NY.demo.netapp.com
Physical Standby	db1_LA	db1_LA.demo.netapp.com

Key factors for deployment consideration

- **How Oracle Standby Database FlexClone Works.** AWS FSx ONTAP FlexClone provides shared copies of the same standby database volumes that are writable. The copies of the volumes are actually pointers that link back to original data blocks until a new write initiates on the clone. ONTAP then allocates new storage blocks for the new writes. Any read IOs are serviced by original data blocks under active replication. Thus, the clone are very storage efficient that can be used for many other use cases with minimal and incremental new storage allocation for new write IOs. This provides tremendous storage cost saving by substantially reducing Active Data Guard storage footprint. NetApp recommends to minimize FlexClone activities in the event of database switching over from primary storage to standby FSx storage in order to maintain Oracle performance at high level.
- **Oracle Software Requirements.** In general, a physical standby database must have the same Database Home version as the primary database including Patch Set Exceptions (PSEs), Critical Patch Updates (CPUs), and Patch Set Updates (PSUs), unless an Oracle Data Guard Standby-First Patch Apply process is in progress (as described in My Oracle Support note 1265700.1 at support.oracle.com)
- **Standby Database Directory Structure Considerations.** If possible, the data files, log files, and control files on the primary and standby systems should have the same names and path names and use Optimal Flexible Architecture (OFA) naming conventions. The archival directories on the standby database should also be identical between sites, including size and structure. This strategy allows other operations such as backups, switchovers, and failovers to execute the same set of steps, reducing the maintenance complexity.
- **Force Logging Mode.** To protect against unlogged direct writes in the primary database that cannot be propagated to the standby database, turn on FORCE LOGGING at the primary database before performing data file backups for standby creation.
- **Database Storage Management.** For operational simplicity, Oracle recommends that when you set up Oracle Automatic Storage Management (Oracle ASM) and Oracle Managed Files (OMF) in an Oracle Data Guard configuration that you set it up symmetrically on the primary and standby database(s).
- **EC2 compute instances.** In these tests and validations, we used an AWS EC2 t2.xlarge instance as the Oracle database compute instance. NetApp recommends using a M5 type EC2 instance as the compute instance for Oracle in production deployment because it is optimized for database workload. You need to size the EC2 instance appropriately for the number of vCPUs and the amount of RAM based on actual workload requirements.
- **FSx storage HA clusters single- or multi-zone deployment.** In these tests and validations, we deployed an FSx HA cluster in a single AWS availability zone. For production deployment, NetApp recommends deploying an FSx HA pair in two different availability zones. An FSx cluster is always provisioned in a HA pair that is sync mirrored in a pair of active-passive file systems to provide storage-level redundancy. Multi-zone deployment further enhances high availability in the event of failure in a single AWS zone.
- **FSx storage cluster sizing.** An Amazon FSx for ONTAP storage file system provides up to 160,000 raw SSD IOPS, up to 4GBps throughput, and a maximum of 192TiB capacity. However, you can size the cluster in terms of provisioned IOPS, throughput, and the storage limit (minimum 1,024 GiB) based on your actual requirements at the time of deployment. The capacity can be adjusted dynamically on the fly without affecting application availability.

Solution deployment

It is assumed that you already have your primary Oracle database deployed in AWS EC2 environment within a VPC as the starting point for setting up Data Guard. The primary database is deployed using Oracle ASM for storage management. Two ASM disk groups - +DATA and +LOGS are created for Oracle data files, log files, and control file etc. For details on Oracle deployment in AWS with ASM, please refer to following technical reports for help.

- [Oracle Database Deployment on EC2 and FSx Best Practices](#)
- [Oracle Database Deployment and Protection in AWS FSx/EC2 with iSCSI/ASM](#)
- [Oracle 19c in Standalone Restart on AWS FSx/EC2 with NFS/ASM](#)

Your primary Oracle database can be running either on an FSx ONTAP or any other storage of choices within the AWS EC2 ecosystem. The following section provides step-by-step deployment procedures for setting up Oracle Data Guard between a primary EC2 DB instance with ASM storage to a standby EC2 DB instance with ASM storage.

Prerequisites for deployment

Deployment requires the following prerequisites.

1. An AWS account has been set up, and the necessary VPC and network segments have been created within your AWS account.
2. From the AWS EC2 console, you need to deploy minimum three EC2 Linux instances, one as the primary Oracle DB instance, one as standby Oracle DB instance, and an clone target DB instance for reporting, dev, and test etc. See the architecture diagram in the previous section for more details about the environment setup. Also review the AWS [User Guide for Linux instances](#) for more information.
3. From the AWS EC2 console, deploy Amazon FSx for ONTAP storage HA clusters to host Oracle volumes that stores the Oracle standby database. If you are not familiar with the deployment of FSx storage, see the documentation [Creating FSx for ONTAP file systems](#) for step-by-step instructions.
4. Steps 2 and 3 can be performed using the following Terraform automation toolkit, which creates an EC2 instance named `ora_01` and an FSx file system named `fsx_01`. Review the instruction carefully and change the variables to suit your environment before execution. The template can be easily revised for your own deployment requirements.

```
git clone https://github.com/NetApp-  
Automation/na_aws_fsx_ec2_deploy.git
```



Ensure that you have allocated at least 50G in EC2 instance root volume in order to have sufficient space to stage Oracle installation files.

Prepare the primary database for Data Guard

In this demonstration, we have setup a primary Oracle database called db1 on the primary EC2 DB instance with two ASM disk groups in standalone Restart configuration with data files in ASM disk group +DATA and flash recovery area in ASM disk group +LOGS. Following illustrates the detailed procedures for setting up primary database for Data Guard. All steps should be executed as database owner - oracle user.

1. Primary database db1 configuration on primary EC2 DB instance ip-172-30-15-45. The ASM disk groups can be on any type of storage within EC2 ecosystem.

```
[oracle@ip-172-30-15-45 ~]$ cat /etc/oratab

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

# A colon, ':', is used as the field terminator.  A new line
terminates
# the entry.  Lines beginning with a pound sign, '#', are comments.
#
# Entries are of the form:
#   $ORACLE_SID:$ORACLE_HOME:<N|Y>:
#
# The first and second fields are the system identifier and home
# directory of the database respectively.  The third field indicates
# to the dbstart utility that the database should , "Y", or should
not,
# "N", be brought up at system boot time.
#
# Multiple entries with the same $ORACLE_SID are not allowed.
#
#
+ASM:/u01/app/oracle/product/19.0.0/grid:N
db1:/u01/app/oracle/product/19.0.0/db1:N

[oracle@ip-172-30-15-45 ~]$
/u01/app/oracle/product/19.0.0/grid/bin/crsctl stat res -t
-----
-----
Name          Target  State          Server          State
details
-----
-----
Local Resources
-----
```

```

-----
ora.DATA.dg
      ONLINE  ONLINE      ip-172-30-15-45      STABLE
ora.LISTENER.lsnr
      ONLINE  ONLINE      ip-172-30-15-45      STABLE
ora.LOGS.dg
      ONLINE  ONLINE      ip-172-30-15-45      STABLE
ora.asm
      ONLINE  ONLINE      ip-172-30-15-45
Started,STABLE
ora.ons
      OFFLINE OFFLINE      ip-172-30-15-45      STABLE
-----
Cluster Resources
-----
ora.cssd
      1      ONLINE  ONLINE      ip-172-30-15-45      STABLE
ora.dbf.db
      1      ONLINE  ONLINE      ip-172-30-15-45
Open,HOME=/u01/app/o
racle/product/19.0.0
/db1,STABLE
ora.diskmon
      1      OFFLINE OFFLINE      STABLE
ora.driver.afd
      1      ONLINE  ONLINE      ip-172-30-15-45      STABLE
ora.evmd
      1      ONLINE  ONLINE      ip-172-30-15-45      STABLE
-----

```

2. From sqlplus, enable forced logging on primary.

```
alter database force logging;
```

3. From sqlplus, enable flashback on primary. Flashback allows easy reinstate primary database as a standby after a failover.

```
alter database flashback on;
```

4. Configure redo transport authentication using Oracle password file - create a pwd file on the primary using orapwd utility if not set and copy over to standby database \$ORACLE_HOME/dbs directory.
5. Create standby redo logs on the primary DB with same size as current online log file. Log groups are one more than online log file groups. The primary database can then quickly transition to the standby role and begin receiving redo data, if necessary.

```
alter database add standby logfile thread 1 size 200M;
```

Validate after standby logs addition:

```
SQL> select group#, type, member from v$logfile;
```

GROUP#	TYPE	MEMBER
3	ONLINE	+DATA/DB1/ONLINELOG/group_3.264.1145821513
2	ONLINE	+DATA/DB1/ONLINELOG/group_2.263.1145821513
1	ONLINE	+DATA/DB1/ONLINELOG/group_1.262.1145821513
4	STANDBY	+DATA/DB1/ONLINELOG/group_4.286.1146082751
4	STANDBY	+LOGS/DB1/ONLINELOG/group_4.258.1146082753
5	STANDBY	+DATA/DB1/ONLINELOG/group_5.287.1146082819
5	STANDBY	+LOGS/DB1/ONLINELOG/group_5.260.1146082821
6	STANDBY	+DATA/DB1/ONLINELOG/group_6.288.1146082825
6	STANDBY	+LOGS/DB1/ONLINELOG/group_6.261.1146082827
7	STANDBY	+DATA/DB1/ONLINELOG/group_7.289.1146082835
7	STANDBY	+LOGS/DB1/ONLINELOG/group_7.262.1146082835

11 rows selected.

6. From sqlplus, create a pfile from spfile for editing.

```
create pfile='/home/oracle/initdb1.ora' from spfile;
```

7. Revise the pfile and add following parameters.

```
DB_NAME=db1
DB_UNIQUE_NAME=db1_NY
LOG_ARCHIVE_CONFIG='DG_CONFIG=(db1_NY,db1_LA) '
LOG_ARCHIVE_DEST_1='LOCATION=USE_DB_RECOVERY_FILE_DEST
VALID_FOR=(ALL_LOGFILES,ALL_ROLES) DB_UNIQUE_NAME=db1_NY '
LOG_ARCHIVE_DEST_2='SERVICE=db1_LA ASYNC
VALID_FOR=(ONLINE_LOGFILES,PRIMARY_ROLE) DB_UNIQUE_NAME=db1_LA '
REMOTE_LOGIN_PASSWORDFILE=EXCLUSIVE
FAL_SERVER=db1_LA
STANDBY_FILE_MANAGEMENT=AUTO
```

8. From sqlplus, create spfile in ASM +DATA directory from revised pfile in /home/oracle directory.

```
create spfile='+DATA' from pfile='/home/oracle/initdb1.ora';
```

9. Locate the newly created spfile under +DATA disk group(using asmcmd utility if necessary). Use srvctl to modify grid to start database from new spfile as shown below.

```

[oracle@ip-172-30-15-45 db1]$ srvctl config database -d db1
Database unique name: db1
Database name: db1
Oracle home: /u01/app/oracle/product/19.0.0/db1
Oracle user: oracle
Spfile: +DATA/DB1/PARAMETERFILE/spfile.270.1145822903
Password file:
Domain: demo.netapp.com
Start options: open
Stop options: immediate
Database role: PRIMARY
Management policy: AUTOMATIC
Disk Groups: DATA
Services:
OSDBA group:
OSOPER group:
Database instance: db1
[oracle@ip-172-30-15-45 db1]$ srvctl modify database -d db1 -spfile
+DATA/DB1/PARAMETERFILE/spfiledb1.ora
[oracle@ip-172-30-15-45 db1]$ srvctl config database -d db1
Database unique name: db1
Database name: db1
Oracle home: /u01/app/oracle/product/19.0.0/db1
Oracle user: oracle
Spfile: +DATA/DB1/PARAMETERFILE/spfiledb1.ora
Password file:
Domain: demo.netapp.com
Start options: open
Stop options: immediate
Database role: PRIMARY
Management policy: AUTOMATIC
Disk Groups: DATA
Services:
OSDBA group:
OSOPER group:
Database instance: db1

```

10. Modify tnsnames.ora to add db_unique_name for name resolution.


```
# tnsnames.ora Network Configuration File:
/u01/app/oracle/product/19.0.0/db1/network/admin/tnsnames.ora
# Generated by Oracle configuration tools.

db1_NY =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = ip-172-30-15-
45.ec2.internal) (PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SID = db1)
    )
  )

db1_LA =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = ip-172-30-15-
67.ec2.internal) (PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SID = db1)
    )
  )

LISTENER_DB1 =
  (ADDRESS = (PROTOCOL = TCP) (HOST = ip-172-30-15-
45.ec2.internal) (PORT = 1521))
```

11. Add data guard service name db1_NY_DGMGRL.demo.netapp for primary database to listener.ora file.

```
#Backup file is /u01/app/oracle/crsdata/ip-172-30-15-45/output/listener.ora.bak.ip-172-30-15-45.oracle line added by Agent
# listener.ora Network Configuration File:
/u01/app/oracle/product/19.0.0/grid/network/admin/listener.ora
# Generated by Oracle configuration tools.
```

```
LISTENER =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS = (PROTOCOL = TCP) (HOST = ip-172-30-15-45.ec2.internal) (PORT = 1521))
      (ADDRESS = (PROTOCOL = IPC) (KEY = EXTPROC1521))
    )
  )
```

```
SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (GLOBAL_DBNAME = db1_NY_DGMGRL.demo.netapp.com)
      (ORACLE_HOME = /u01/app/oracle/product/19.0.0/db1)
      (SID_NAME = db1)
    )
  )
```

```
ENABLE_GLOBAL_DYNAMIC_ENDPOINT_LISTENER=ON # line added by Agent
VALID_NODE_CHECKING_REGISTRATION_LISTENER=ON # line added by Agent
```

1. Shutdown and restart database with srvctl and validate that data guard parameters are now active.

```
srvctl stop database -d db1
```

```
srvctl start database -d db1
```

This completes primary database setup for Data Guard.

Prepare standby database and activate Data Guard

Oracle Data Guard requires OS kernel configuration and Oracle software stacks including patch sets on standby EC2 DB instance to match with primary EC2 DB instance. For easy management and simplicity, the standby EC2 DB instance database storage configuration ideally should match with the primary EC2 DB instance as well, such as the name, number and size of ASM disk groups. Following are detail procedures for setting up the standby EC2 DB instance for Data Guard. All commands should be executed as oracle owner user id.

1. First, review the configuration of the primary database on primary EC2 instance. In this demonstration, we have setup a primary Oracle database called db1 on the primary EC2 DB instance with two ASM disk groups +DATA and +LOGS in standalone Restart configuration. The primary ASM disk groups may be on any type of storage within EC2 ecosystem.
2. Follow procedures in documentation [TR-4965: Oracle Database Deployment and Protection in AWS FSx/EC2 with iSCSI/ASM](#) to install and configure grid and Oracle on standby EC2 DB instance to match with primary database. The database storage should be provisioned and allocated to standby EC2 DB instance from FSx ONTAP with same storage capacity as primary EC2 DB instance.



Stop at step 10 in Oracle database installation section. The standby database will be instantiated from primary database using dbca database duplication function.

3. Once Oracle software is installed and configured, from standby \$ORACLE_HOME dbs directory, copy oracle password from primary database.

```
scp
oracle@172.30.15.45:/u01/app/oracle/product/19.0.0/db1/dbs/orapwdb1
.
```

4. Create tnsnames.ora file with following entries.

```
# tnsnames.ora Network Configuration File:
/u01/app/oracle/product/19.0.0/db1/network/admin/tnsnames.ora
# Generated by Oracle configuration tools.

db1_NY =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = ip-172-30-15-
45.ec2.internal) (PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SID = db1)
    )
  )

db1_LA =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = ip-172-30-15-
67.ec2.internal) (PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SID = db1)
    )
  )
```

5. Add DB data guard service name to listener.ora file.

```

#Backup file is /u01/app/oracle/crsdata/ip-172-30-15-
67/output/listener.ora.bak.ip-172-30-15-67.oracle line added by
Agent
# listener.ora Network Configuration File:
/u01/app/oracle/product/19.0.0/grid/network/admin/listener.ora
# Generated by Oracle configuration tools.

LISTENER =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS = (PROTOCOL = TCP) (HOST = ip-172-30-15-
67.ec2.internal) (PORT = 1521))
      (ADDRESS = (PROTOCOL = IPC) (KEY = EXTPROC1521))
    )
  )

SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (GLOBAL_DBNAME = db1_LA_DGMGRL.demo.netapp.com)
      (ORACLE_HOME = /u01/app/oracle/product/19.0.0/db1)
      (SID_NAME = db1)
    )
  )

ENABLE_GLOBAL_DYNAMIC_ENDPOINT_LISTENER=ON # line added
by Agent
VALID_NODE_CHECKING_REGISTRATION_LISTENER=ON # line added
by Agent

```

6. Set oracle home and path.

```
export ORACLE_HOME=/u01/app/oracle/product/19.0.0/db1
```

```
export PATH=$PATH:$ORACLE_HOME/bin
```

7. Use dbca to instantiate standby database from primary database db1.

```

[oracle@ip-172-30-15-67 bin]$ dbca -silent -createDuplicateDB
-gdbName db1 -primaryDBConnectionString ip-172-30-15-
45.ec2.internal:1521/db1_NY.demo.netapp.com -sid db1 -initParams
fal_server=db1_NY -createAsStandby -dbUniqueName db1_LA
Enter SYS user password:

Prepare for db operation
22% complete
Listener config step
44% complete
Auxiliary instance creation
67% complete
RMAN duplicate
89% complete
Post duplicate database operations
100% complete

Look at the log file
"/u01/app/oracle/cfgtoollogs/dbca/db1_LA/db1_LA.log" for further
details.

```

8. Validate duplicated standby database. Newly duplicated standby database open in READ ONLY mode initially.

```

[oracle@ip-172-30-15-67 bin]$ export ORACLE_SID=db1
[oracle@ip-172-30-15-67 bin]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Wed Aug 30 18:25:46
2023
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 from v$database;

NAME          OPEN_MODE
-----
DB1           READ ONLY

```

```
SQL> show parameter name
```

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

cdb_cluster_name	string	
cell_offloadgroup_name	string	
db_file_name_convert	string	
db_name	string	db1
db_unique_name	string	db1_LA
global_names	boolean	FALSE
instance_name	string	db1
lock_name_space	string	
log_file_name_convert	string	
pdb_file_name_convert	string	
processor_group_name	string	

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

service_names	string	
db1_LA.demo.netapp.com		

```
SQL>
```

```
SQL> show parameter log_archive_config
```

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

log_archive_config	string	
DG_CONFIG=(db1_NY,db1_LA)		

```
SQL> show parameter fal_server
```

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

fal_server	string	db1_NY

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

NAME

+DATA/DB1_LA/DATAFILE/system.261.1146248215
+DATA/DB1_LA/DATAFILE/sysaux.262.1146248231
+DATA/DB1_LA/DATAFILE/undotbs1.263.1146248247
+DATA/DB1_LA/03C5C01A66EE9797E0632D0F1EAC5F59/DATAFILE/system.264.11

```
46248253
+DATA/DB1_LA/03C5C01A66EE9797E0632D0F1EAC5F59/DATAFILE/sysaux.265.11
46248261
+DATA/DB1_LA/DATAFILE/users.266.1146248267
+DATA/DB1_LA/03C5C01A66EE9797E0632D0F1EAC5F59/DATAFILE/undotbs1.267.
1146248269
+DATA/DB1_LA/03C5EFD07C41A1FAE0632D0F1EAC9BD8/DATAFILE/system.268.11
46248271
+DATA/DB1_LA/03C5EFD07C41A1FAE0632D0F1EAC9BD8/DATAFILE/sysaux.269.11
46248279
+DATA/DB1_LA/03C5EFD07C41A1FAE0632D0F1EAC9BD8/DATAFILE/undotbs1.270.
1146248285
+DATA/DB1_LA/03C5EFD07C41A1FAE0632D0F1EAC9BD8/DATAFILE/users.271.114
6248293
```

NAME

```
-----
-----
+DATA/DB1_LA/03C5F0DDF35CA2B6E0632D0F1EAC8B6B/DATAFILE/system.272.11
46248295
+DATA/DB1_LA/03C5F0DDF35CA2B6E0632D0F1EAC8B6B/DATAFILE/sysaux.273.11
46248301
+DATA/DB1_LA/03C5F0DDF35CA2B6E0632D0F1EAC8B6B/DATAFILE/undotbs1.274.
1146248309
+DATA/DB1_LA/03C5F0DDF35CA2B6E0632D0F1EAC8B6B/DATAFILE/users.275.114
6248315
+DATA/DB1_LA/03C5F1C9B142A2F1E0632D0F1EACF21A/DATAFILE/system.276.11
46248317
+DATA/DB1_LA/03C5F1C9B142A2F1E0632D0F1EACF21A/DATAFILE/sysaux.277.11
46248323
+DATA/DB1_LA/03C5F1C9B142A2F1E0632D0F1EACF21A/DATAFILE/undotbs1.278.
1146248331
+DATA/DB1_LA/03C5F1C9B142A2F1E0632D0F1EACF21A/DATAFILE/users.279.114
6248337
```

19 rows selected.

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

NAME

```
-----
-----
+DATA/DB1_LA/CONTROLFILE/current.260.1146248209
+LOGS/DB1_LA/CONTROLFILE/current.257.1146248209
```

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



```
NAME
```

```
-----  
-----  
+DATA/DB1_LA/TEMPFILE/temp.287.1146248371  
+DATA/DB1_LA/03C5C01A66EE9797E0632D0F1EAC5F59/TEMPFILE/temp.288.1146  
248375  
+DATA/DB1_LA/03C5EFD07C41A1FAE0632D0F1EAC9BD8/TEMPFILE/temp.290.1146  
248463  
+DATA/DB1_LA/03C5F0DDF35CA2B6E0632D0F1EAC8B6B/TEMPFILE/temp.291.1146  
248463  
+DATA/DB1_LA/03C5F1C9B142A2F1E0632D0F1EACF21A/TEMPFILE/temp.292.1146  
248463
```

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

```
GROUP# TYPE MEMBER  
-----  
-----  
1 ONLINE +LOGS/DB1_LA/ONLINELOG/group_1.259.1146248349  
1 ONLINE +DATA/DB1_LA/ONLINELOG/group_1.280.1146248347  
2 ONLINE +DATA/DB1_LA/ONLINELOG/group_2.281.1146248351  
2 ONLINE +LOGS/DB1_LA/ONLINELOG/group_2.258.1146248353  
3 ONLINE +DATA/DB1_LA/ONLINELOG/group_3.282.1146248355  
3 ONLINE +LOGS/DB1_LA/ONLINELOG/group_3.260.1146248355  
4 STANDBY +DATA/DB1_LA/ONLINELOG/group_4.283.1146248357  
4 STANDBY +LOGS/DB1_LA/ONLINELOG/group_4.261.1146248359  
5 STANDBY +DATA/DB1_LA/ONLINELOG/group_5.284.1146248361  
5 STANDBY +LOGS/DB1_LA/ONLINELOG/group_5.262.1146248363  
6 STANDBY +LOGS/DB1_LA/ONLINELOG/group_6.263.1146248365  
6 STANDBY +DATA/DB1_LA/ONLINELOG/group_6.285.1146248365  
7 STANDBY +LOGS/DB1_LA/ONLINELOG/group_7.264.1146248369  
7 STANDBY +DATA/DB1_LA/ONLINELOG/group_7.286.1146248367
```

```
14 rows selected.
```

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

```
NAME OPEN_MODE  
-----  
DB1 READ ONLY
```

- Restart standby database in mount stage and execute following command to activate standby database managed recovery.

```
alter database recover managed standby database disconnect from
session;
```

```
SQL> shutdown immediate;
Database closed.
Database dismounted.
ORACLE instance shut down.
SQL> startup mount;
ORACLE instance started.
```

```
Total System Global Area 8053062944 bytes
Fixed Size                  9182496 bytes
Variable Size               1291845632 bytes
Database Buffers            6744440832 bytes
Redo Buffers                 7593984 bytes
```

```
Database mounted.
```

```
SQL> alter database recover managed standby database disconnect from
session;
```

```
Database altered.
```

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

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

ROLE	THREAD#	SEQUENCE#	ACTION
recovery apply slave	0	0	IDLE
recovery apply slave	0	0	IDLE
recovery apply slave	0	0	IDLE
recovery apply slave	0	0	IDLE
recovery logmerger	1	30	APPLYING_LOG
RFS ping	1	30	IDLE
RFS async	1	30	IDLE
archive redo	0	0	IDLE
archive redo	0	0	IDLE
archive redo	0	0	IDLE
gap manager	0	0	IDLE

ROLE	THREAD#	SEQUENCE#	ACTION
managed recovery	0	0	IDLE
redo transport monitor	0	0	IDLE
log writer	0	0	IDLE
archive local	0	0	IDLE
redo transport timer	0	0	IDLE

```
16 rows selected.
```

```
SQL>
```

This completes the Data Guard protection setup for db1 from primary to standby with managed standby recovery enabled.

Setup Data Guard Broker

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

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

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

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

```
[oracle@ip-172-30-15-45 db1]$ dgmgrl sys@db1_NY
DGMGRL for Linux: Release 19.0.0.0.0 - Production on Wed Aug 30
19:34:14 2023
Version 19.18.0.0.0

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

Welcome to DGMGRL, type "help" for information.
Password:
Connected to "db1_NY"
Connected as SYSDBA.
```

3. Create and enable Data Guard Broker configuration.

```
DGMGRL> create configuration dg_config as primary database is db1_NY
connect identifier is db1_NY;
Configuration "dg_config" created with primary database "db1_ny"
DGMGRL> add database db1_LA as connect identifier is db1_LA;
Database "db1_la" added
DGMGRL> enable configuration;
Enabled.
DGMGRL> show configuration;

Configuration - dg_config

Protection Mode: MaxPerformance
Members:
db1_ny - Primary database
db1_la - Physical standby database

Fast-Start Failover: Disabled

Configuration Status:
SUCCESS (status updated 28 seconds ago)
```

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

```
DGMGRL> show database db1_ny;
```

```
Database - db1_ny
```

```
Role:                PRIMARY
Intended State:      TRANSPORT-ON
Instance(s):        db1
```

```
Database Status:
SUCCESS
```

```
DGMGRL> show database db1_la;
```

```
Database - db1_la
```

```
Role:                PHYSICAL STANDBY
Intended State:      APPLY-ON
Transport Lag:       0 seconds (computed 1 second ago)
Apply Lag:           0 seconds (computed 1 second ago)
Average Apply Rate: 2.00 KByte/s
Real Time Query:    OFF
Instance(s):        db1
```

```
Database Status:
SUCCESS
```

```
DGMGRL>
```

In the event of a failure, Data Guard Broker can be used to failover primary database to standby instantaneously.

Clone standby database for other use cases

The key benefit of staging standby database on AWS FSx ONTAP in Data Guard is that it can be FlexCloned to serve many other use cases with minimal additional storage investment. In the following section, we demonstrate how to snapshot and clone the mounted and under recovery standby database volumes on FSx ONTAP for other purposes, such as DEV, TEST, REPORT, etc., using the NetApp SnapCenter tool.

Following are high level procedures to clone a READ/WRITE database from the managed physical standby database in Data Guard using SnapCenter. For detail instructions on how to setup and configure SnapCenter, please refer to [Hybrid Cloud Database Solutions with SnapCenter](#) relevant Oracle sections.

1. We begin with creating a test table and inserting a row into the test table on primary database. We will then validate if the transaction traverse down to standby and finally the clone.

```
[oracle@ip-172-30-15-45 db1]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Thu Aug 31 16:35:53
2023
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 session set container=db1_pdb1;

Session altered.

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

Table created.

SQL> insert into test values(1, sysdate, 'a test transaction on
primary database db1 and ec2 db host: ip-172-30-15-
45.ec2.internal');

1 row created.

SQL> commit;

Commit complete.
```

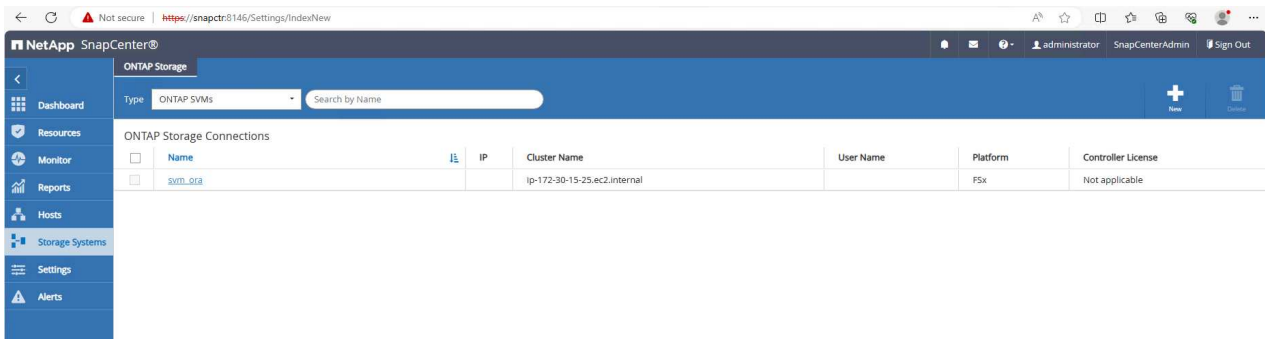
```
SQL> select * from test;
```

```
          ID
-----
DT
-----
EVENT
-----
          1
31-AUG-23 04.49.29.000000 PM
a test transaction on primary database db1 and ec2 db host: ip-172-
30-15-45.ec2.
internal
```


```
SQL> select instance_name, host_name from v$instance;
```

```
INSTANCE_NAME
-----
HOST_NAME
-----
db1
ip-172-30-15-45.ec2.internal
```


2. Add FSx storage cluster to Storage Systems in SnapCenter with FSx cluster management IP and fsxadmin credential.



3. Add AWS ec2-user to Credential in Settings.




4. Add standby EC2 DB instance and clone EC2 DB instance to Hosts.

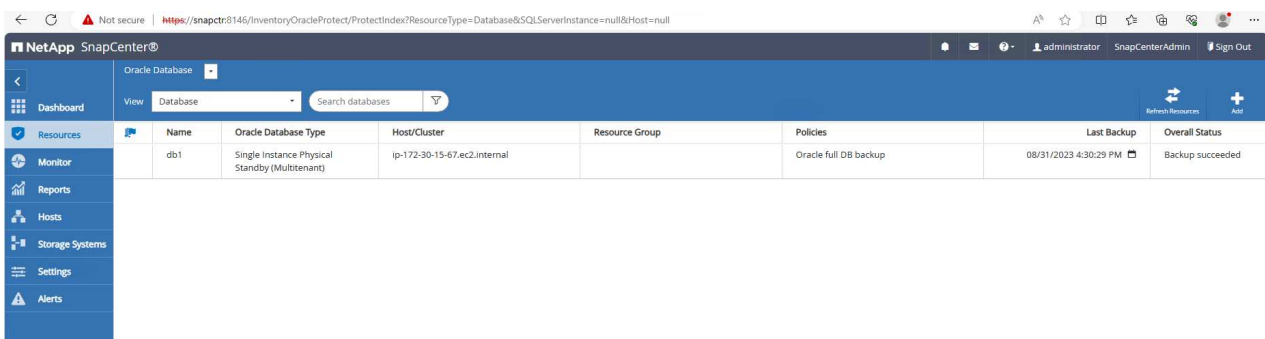


The clone EC2 DB instance should have similar Oracle software stacks installed and configured. In our test case, the grid infrastructure and Oracle 19C installed and configured but no database created.

5. Create a backup policy that is tailored for offline/mount full database backup.




6. Apply backup policy to protect standby database in Resources tab.




7. Click on database name to open the database backups page. Select a backup to be used for

database clone and click on Clone button to launch clone workflow.



8. Select Complete Database Clone and name the clone instance SID.



9. Select the clone host, which hosts the cloned database from standby DB. Accept the default for data files, control files, and redo logs. Two ASM disk groups will be created on the clone host that are corresponding to the disk groups on standby database.

x
Clone from db1

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Select the host to create a clone

Clone host

Datafile locations i

+SC_2090922_db1dev

+SC_2342319_db1dev

Control files i

+SC_2090922_db1dev/db1dev/control/control01.ctl x

+SC_2090922_db1dev/db1dev/control/control02.ctl x

Redo logs i

Group	Size	Unit	Number of files
▶ RedoGroup 1 x	200	MB	2 +
▶ RedoGroup 2 x	200	MB	2 +
▶ RedoGroup 3 x	200	MB	2 +

10. No database credentials are needed for OS based authentication. Match Oracle home setting with what is configured on the clone EC2 database instance.

Clone from db1 x

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Database Credentials for the clone

Credential name for sys user + ⓘ

ASM instance Credential name + ⓘ

Database port

ASM Port

Oracle Home Settings ⓘ

Oracle Home

Oracle OS User

Oracle OS Group

11. Change clone database parameters if needed and specify scripts to run before cloen if any.

Clone from db1
✕

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps
- 6 Notification
- 7 Summary

Specify scripts to run before clone operation ?

Prescript full path

Arguments

Script timeout secs

Database Parameter settings

audit_file_dest	/u01/app/oracle/admin/db1dev_LA/adump	✕
audit_trail	DB	✕
open_cursors	300	✕
pga_aggregate_target	2684354560	✕

12. Enter SQL to run after clone. In the demo, we executed commands to turn off database archive mode for a dev/test/report database.

Clone from db1 ✕

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps**
- 6 Notification
- 7 Summary

Until Cancel recovery will be performed for Physical Standby Dataguard/Active Dataguard database.

Create new DBID ⓘ

Create tempfile for temporary tablespace ⓘ

Enter SQL queries to apply when clone is created

shutdown immediate ; startup mount ; alter database noarchivelog ; alter database open ; + Reset

Enter scripts to run after clone operation ⓘ

Previous Next

13. Configure email notification if desired.

Clone from db1 ×

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Provide email settings ⓘ

Email preference

From

To

Subject

Attach job report

14. Review the summary, click `Finish` to start the clone.

x
Clone from db1


- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps
- 6 Notification
- 7 Summary

Summary


Clone from backup	db1_LA_08-31-2023_17.42.01.6804_0
Clone SID	db1dev
Clone server	ip-172-30-15-126.ec2.internal
Exclude PDBs	none
Oracle home	/u01/app/oracle/product/19.0.0/dev
Oracle OS user	oracle
Oracle OS group	oinstall
Datafile mountpaths	+SC_2090922_db1dev +SC_2342319_db1dev
Control files	+SC_2090922_db1dev/db1dev/control/control01.ctl +SC_2090922_db1dev/db1dev/control/control02.ctl
Redo groups	RedoGroup =1 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redo01_01.log RedoGroup =1 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redo01_02.log RedoGroup =2 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redo02_01.log RedoGroup =2 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redo02_02.log RedoGroup =3 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redo03_01.log RedoGroup =3 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redo03_02.log RedoGroup =4 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redo04_01.log RedoGroup =4 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redo04_02.log RedoGroup =5 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redo05_01.log RedoGroup =5 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redo05_02.log RedoGroup =6 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redo06_01.log RedoGroup =6 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redo06_02.log

Previous
Finish

15. Monitor clone job in Monitor tab. We observed that it took around 8 minutes to clone a database about 300GB in database volume size.



16. Validate the clone database from SnapCenter, which is immediately registered in Resources tab right after clone operation.



17. Query the clone database from clone EC2 instance. We validated that test transaction that occurred in primary database had traversed down to clone database.

```
[oracle@ip-172-30-15-126 ~]$ export
ORACLE_HOME=/u01/app/oracle/product/19.0.0/dev
[oracle@ip-172-30-15-126 ~]$ export ORACLE_SID=db1dev
[oracle@ip-172-30-15-126 ~]$ export PATH=$PATH:$ORACLE_HOME/bin
[oracle@ip-172-30-15-126 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Wed Sep 6 16:41:41 2023
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
DB1DEV	READ WRITE	NOARCHIVELOG

```
SQL> select instance_name, host_name from v$instance;
```

INSTANCE_NAME	HOST_NAME
db1dev	ip-172-30-15-126.ec2.internal

```
SQL> alter session set container=db1_pdb1;
```

```
Session altered.
```

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

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

```
1
31-AUG-23 04.49.29.000000 PM
a test transaction on primary database db1 and ec2 db host: ip-172-
30-15-45.ec2.
internal

SQL>
```

This completes the clone and validation of a new Oracle database from standby database in Data Guard on FSx storage for DEV, TEST, REPORT or any other use cases. Multiple Oracle databases can be cloned off the same standby database in Data Guard.

Where to find additional information

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

- Data Guard Concepts and Administration

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

- WP-7357: Oracle Database Deployment on EC2 and FSx Best Practices

[Introduction](#)

- Amazon FSx for NetApp ONTAP

<https://aws.amazon.com/fsx/netapp-ontap/>

- Amazon EC2

https://aws.amazon.com/pm/ec2/?trk=36c6da98-7b20-48fa-8225-4784bced9843&sc_channel=ps&s_kwcid=AL!4422!3!467723097970!e!!g!!aws%20ec2&ef_id=Cj0KCQiA54KfBhCKARIsAJzSrdqwQrghn6l71jiWzSeaT9Uh1-vY-VfhJixF-xnv5rWwn2S7RqZOTQ0aAh7eEALw_wcB:G:s&s_kwcid=AL!4422!3!467723097970!e!!g!!aws%20ec2

TR-4973: Quick Recovery and Clone of Oracle VLDB with Incremental Merge on AWS FSx ONTAP

Allen Cao, Niyaz Mohamed, NetApp

This solution provides overview and details for quick recovery and clone of Oracle VLDB deployed to AWS EC2 compute instance with NFS mount on FSx ONTAP to staging a standby data file copy to be 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 AWS, you could set up a similar database protection structure using resources such as AWS FSx ONTAP for staging your standby database image copy. In this documentation, we demonstrate how to provision and export an NFS file system from AWS FSx ONTAP 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.

Better yet, we also show how you could leverage NetApp FlexClone to create a copy of the same staging NFS file system for other use cases such as standing up a dev/test Oracle environment with this same standby database image copy without additional storage investment.

This solution addresses the following use cases:

- An Oracle VLDB image copy incremental merge via RMAN on NFS mount point off AWS FSx ONTAP storage.
- Quick recovery of an Oracle VLDB by switching to database image copy on FSx ONTAP storage in the event of failure.
- Clone FSx ONTAP NFS file system volume storing an Oracle VLDB image copy to be used for standing up another database instance for other use cases.

Audience

This solution is intended for the following people:


- A DBA who set up Oracle VLDB image copy incremental merge via RMAN in AWS for faster database recovery.
- A database solution architect who tests Oracle workloads in the AWS public cloud.
- A storage administrator who manages Oracle databases deployed to AWS FSx ONTAP storage.
- An application owner who would like to stand up Oracle databases in AWS FSx/EC2 environment.

Solution test and validation environment

The testing and validation of this solution was performed in an AWS FSx ONTAP and EC2 environment that might not match the final deployment environment. For more information, see the section [Key factors for deployment consideration](#).

Architecture

Oracle VLDB Incremental Merge via RMAN on AWS FSxN



Hardware and software components

Hardware

FSx ONTAP storage	Current version offered by AWS	One FSx HA cluster in the same VPC and availability zone
EC2 instance for compute	t2.xlarge/4vCPU/16G	Two EC2 T2 xlarge EC2 instances, one as primary DB server and the other as a clone DB server

Software

RedHat Linux	RHEL-8.6.0_HVM-20220503-x86_64-2-Hourly2-GP2	Deployed RedHat subscription for testing
Oracle Grid Infrastructure	Version 19.18	Applied RU patch p34762026_190000_Linux-x86-64.zip
Oracle Database	Version 19.18	Applied RU patch p34765931_190000_Linux-x86-64.zip
Oracle OPatch	Version 12.2.0.1.36	Latest patch p6880880_190000_Linux-x86-64.zip

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 FSx file system, which has 4GBps throughput, 160,000 raw SSD IOPS, and 192TiB capacity limit. For deployment over the thresholds,

multiple FSx file systems 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 FSx database backup image copy. To minimize the data loss, Oracle flash recovery area can be setup on FSx NFS mount point and archived logs are backed up to FSx NFS mount along with database image copy.
- **Running Oracle VLDB off FSx NFS file system.** Unlike other bulk storage for database backup, AWS FSx ONTAP 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 FSx ONTAP 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.
- **FlexClone Oracle VLDB image copy of NFS volume for other use cases.** AWS FSx ONTAP FlexClone provides shared copies of the same NFS data volume that are writable. Thus, they can be used for many other use cases while still maintaining the integrity of staging Oracle VLDB image copy even when Oracle database is switched over. This provides tremendous storage cost saving by substantially reducing VLDB storage footprint. NetApp recommends to minimize FlexClone activities in the event of database switching over from primary storage to database image copy in order to maintain Oracle performance at high level.
- **EC2 compute instances.** In these tests and validations, we used an AWS EC2 t2.xlarge instance as the Oracle database compute instance. NetApp recommends using an M5 type EC2 instance as the compute instance for Oracle in production deployment because it is optimized for database workload. You need to size the EC2 instance appropriately for the number of vCPUs and the amount of RAM based on actual workload requirements.
- **FSx storage HA clusters single- or multi-zone deployment.** In these tests and validations, we deployed an FSx HA cluster in a single AWS availability zone. For production deployment, NetApp recommends deploying an FSx HA pair in two different availability zones. An FSx HA cluster is always provisioned in a HA pair that is sync mirrored in a pair of active-passive file systems to provide storage-level redundancy. Multi-zone deployment further enhances high availability in the event of failure in a single AWS zone.
- **FSx storage cluster sizing.** An Amazon FSx for ONTAP storage file system provides up to 160,000 raw SSD IOPS, up to 4GBps throughput, and a maximum of 192TiB capacity. However, you can size the cluster in terms of provisioned IOPS, throughput, and the storage limit (minimum 1,024 GiB) based on your actual requirements at the time of deployment. The capacity can be adjusted dynamically on the fly without affecting application availability.
- **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 FSx file systems deployment for a VLDB, dNFS multi-path to different FSx NFS file systems should be properly configured.

Solution deployment

It is assumed that you already have your Oracle VLDB deployed in AWS EC2 environment within a VPC. If you need help on Oracle deployment in AWS, please refer to following technical reports for help.

- [Oracle Database Deployment on EC2 and FSx Best Practices](#)
- [Oracle Database Deployment and Protection in AWS FSx/EC2 with iSCSI/ASM](#)
- [Oracle 19c in Standalone Restart on AWS FSx/EC2 with NFS/ASM](#)

Your Oracle VLDB can be running either on a FSx ONTAP or any other storage of choices within the AWS EC2

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 AWS FSx ONTAP storage.

Prerequisites for deployment

Deployment requires the following prerequisites.

1. An AWS account has been set up, and the necessary VPC and network segments have been created within your AWS account.
2. From the AWS EC2 console, you must deploy two EC2 Linux instances, one as the primary Oracle DB server and an optional alternative clone target DB server. See the architecture diagram in the previous section for more details about the environment setup. Also review the [User Guide for Linux instances](#) for more information.
3. From the AWS EC2 console, deploy Amazon FSx for ONTAP storage HA clusters to host the NFS volumes that stores the Oracle database standby image copy. If you are not familiar with the deployment of FSx storage, see the documentation [Creating FSx for ONTAP file systems](#) for step-by-step instructions.
4. Steps 2 and 3 can be performed using the following Terraform automation toolkit, which creates an EC2 instance named `ora_01` and an FSx file system named `fsx_01`. Review the instruction carefully and change the variables to suit your environment before execution. The template can be easily revised for your own deployment requirements.

```
git clone https://github.com/NetApp-
Automation/na_aws_fsx_ec2_deploy.git
```



Ensure that you have allocated at least 50G in EC2 instance root volume in order to have sufficient space to stage Oracle installation files.

Provision and export NFS volume to be mounted to EC2 DB instance host

In this demonstration, we will show how to provision an NFS volume from the command line by login to an FSx cluster via ssh as fsxadmin user through FSx cluster management IP. Alternatively, the volume can be allocated using the AWS FSx console as well. Repeat the procedures on other FSx file systems if more than one FSx file system are set up to accommodate the size of the database.


1. First, provision NFS volume via CLI by logging to the FSx cluster through SSH as the fsxadmin user. Change to your FSx cluster management IP address, which can be retrieved from AWS FSx ONTAP UI console.

```
ssh fsxadmin@172.30.15.53
```

2. Create NFS volume the same size as your primary storage for storing primary Oracle VLDB database data files image copy.

```
vol create -volume ora_01_copy -aggregate aggr1 -size 100G -state  
online -type RW -junction-path /ora_01_copy -snapshot-policy none  
-tiering-policy snapshot-only
```

3. Alternatively, the volume can be provisioned from AWS FSx console UI with options: storage efficiency Enabled, security style Unix , Snapshot policy None, and Storage tiering Snapshot Only as show below.



4. Create a customized snapshot policy for oracle database with a daily schedule and 30 days retention. You should adjust the policy to fit your specific needs in terms of snapshot frequency and retention window.

```
snapshot policy create -policy oracle -enabled true -schedule1 daily
-count1 30
```

Apply policy to provisioned NFS volume for RMAN incremental backup and merge.

```
vol modify -volume ora_01_copy -snapshot-policy oracle
```

5. Login to EC2 instance as ec2-user and create a directory /nfsfsxn. Create additional mount point directories for additional FSx file systems.

```
sudo mkdir /nfsfsxn
```

6. Mount the FSx ONTAP NFS volume to EC2 DB instance host. Change to your FSx virtual server NFS lif address. The NFS lif address can be retrieved from FSx ONTAP UI console.

```
sudo mount 172.30.15.19:/ora_01_copy /nfsfsxn -o  
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=262144,wsiz=262144,noi  
tr
```

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

```
sudo chown oracle:oinstall /nfsfsxn
```

Setup Oracle RMAN incremental merge to image copy on FSx

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 DB server EC2 instance as oracle user
2. Create an oracopy directory under mount point /nfsfsxn to store oracle data files image copies and archlog directory for Oracle flash recovery area.

```
mkdir /nfsfsxn/oracopy
```

```
mkdir /nfsfsxn/archlog
```

3. Login to Oracle database via sqlplus, enable block change tracking for faster incremental backup and change Oracle flash recovery area to FSxN 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 FSxN NFS mount for recovery.

```
sqlplus / as sysdba
```

From sqlplus prompt, execute following command.

```
alter database enable block change tracking using file  
'/nfsfsxn/oracopy/bct_db1.ctf'
```

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

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 '/nfsfsxn/oracopy/%U';  
  allocate channel c2 device type disk format '/nfsfsxn/oracopy/%U';  
  allocate channel c3 device type disk format '/nfsfsxn/oracopy/%U';  
  allocate channel c4 device type disk format '/nfsfsxn/oracopy/%U';  
  delete obsolete;  
  sql 'alter system archive log current';  
  recover copy of database with tag 'OraCopyBKUPonFSxN_level_0';  
  backup incremental level 1 copies=1 for recover of copy with tag  
'OraCopyBKUPonFSxN_level_0' database;  
}
```

5. At EC2 DB 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@ip-172-30-15-99 ~]$ rman target / nocatalog;
```

```
Recovery Manager: Release 19.0.0.0.0 - Production on Wed May 24  
17:44:49 2023  
Version 19.18.0.0.0
```

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

```
connected to target database: DB1 (DBID=1730530050)  
using target database control file instead of recovery catalog
```

```
RMAN>
```

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
   '/nfsfsxn/oracopy/%U';
4>  allocate channel c2 device type disk format
   '/nfsfsxn/oracopy/%U';
5>  allocate channel c3 device type disk format
   '/nfsfsxn/oracopy/%U';
6>  allocate channel c4 device type disk format
   '/nfsfsxn/oracopy/%U';
7>  delete obsolete;
8>  sql 'alter system archive log current';
9>  recover copy of database with tag 'OraCopyBKUPonFSxN_level_0';
10> backup incremental level 1 copies=1 for recover of copy with
tag 'OraCopyBKUPonFSxN_level_0' database;
11> }

allocated channel: c1
channel c1: SID=411 device type=DISK

allocated channel: c2
channel c2: SID=146 device type=DISK

allocated channel: c3
channel c3: SID=402 device type=DISK

allocated channel: c4
channel c4: SID=37 device type=DISK

Starting recover at 17-MAY-23
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 5 found to recover
no copy of datafile 6 found to recover
no copy of datafile 7 found to recover
.
.
Finished recover at 17-MAY-23

Starting backup at 17-MAY-23
channel c1: starting incremental level 1 datafile backup set
channel c1: specifying datafile(s) in backup set
input datafile file number=00022
name=+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.287.113
7018311

```

```
input datafile file number=00026
name=+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.291.113
7018481
input datafile file number=00030
name=+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.295.113
7018787
input datafile file number=00011
name=+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/undotbs1.27
1.1136668041
input datafile file number=00035
name=+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.300.113
7019181
channel c1: starting piece 1 at 17-MAY-23
channel c2: starting incremental level 1 datafile backup set
channel c2: specifying datafile(s) in backup set
input datafile file number=00023
name=+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.288.113
7018359
input datafile file number=00027
name=+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.292.113
7018523
input datafile file number=00031
name=+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.296.113
7018837
input datafile file number=00009
name=+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/system.272.
1136668041
input datafile file number=00034
name=+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.299.113
7019117
.
.
Finished backup at 17-MAY-23

Starting Control File and SPFILE Autobackup at 17-MAY-23
piece
handle=+LOGS/DB1/AUTOBACKUP/2023_05_17/s_1137095435.367.1137095435
comment=NONE
Finished Control File and SPFILE Autobackup at 17-MAY-23
released channel: c1
released channel: c2
released channel: c3
released channel: c4

RMAN> **end-of-file**
```

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

```

RMAN> list copy of database tag 'OraCopyBKUPonFSxN_level_0';

List of Datafile Copies
=====

Key          File S Completion Time Ckp SCN      Ckp Time      Sparse
-----
19           1    A 17-MAY-23      3009819      17-MAY-23     NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
      SYSTEM_FNO-1_0h1sd7ae
      Tag: ORACOPYBKUPONFSXN_LEVEL_0

20           3    A 17-MAY-23      3009826      17-MAY-23     NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
      SYSAUX_FNO-3_0i1sd7at
      Tag: ORACOPYBKUPONFSXN_LEVEL_0

21           4    A 17-MAY-23      3009830      17-MAY-23     NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
      UNDOTBS1_FNO-4_0j1sd7b4
      Tag: ORACOPYBKUPONFSXN_LEVEL_0

27           5    A 17-MAY-23      2383520      12-MAY-23     NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
      SYSTEM_FNO-5_0p1sd7cf
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 2, PDB Name: PDB$SEED

26           6    A 17-MAY-23      2383520      12-MAY-23     NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
      SYSAUX_FNO-6_0o1sd7c8
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 2, PDB Name: PDB$SEED

34           7    A 17-MAY-23      3009907      17-MAY-23     NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-
      7_101sd7dl
      Tag: ORACOPYBKUPONFSXN_LEVEL_0

33           8    A 17-MAY-23      2383520      12-MAY-23     NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
      UNDOTBS1_FNO-8_0v1sd7di
      Tag: ORACOPYBKUPONFSXN_LEVEL_0

```

```

Container ID: 2, PDB Name: PDB$SEED

28      9      A 17-MAY-23      3009871      17-MAY-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
SYSTEM_FNO-9_0q1sd7cm
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 3, PDB Name: DB1_PDB1

22      10     A 17-MAY-23      3009849      17-MAY-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
SYSAUX_FNO-10_0k1sd7bb
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 3, PDB Name: DB1_PDB1

25      11     A 17-MAY-23      3009862      17-MAY-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
UNDOTBS1_FNO-11_0n1sd7c1
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 3, PDB Name: DB1_PDB1

35      12     A 17-MAY-23      3009909      17-MAY-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-
12_111sd7dm
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 3, PDB Name: DB1_PDB1

29      13     A 17-MAY-23      3009876      17-MAY-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
SYSTEM_FNO-13_0r1sd7ct
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 4, PDB Name: DB1_PDB2

23      14     A 17-MAY-23      3009854      17-MAY-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
SYSAUX_FNO-14_0l1sd7bi
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 4, PDB Name: DB1_PDB2

31      15     A 17-MAY-23      3009900      17-MAY-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
UNDOTBS1_FNO-15_0t1sd7db
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 4, PDB Name: DB1_PDB2

36      16     A 17-MAY-23      3009911      17-MAY-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-

```



```

16_121sd7dn
    Tag: ORACOPYBKUPONFSXN_LEVEL_0
    Container ID: 4, PDB Name: DB1_PDB2

30      17      A 17-MAY-23      3009895      17-MAY-23      NO
    Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
SYSTEM_FNO-17_0s1sd7d4
    Tag: ORACOPYBKUPONFSXN_LEVEL_0
    Container ID: 5, PDB Name: DB1_PDB3

24      18      A 17-MAY-23      3009858      17-MAY-23      NO
    Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
SYSAUX_FNO-18_0m1sd7bq
    Tag: ORACOPYBKUPONFSXN_LEVEL_0
    Container ID: 5, PDB Name: DB1_PDB3

32      19      A 17-MAY-23      3009903      17-MAY-23      NO
    Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
UNDOTBS1_FNO-19_0u1sd7de
    Tag: ORACOPYBKUPONFSXN_LEVEL_0
    Container ID: 5, PDB Name: DB1_PDB3

37      20      A 17-MAY-23      3009914      17-MAY-23      NO
    Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-
20_131sd7do
    Tag: ORACOPYBKUPONFSXN_LEVEL_0
    Container ID: 5, PDB Name: DB1_PDB3

4       21      A 17-MAY-23      3009019      17-MAY-23      NO
    Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
21_021sd6pv
    Tag: ORACOPYBKUPONFSXN_LEVEL_0
    Container ID: 3, PDB Name: DB1_PDB1

5       22      A 17-MAY-23      3009419      17-MAY-23      NO
    Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
22_031sd6r2
    Tag: ORACOPYBKUPONFSXN_LEVEL_0
    Container ID: 3, PDB Name: DB1_PDB1

6       23      A 17-MAY-23      3009460      17-MAY-23      NO
    Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
23_041sd6s5
    Tag: ORACOPYBKUPONFSXN_LEVEL_0
    Container ID: 3, PDB Name: DB1_PDB1

```

7	24	A	17-MAY-23	3009473	17-MAY-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-						
24_051sd6t9						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
8	25	A	17-MAY-23	3009502	17-MAY-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-						
25_061sd6uc						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
9	26	A	17-MAY-23	3009548	17-MAY-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-						
26_071sd6vf						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
10	27	A	17-MAY-23	3009576	17-MAY-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-						
27_081sd70i						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
11	28	A	17-MAY-23	3009590	17-MAY-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-						
28_091sd71l						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
12	29	A	17-MAY-23	3009619	17-MAY-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-						
29_0a1sd72o						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
13	30	A	17-MAY-23	3009648	17-MAY-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-						
30_0b1sd73r						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
14	31	A	17-MAY-23	3009671	17-MAY-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-						
31_0c1sd74u						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						

```

Container ID: 3, PDB Name: DB1_PDB1

15      32      A 17-MAY-23      3009729      17-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
32_0d1sd762
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 3, PDB Name: DB1_PDB1

16      33      A 17-MAY-23      3009743      17-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
33_0e1sd775
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 3, PDB Name: DB1_PDB1

17      34      A 17-MAY-23      3009771      17-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
34_0f1sd788
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 3, PDB Name: DB1_PDB1

18      35      A 17-MAY-23      3009805      17-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
35_0g1sd79b
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 3, PDB Name: DB1_PDB1

```

RMAN>

- Report schema from Oracle RMAN command prompt to observe that current active database data files are in primary storage ASM +DATA disk group.

```

RMAN> report schema;

Report of database schema for database with db_unique_name DB1

List of Permanent Datafiles
=====
File Size(MB) Tablespace          RB segs Datafile Name
-----
1      1060      SYSTEM          YES
+DATA/DB1/DATAFILE/system.257.1136666315
3      810       SYSAUX          NO
+DATA/DB1/DATAFILE/sysaux.258.1136666361
4      675       UNDOTBS1        YES
+DATA/DB1/DATAFILE/undotbs1.259.1136666385

```

```

5      400      PDB$SEED:SYSTEM      NO
+DATA/DB1/86B637B62FE07A65E053F706E80A27CA/DATAFILE/system.266.11366
67165
6      460      PDB$SEED:SYSAUX      NO
+DATA/DB1/86B637B62FE07A65E053F706E80A27CA/DATAFILE/sysaux.267.11366
67165
7      5        USERS      NO
+DATA/DB1/DATAFILE/users.260.1136666387
8      230      PDB$SEED:UNDOTBS1    NO
+DATA/DB1/86B637B62FE07A65E053F706E80A27CA/DATAFILE/undotbs1.268.113
6667165
9      400      DB1_PDB1:SYSTEM      YES
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/system.272.11366
68041
10     490      DB1_PDB1:SYSAUX      NO
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/sysaux.273.11366
68041
11     465      DB1_PDB1:UNDOTBS1    YES
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/undotbs1.271.113
6668041
12     5        DB1_PDB1:USERS      NO
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/users.275.113666
8057
13     400      DB1_PDB2:SYSTEM      YES
+DATA/DB1/FB867EA89ECF81C0E053630F1EACB901/DATAFILE/system.277.11366
68057
14     470      DB1_PDB2:SYSAUX      NO
+DATA/DB1/FB867EA89ECF81C0E053630F1EACB901/DATAFILE/sysaux.278.11366
68057
15     235      DB1_PDB2:UNDOTBS1    YES
+DATA/DB1/FB867EA89ECF81C0E053630F1EACB901/DATAFILE/undotbs1.276.113
6668057
16     5        DB1_PDB2:USERS      NO
+DATA/DB1/FB867EA89ECF81C0E053630F1EACB901/DATAFILE/users.280.113666
8071
17     400      DB1_PDB3:SYSTEM      YES
+DATA/DB1/FB867F8A4D4F821CE053630F1EAC69CC/DATAFILE/system.282.11366
68073
18     470      DB1_PDB3:SYSAUX      NO
+DATA/DB1/FB867F8A4D4F821CE053630F1EAC69CC/DATAFILE/sysaux.283.11366
68073
19     235      DB1_PDB3:UNDOTBS1    YES
+DATA/DB1/FB867F8A4D4F821CE053630F1EAC69CC/DATAFILE/undotbs1.281.113
6668073
20     5        DB1_PDB3:USERS      NO
+DATA/DB1/FB867F8A4D4F821CE053630F1EAC69CC/DATAFILE/users.285.113666

```

8087

21 4096 DB1_PDB1:SOE NO
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.286.11370182
39

22 4096 DB1_PDB1:SOE NO
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.287.11370183
11

23 4096 DB1_PDB1:SOE NO
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.288.11370183
59

24 4096 DB1_PDB1:SOE NO
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.289.11370184
05

25 4096 DB1_PDB1:SOE NO
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.290.11370184
43

26 4096 DB1_PDB1:SOE NO
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.291.11370184
81

27 4096 DB1_PDB1:SOE NO
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.292.11370185
23

28 4096 DB1_PDB1:SOE NO
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.293.11370187
07

29 4096 DB1_PDB1:SOE NO
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.294.11370187
45

30 4096 DB1_PDB1:SOE NO
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.295.11370187
87

31 4096 DB1_PDB1:SOE NO
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.296.11370188
37

32 4096 DB1_PDB1:SOE NO
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.297.11370189
35

33 4096 DB1_PDB1:SOE NO
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.298.11370190
77

34 4096 DB1_PDB1:SOE NO
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.299.11370191
17

35 4096 DB1_PDB1:SOE NO
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.300.11370191
81

List of Temporary Files

=====

File	Size (MB)	Tablespace	Maxsize (MB)	Tempfile Name
1	123	TEMP	32767	+DATA/DB1/TEMPFILE/temp.265.1136666447
2	123	PDB\$SEED:TEMP	32767	+DATA/DB1/FB864A929AEB79B9E053630F1EAC7046/TEMPFILE/temp.269.1136667185
3	10240	DB1_PDB1:TEMP	32767	+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/TEMPFILE/temp.274.1136668051
4	123	DB1_PDB2:TEMP	32767	+DATA/DB1/FB867EA89ECF81C0E053630F1EACB901/TEMPFILE/temp.279.1136668067
5	123	DB1_PDB3:TEMP	32767	+DATA/DB1/FB867F8A4D4F821CE053630F1EAC69CC/TEMPFILE/temp.284.1136668081

RMAN>

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

```
[oracle@ip-172-30-15-99 ~]$ ls -l /nfsfsxn/oracopy/
total 70585148
-rw-r----- 1 oracle asm 4294975488 May 17 18:09 data_D-DB1_I-
1730530050_TS-SOE_FNO-21_021sd6pv
-rw-r----- 1 oracle asm 4294975488 May 17 18:10 data_D-DB1_I-
1730530050_TS-SOE_FNO-22_031sd6r2
-rw-r----- 1 oracle asm 4294975488 May 17 18:10 data_D-DB1_I-
1730530050_TS-SOE_FNO-23_041sd6s5
-rw-r----- 1 oracle asm 4294975488 May 17 18:11 data_D-DB1_I-
1730530050_TS-SOE_FNO-24_051sd6t9
-rw-r----- 1 oracle asm 4294975488 May 17 18:11 data_D-DB1_I-
1730530050_TS-SOE_FNO-25_061sd6uc
-rw-r----- 1 oracle asm 4294975488 May 17 18:12 data_D-DB1_I-
1730530050_TS-SOE_FNO-26_071sd6vf
-rw-r----- 1 oracle asm 4294975488 May 17 18:13 data_D-DB1_I-
1730530050_TS-SOE_FNO-27_081sd70i
-rw-r----- 1 oracle asm 4294975488 May 17 18:13 data_D-DB1_I-
1730530050_TS-SOE_FNO-28_091sd71l
-rw-r----- 1 oracle asm 4294975488 May 17 18:14 data_D-DB1_I-
1730530050_TS-SOE_FNO-29_0a1sd72o
-rw-r----- 1 oracle asm 4294975488 May 17 18:14 data_D-DB1_I-
```

```
1730530050_TS-SOE_FNO-30_0b1sd73r
-rw-r----- 1 oracle asm 4294975488 May 17 18:15 data_D-DB1_I-
1730530050_TS-SOE_FNO-31_0c1sd74u
-rw-r----- 1 oracle asm 4294975488 May 17 18:16 data_D-DB1_I-
1730530050_TS-SOE_FNO-32_0d1sd762
-rw-r----- 1 oracle asm 4294975488 May 17 18:16 data_D-DB1_I-
1730530050_TS-SOE_FNO-33_0e1sd775
-rw-r----- 1 oracle asm 4294975488 May 17 18:17 data_D-DB1_I-
1730530050_TS-SOE_FNO-34_0f1sd788
-rw-r----- 1 oracle asm 4294975488 May 17 18:17 data_D-DB1_I-
1730530050_TS-SOE_FNO-35_0g1sd79b
-rw-r----- 1 oracle asm 513810432 May 17 18:18 data_D-DB1_I-
1730530050_TS-SYSAUX_FNO-10_0k1sd7bb
-rw-r----- 1 oracle asm 492838912 May 17 18:18 data_D-DB1_I-
1730530050_TS-SYSAUX_FNO-14_0l1sd7bi
-rw-r----- 1 oracle asm 492838912 May 17 18:18 data_D-DB1_I-
1730530050_TS-SYSAUX_FNO-18_0m1sd7bq
-rw-r----- 1 oracle asm 849354752 May 17 18:18 data_D-DB1_I-
1730530050_TS-SYSAUX_FNO-3_0i1sd7at
-rw-r----- 1 oracle asm 482353152 May 17 18:18 data_D-DB1_I-
1730530050_TS-SYSAUX_FNO-6_0o1sd7c8
-rw-r----- 1 oracle asm 1111498752 May 17 18:18 data_D-DB1_I-
1730530050_TS-SYSTEM_FNO-1_0h1sd7ae
-rw-r----- 1 oracle asm 419438592 May 17 18:19 data_D-DB1_I-
1730530050_TS-SYSTEM_FNO-13_0r1sd7ct
-rw-r----- 1 oracle asm 419438592 May 17 18:19 data_D-DB1_I-
1730530050_TS-SYSTEM_FNO-17_0s1sd7d4
-rw-r----- 1 oracle asm 419438592 May 17 18:19 data_D-DB1_I-
1730530050_TS-SYSTEM_FNO-5_0p1sd7cf
-rw-r----- 1 oracle asm 419438592 May 17 18:19 data_D-DB1_I-
1730530050_TS-SYSTEM_FNO-9_0q1sd7cm
-rw-r----- 1 oracle asm 487596032 May 17 18:18 data_D-DB1_I-
1730530050_TS-UNDOTBS1_FNO-11_0n1sd7c1
-rw-r----- 1 oracle asm 246423552 May 17 18:19 data_D-DB1_I-
1730530050_TS-UNDOTBS1_FNO-15_0t1sd7db
-rw-r----- 1 oracle asm 246423552 May 17 18:19 data_D-DB1_I-
1730530050_TS-UNDOTBS1_FNO-19_0u1sd7de
-rw-r----- 1 oracle asm 707796992 May 17 18:18 data_D-DB1_I-
1730530050_TS-UNDOTBS1_FNO-4_0j1sd7b4
-rw-r----- 1 oracle asm 241180672 May 17 18:19 data_D-DB1_I-
1730530050_TS-UNDOTBS1_FNO-8_0v1sd7di
-rw-r----- 1 oracle asm 5251072 May 17 18:19 data_D-DB1_I-
1730530050_TS-USERS_FNO-12_1l1sd7dm
-rw-r----- 1 oracle asm 5251072 May 17 18:19 data_D-DB1_I-
1730530050_TS-USERS_FNO-16_1t1sd7dn
-rw-r----- 1 oracle asm 5251072 May 17 18:19 data_D-DB1_I-
```

```
1730530050_TS-USERS_FNO-20_131sd7do  
-rw-r----- 1 oracle asm      5251072 May 17 18:19 data_D-DB1_I-  
1730530050_TS-USERS_FNO-7_101sd7d1
```

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

Switch Oracle DB 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 FSx ONTAP 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 database host instance is intact and database control file, archived and current logs are all available for recovery.

1. Login to EC2 DB server host as oracle user and create a test table before switch over.

```
[ec2-user@ip-172-30-15-99 ~]$ sudo su
[root@ip-172-30-15-99 ec2-user]# su - oracle
Last login: Thu May 18 14:22:34 UTC 2023
[oracle@ip-172-30-15-99 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Thu May 18 14:30:36
2023
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  DB1_PDB1                                READ WRITE NO
          4  DB1_PDB2                                READ WRITE NO
          5  DB1_PDB3                                READ WRITE NO

SQL> alter session set container=db1_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
18-MAY-23 02.35.37.000000 PM
test oracle incremental merge switch to copy

SQL>
```

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 1.2885E+10 bytes
Fixed Size                  9177880 bytes
Variable Size               1778384896 bytes
Database Buffers           1.1073E+10 bytes
Redo Buffers                24375296 bytes
Database mounted.
SQL>
```

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

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

```
datafile 1 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-1_0h1sd7ae"
datafile 3 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-3_0i1sd7at"
```

datafile 4 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-4_0j1sd7b4"
datafile 5 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-5_0p1sd7cf"
datafile 6 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-6_0o1sd7c8"
datafile 7 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-7_101sd7d1"
datafile 8 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-8_0v1sd7di"
datafile 9 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-9_0q1sd7cm"
datafile 10 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-10_0k1sd7bb"
datafile 11 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-11_0n1sd7c1"
datafile 12 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-12_111sd7dm"
datafile 13 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-13_0r1sd7ct"
datafile 14 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-14_0l1sd7bi"
datafile 15 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-15_0t1sd7db"
datafile 16 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-16_121sd7dn"
datafile 17 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-17_0s1sd7d4"
datafile 18 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-18_0m1sd7bq"
datafile 19 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-19_0u1sd7de"
datafile 20 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-20_131sd7do"
datafile 21 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-21_021sd6pv"
datafile 22 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-22_031sd6r2"
datafile 23 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-23_041sd6s5"
datafile 24 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-24_051sd6t9"
datafile 25 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-25_061sd6uc"
datafile 26 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-26_071sd6vf"

```
datafile 27 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-27_081sd70i"
datafile 28 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-28_091sd711"
datafile 29 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-29_0a1sd72o"
datafile 30 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-30_0b1sd73r"
datafile 31 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-31_0c1sd74u"
datafile 32 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-32_0d1sd762"
datafile 33 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-33_0e1sd775"
datafile 34 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-34_0f1sd788"
datafile 35 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-35_0g1sd79b"
```

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

```
RMAN> recover database;

Starting recover at 18-MAY-23
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: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-9_0q1sd7cm
destination for restore of datafile 00023: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-23_041sd6s5
destination for restore of datafile 00027: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-27_081sd70i
destination for restore of datafile 00031: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-31_0c1sd74u
destination for restore of datafile 00034: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-34_0f1sd788
channel ORA_DISK_1: reading from backup piece
/nfsfsxn/oracopy/321sfous_98_1_1
channel ORA_DISK_1: piece handle=/nfsfsxn/oracopy/321sfous_98_1_1
tag=ORACOPYBKUPONFSXN_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: /nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SYSAUX_FNO-10_0k1sd7bb
destination for restore of datafile 00021: /nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SOE_FNO-21_021sd6pv
destination for restore of datafile 00025: /nfsfsxn/oracopy/data_D-
DB1_I-1730530050_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: /nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-USERS_FNO-16_121sd7dn
channel ORA_DISK_1: reading from backup piece
/nfsfsxn/oracopy/3i1sfov0_114_1_1
channel ORA_DISK_1: piece handle=/nfsfsxn/oracopy/3i1sfov0_114_1_1
tag=ORACOPYBKUPONFSXN_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: /nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-USERS_FNO-20_131sd7do
channel ORA_DISK_1: reading from backup piece
/nfsfsxn/oracopy/3j1sfov0_115_1_1
channel ORA_DISK_1: piece handle=/nfsfsxn/oracopy/3j1sfov0_115_1_1
tag=ORACOPYBKUPONFSXN_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 18-MAY-23

RMAN> alter database open;

Statement processed

RMAN>
```

5. Check database structure from sqlplus after recovery to observe that all database data files with exception of control, temp, and current log files are now switched over to copy on FSx ONTAP 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;
```

NAME

```
-----
-----
+DATA/DB1/CONTROLFILE/current.261.1136666435
+DATA/DB1/FB864A929AEB79B9E053630F1EAC7046/TEMPFILE/temp.269.1136667
185
+DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/TEMPFILE/temp.274.1136668
051
+DATA/DB1/FB867EA89ECF81C0E053630F1EACB901/TEMPFILE/temp.279.1136668
067
+DATA/DB1/FB867F8A4D4F821CE053630F1EAC69CC/TEMPFILE/temp.284.1136668
081
+DATA/DB1/ONLINELOG/group_1.262.1136666437
+DATA/DB1/ONLINELOG/group_2.263.1136666437
+DATA/DB1/ONLINELOG/group_3.264.1136666437
+DATA/DB1/TEMPFILE/temp.265.1136666447
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-21_021sd6pv
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-22_031sd6r2
```

NAME

```
-----
-----
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-23_041sd6s5
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-24_051sd6t9
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-25_061sd6uc
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-26_071sd6vf
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-27_081sd70i
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-28_091sd71l
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-29_0a1sd72o
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-30_0b1sd73r
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-31_0c1sd74u
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-32_0d1sd762
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-33_0e1sd775
```

NAME

```
-----  
-----  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-34_0f1sd788  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-35_0g1sd79b  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-10_0k1sd7bb  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-14_0l1sd7bi  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-18_0m1sd7bq  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-3_0i1sd7at  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-6_0o1sd7c8  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-13_0r1sd7ct  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-17_0s1sd7d4  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-1_0h1sd7ae  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-5_0p1sd7cf
```

NAME

```
-----  
-----  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-9_0q1sd7cm  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-11_0n1sd7c1  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-15_0t1sd7db  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-19_0u1sd7de  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-4_0j1sd7b4  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-8_0v1sd7di  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-12_1l1sd7dm  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-16_121sd7dn  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-20_131sd7do  
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-7_101sd7dl
```

43 rows selected.

SQL>

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

```

SQL> show pdbs

  CON_ID  CON_NAME                                OPEN MODE  RESTRICTED
-----  -
      2  PDB$SEED                                READ ONLY  NO
      3  DB1_PDB1                                READ WRITE NO
      4  DB1_PDB2                                READ WRITE NO
      5  DB1_PDB3                                READ WRITE NO
SQL> alter session set container=db1_pdb1;

Session altered.

SQL> select * from test;

      ID
-----
DT
-----
EVENT
-----
      1
18-MAY-23 02.35.37.000000 PM
test oracle incremental merge switch to copy

SQL>

```

7. You could run the Oracle database in FSx NFS mount for an extended period without a performance penalty because FSx ONTAP is redundant production-grade storage that delivers high performance. When the primary storage issue is fixed, you can swing back to it by reversing the incremental backup merge processes with minimal downtime.

Oracle DB recovery from image copy to different EC2 DB instance host

In a failure when both primary storage and EC2 DB instance host are lost, the recovery can not be conducted from the original server. Fortunately, you still have an Oracle database backup image copy on the redundant FSxN NFS file system. You could quickly provision another identical EC2 DB instance and easily mount the image copy of your VLDB to the new EC2 DB host via NFS to run recovery. In this section, we will demonstrate the step-by-step procedures for doing so.

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

```
[oracle@ip-172-30-15-99 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Tue May 30 17:21:05
2023
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 DB1_PDB1                                READ WRITE NO
          4 DB1_PDB2                                READ WRITE NO
          5 DB1_PDB3                                READ WRITE NO

SQL> alter session set container=db1_pdb1;

Session altered.

SQL> insert into test values(2, sysdate, 'test recovery on a new EC2
instance host with image copy on FSxN');

1 row created.

SQL> commit;

Commit complete.

SQL> select * from test;
```

```

          ID
-----
DT
-----
EVENT
-----
          1
18-MAY-23 02.35.37.000000 PM
test oracle incremental merge switch to copy

          2
30-MAY-23 05.23.11.000000 PM
test recovery on a new EC2 instance host with image copy on FSxN

SQL>

```

2. As oracle user, run RMAN incremental backup and merge to flush the transaction to backup set on FSxN 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: DB1 (DBID=1730530050)
using target database control file instead of recovery catalog

RMAN> @rman_bkup_merge.cmd

```

3. Shutdown primary EC2 DB instance host to simulate a total failure of storage and DB server host.
4. Provision a new EC2 DB instance host ora_02 with same OS and version via AWS EC2 console. Configure OS kernel with same patches as primary EC2 DB server host, Oracle preinstall RPM, and add swap space to the host as well. Install same version and patches of Oracle as in primary EC2 DB server host with software only option. These tasks can be automated with NetApp automation toolkit as available from below links.

Toolkit: [na_oracle19c_deploy](#)

Documentation: [Automated Deployment of Oracle19c for ONTAP on NFS](#)

5. Configure oracle environment similiarly to primary EC2 DB instance host ora_01, such as oratab, oralnst.loc, and oracle user .bash_profile. It is a good practice to backup those files to FSxN NFS mount point.
6. The Oracle database backup image copy on FSxN NFS mount is stored on a FSx cluster that spans AWS availability zones for redundancy, high avilability, and high performance. The NFS file system can be easily mounted to a new server as far as the networking is reachable. The following procedures mount the image copy of an Oracle VLDB backup to newly provsioned EC2 DB instance host for recovery.

As ec2-user, create the mount point.

```
sudo mkdir /nfsfsxn
```

As ec2-user, mount the NFS volume that stored Oracle VLDB backup image copy.

```
sudo mount 172.30.15.19:/ora_01_copy /nfsfsxn -o
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=262144,wsiz=262144,noi
tr
```

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

```
[ec2-user@ip-172-30-15-124 ~]$ ls -ltr /nfsfsxn/oracopy
total 78940700
-rw-r-----. 1 oracle 54331 482353152 May 26 18:45 data_D-DB1_I-
1730530050_TS-SYSAUX_FNO-6_4m1t508t
-rw-r-----. 1 oracle 54331 419438592 May 26 18:45 data_D-DB1_I-
1730530050_TS-SYSTEM_FNO-5_4q1t509n
-rw-r-----. 1 oracle 54331 241180672 May 26 18:45 data_D-DB1_I-
1730530050_TS-UNDOTBS1_FNO-8_4t1t50a6
-rw-r-----. 1 oracle 54331 450560 May 30 15:29 6b1tf6b8_203_1_1
-rw-r-----. 1 oracle 54331 663552 May 30 15:29 6c1tf6b8_204_1_1
-rw-r-----. 1 oracle 54331 122880 May 30 15:29 6d1tf6b8_205_1_1
-rw-r-----. 1 oracle 54331 507904 May 30 15:29 6e1tf6b8_206_1_1
-rw-r-----. 1 oracle 54331 4259840 May 30 15:29 6f1tf6b9_207_1_1
-rw-r-----. 1 oracle 54331 9060352 May 30 15:29 6h1tf6b9_209_1_1
-rw-r-----. 1 oracle 54331 442368 May 30 15:29 6i1tf6b9_210_1_1
-rw-r-----. 1 oracle 54331 475136 May 30 15:29 6j1tf6bb_211_1_1
-rw-r-----. 1 oracle 54331 48660480 May 30 15:29 6g1tf6b9_208_1_1
-rw-r-----. 1 oracle 54331 589824 May 30 15:29 6l1tf6bb_213_1_1
-rw-r-----. 1 oracle 54331 606208 May 30 15:29 6m1tf6bb_214_1_1
-rw-r-----. 1 oracle 54331 368640 May 30 15:29 6o1tf6bb_216_1_1
-rw-r-----. 1 oracle 54331 368640 May 30 15:29 6p1tf6bc_217_1_1
-rw-r-----. 1 oracle 54331 57344 May 30 15:29 6r1tf6bc_219_1_1
-rw-r-----. 1 oracle 54331 57344 May 30 15:29 6s1tf6bc_220_1_1
-rw-r-----. 1 oracle 54331 57344 May 30 15:29 6t1tf6bc_221_1_1
```

```
-rw-r-----. 1 oracle 54331 4294975488 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SOE_FNO-23_3q1t4ut3  
-rw-r-----. 1 oracle 54331 4294975488 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SOE_FNO-21_3o1t4ut2  
-rw-r-----. 1 oracle 54331 4294975488 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SOE_FNO-27_461t4vt7  
-rw-r-----. 1 oracle 54331 4294975488 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SOE_FNO-25_3s1t4v1a  
-rw-r-----. 1 oracle 54331 4294975488 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SOE_FNO-22_3p1t4ut3  
-rw-r-----. 1 oracle 54331 4294975488 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SOE_FNO-31_4a1t5015  
-rw-r-----. 1 oracle 54331 4294975488 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SOE_FNO-29_481t4vt7  
-rw-r-----. 1 oracle 54331 4294975488 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SOE_FNO-34_4d1t5058  
-rw-r-----. 1 oracle 54331 4294975488 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SOE_FNO-26_451t4vt7  
-rw-r-----. 1 oracle 54331 4294975488 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SOE_FNO-24_3r1t4ut3  
-rw-r-----. 1 oracle 54331 555753472 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SYSAUX_FNO-10_4i1t5083  
-rw-r-----. 1 oracle 54331 429924352 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SYSTEM_FNO-9_4n1t509m  
-rw-r-----. 1 oracle 54331 4294975488 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SOE_FNO-30_491t5014  
-rw-r-----. 1 oracle 54331 4294975488 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SOE_FNO-28_471t4vt7  
-rw-r-----. 1 oracle 54331 4294975488 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SOE_FNO-35_4e1t5059  
-rw-r-----. 1 oracle 54331 4294975488 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SOE_FNO-32_4b1t501u  
-rw-r-----. 1 oracle 54331 487596032 May 30 17:26 data_D-DB1_I-  
1730530050_TS-UNDOTBS1_FNO-11_411t508t  
-rw-r-----. 1 oracle 54331 4294975488 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SOE_FNO-33_4c1t501v  
-rw-r-----. 1 oracle 54331 5251072 May 30 17:26 data_D-DB1_I-  
1730530050_TS-USERS_FNO-12_4v1t50aa  
-rw-r-----. 1 oracle 54331 1121984512 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SYSTEM_FNO-1_4f1t506m  
-rw-r-----. 1 oracle 54331 707796992 May 30 17:26 data_D-DB1_I-  
1730530050_TS-UNDOTBS1_FNO-4_4h1t5083  
-rw-r-----. 1 oracle 54331 534781952 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SYSAUX_FNO-14_4j1t508s  
-rw-r-----. 1 oracle 54331 429924352 May 30 17:26 data_D-DB1_I-  
1730530050_TS-SYSTEM_FNO-13_4o1t509m
```

```

-rw-r-----. 1 oracle 54331 429924352 May 30 17:26 data_D-DB1_I-
1730530050_TS-SYSTEM_FNO-17_4p1t509m
-rw-r-----. 1 oracle 54331 534781952 May 30 17:26 data_D-DB1_I-
1730530050_TS-SYSAUX_FNO-18_4k1t508t
-rw-r-----. 1 oracle 54331 1027612672 May 30 17:26 data_D-DB1_I-
1730530050_TS-SYSAUX_FNO-3_4g1t506m
-rw-r-----. 1 oracle 54331 5251072 May 30 17:26 data_D-DB1_I-
1730530050_TS-USERS_FNO-7_4u1t50a6
-rw-r-----. 1 oracle 54331 246423552 May 30 17:26 data_D-DB1_I-
1730530050_TS-UNDOTBS1_FNO-15_4r1t50a6
-rw-r-----. 1 oracle 54331 5251072 May 30 17:26 data_D-DB1_I-
1730530050_TS-USERS_FNO-16_501t50ad
-rw-r-----. 1 oracle 54331 246423552 May 30 17:26 data_D-DB1_I-
1730530050_TS-UNDOTBS1_FNO-19_4s1t50a6
-rw-r-----. 1 oracle 54331 5251072 May 30 17:26 data_D-DB1_I-
1730530050_TS-USERS_FNO-20_511t50ad
-rw-r-----. 1 oracle 54331 2318712832 May 30 17:32 721tfd6b_226_1_1
-rw-r-----. 1 oracle 54331 1813143552 May 30 17:33 701tfd6a_224_1_1
-rw-r-----. 1 oracle 54331 966656 May 30 17:33 731tfdic_227_1_1
-rw-r-----. 1 oracle 54331 5980160 May 30 17:33 751tfdij_229_1_1
-rw-r-----. 1 oracle 54331 458752 May 30 17:33 761tfdin_230_1_1
-rw-r-----. 1 oracle 54331 458752 May 30 17:33 771tfdiq_231_1_1
-rw-r-----. 1 oracle 54331 11091968 May 30 17:33 741tfdij_228_1_1
-rw-r-----. 1 oracle 54331 401408 May 30 17:33 791tfdit_233_1_1
-rw-r-----. 1 oracle 54331 2070708224 May 30 17:33 6v1tfd6a_223_1_1
-rw-r-----. 1 oracle 54331 376832 May 30 17:33 7a1tfdit_234_1_1
-rw-r-----. 1 oracle 54331 1874903040 May 30 17:33 711tfd6b_225_1_1
-rw-r-----. 1 oracle 54331 303104 May 30 17:33 7c1tfdiu_236_1_1
-rw-r-----. 1 oracle 54331 319488 May 30 17:33 7d1tfdi_237_1_1
-rw-r-----. 1 oracle 54331 57344 May 30 17:33 7f1tfdi_239_1_1
-rw-r-----. 1 oracle 54331 57344 May 30 17:33 7g1tfdi_240_1_1
-rw-r-----. 1 oracle 54331 57344 May 30 17:33 7h1tfdi_241_1_1
-rw-r--r--. 1 oracle 54331 12720 May 30 17:33 db1_ctl.sql
-rw-r-----. 1 oracle 54331 11600384 May 30 17:54 bct_db1.ctf

```

8. Verify the available Oracle archived logs on the FSxN NFS mount for recovery and note the last log file log sequence number. In this case, it is 175. Our recovery point is up to log sequence number 176.

```

[ec2-user@ip-172-30-15-124 ~]$ ls -ltr
/nfsfsxn/archlog/DB1/archivelog/2023_05_30
total 5714400
-r--r-----. 1 oracle 54331 321024 May 30 14:59
o1_mf_1_140__003t9mvn_.arc
-r--r-----. 1 oracle 54331 48996352 May 30 15:29
o1_mf_1_141__01t9qf6r_.arc
-r--r-----. 1 oracle 54331 167477248 May 30 15:44

```

```
o1_mf_1_142__02n3x2qb_.arc
-r--r----- . 1 oracle 54331 165684736 May 30 15:46
o1_mf_1_143__02rotwyb_.arc
-r--r----- . 1 oracle 54331 165636608 May 30 15:49
o1_mf_1_144__02x563wh_.arc
-r--r----- . 1 oracle 54331 168408064 May 30 15:51
o1_mf_1_145__031kg2co_.arc
-r--r----- . 1 oracle 54331 169446400 May 30 15:54
o1_mf_1_146__035xpcdt_.arc
-r--r----- . 1 oracle 54331 167595520 May 30 15:56
o1_mf_1_147__03bds8qf_.arc
-r--r----- . 1 oracle 54331 169270272 May 30 15:59
o1_mf_1_148__03gyt7rx_.arc
-r--r----- . 1 oracle 54331 170712576 May 30 16:01
o1_mf_1_149__03mfxl7v_.arc
-r--r----- . 1 oracle 54331 170744832 May 30 16:04
o1_mf_1_150__03qzz0ty_.arc
-r--r----- . 1 oracle 54331 169380864 May 30 16:06
o1_mf_1_151__03wgxdry_.arc
-r--r----- . 1 oracle 54331 169833984 May 30 16:09
o1_mf_1_152__040y85v3_.arc
-r--r----- . 1 oracle 54331 165134336 May 30 16:20
o1_mf_1_153__04ox946w_.arc
-r--r----- . 1 oracle 54331 169929216 May 30 16:22
o1_mf_1_154__04rbv7n8_.arc
-r--r----- . 1 oracle 54331 171903488 May 30 16:23
o1_mf_1_155__04tvlyvn_.arc
-r--r----- . 1 oracle 54331 179061248 May 30 16:25
o1_mf_1_156__04xgfjtl_.arc
-r--r----- . 1 oracle 54331 173593088 May 30 16:26
o1_mf_1_157__04zyg8hw_.arc
-r--r----- . 1 oracle 54331 175999488 May 30 16:27
o1_mf_1_158__052gp9mt_.arc
-r--r----- . 1 oracle 54331 179092992 May 30 16:29
o1_mf_1_159__0551wk7s_.arc
-r--r----- . 1 oracle 54331 175524352 May 30 16:30
o1_mf_1_160__057l46my_.arc
-r--r----- . 1 oracle 54331 173949440 May 30 16:32
o1_mf_1_161__05b2dmwp_.arc
-r--r----- . 1 oracle 54331 184166912 May 30 16:33
o1_mf_1_162__05drbj8n_.arc
-r--r----- . 1 oracle 54331 173026816 May 30 16:35
o1_mf_1_163__05h8lm1h_.arc
-r--r----- . 1 oracle 54331 174286336 May 30 16:36
o1_mf_1_164__05krsqmh_.arc
-r--r----- . 1 oracle 54331 166092288 May 30 16:37
```

```

o1_mf_1_165__05n378pw_.arc
-r--r-----. 1 oracle 54331 177640960 May 30 16:39
o1_mf_1_166__05pmg74l_.arc
-r--r-----. 1 oracle 54331 173972992 May 30 16:40
o1_mf_1_167__05s3o01r_.arc
-r--r-----. 1 oracle 54331 178474496 May 30 16:41
o1_mf_1_168__05vmwt34_.arc
-r--r-----. 1 oracle 54331 177694208 May 30 16:43
o1_mf_1_169__05y45qdd_.arc
-r--r-----. 1 oracle 54331 170814976 May 30 16:44
o1_mf_1_170__060kgh33_.arc
-r--r-----. 1 oracle 54331 177325056 May 30 16:46
o1_mf_1_171__063ltvgv_.arc
-r--r-----. 1 oracle 54331 164455424 May 30 16:47
o1_mf_1_172__065d94fq_.arc
-r--r-----. 1 oracle 54331 178252288 May 30 16:48
o1_mf_1_173__067wnwy8_.arc
-r--r-----. 1 oracle 54331 170579456 May 30 16:50
o1_mf_1_174__06b9zdh8_.arc
-r--r-----. 1 oracle 54331 93928960 May 30 17:26
o1_mf_1_175__08c7jc2b_.arc
[ec2-user@ip-172-30-15-124 ~]$

```

9. As oracle user, set ORACLE_HOME variable to current Oracle installation on new EC2 instance DB host ora_02, ORACLE_SID to primary Oracle instance SID. In this case, it is db1.
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 FSxN NFS mount path as defined in primary Oracle VLDB instance. flash recovery area configuration is demonstrated in section Setup Oracle RMAN incremental merge to image copy on FSx. Set the Oracle control file to FSx ONTAP NFS file system.

```
vi $ORACLE_HOME/dbs/initdb1.ora
```

With following example entries:

```
*.audit_file_dest='/u01/app/oracle/admin/db1/adump'  
*.audit_trail='db'  
*.compatible='19.0.0'  
*.control_files=('/nfsfsxn/oracopy/db1.ctl')  
*.db_block_size=8192  
*.db_create_file_dest='/nfsfsxn/oracopy/'  
*.db_domain='demo.netapp.com'  
*.db_name='db1'  
*.db_recovery_file_dest_size=85899345920  
*.db_recovery_file_dest='/nfsfsxn/archlog/'  
*.diagnostic_dest='/u01/app/oracle'  
*.dispatchers='(PROTOCOL=TCP) (SERVICE=db1XDB)'  
*.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 DB server in the case of discrepancy.

11. As oracle user, launch RMAN to run Oracle recovery on a new EC2 DB instance host.


```
[oracle@ip-172-30-15-124 dbs]$ rman target / nocatalog;
```

```
Recovery Manager: Release 19.0.0.0.0 - Production on Wed May 31  
00:56:07 2023  
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 12884900632 bytes
```

```
Fixed Size 9177880 bytes
```

```
Variable Size 1778384896 bytes
```

```
Database Buffers 11072962560 bytes
```

```
Redo Buffers 24375296 bytes
```

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

```
RMAN> set dbid = 1730530050;
```

```
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 31-MAY-23
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=2 device type=DISK

recovery area destination: /nfsfsxn/archlog
database name (or database unique name) used for search: DB1
channel ORA_DISK_1: AUTOBACKUP
/nfsfsxn/archlog/DB1/autobackup/2023_05_30/o1_mf_s_1138210401__08qlx
rrr_.bkp found in the recovery area
channel ORA_DISK_1: looking for AUTOBACKUP on day: 20230531
channel ORA_DISK_1: looking for AUTOBACKUP on day: 20230530
channel ORA_DISK_1: restoring control file from AUTOBACKUP
/nfsfsxn/archlog/DB1/autobackup/2023_05_30/o1_mf_s_1138210401__08qlx
rrr_.bkp
channel ORA_DISK_1: control file restore from AUTOBACKUP complete
output file name=/nfsfsxn/oracopy/db1.ctl
Finished restore at 31-MAY-23

```

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

```

RMAN> restore spfile to pfile '/tmp/archive/initdb1.ora' from
autobackup;

Starting restore at 31-MAY-23
using channel ORA_DISK_1

recovery area destination: /nfsfsxn/archlog
database name (or database unique name) used for search: DB1
channel ORA_DISK_1: AUTOBACKUP
/nfsfsxn/archlog/DB1/autobackup/2023_05_30/o1_mf_s_1138210401__08qlx
rrr_.bkp found in the recovery area
channel ORA_DISK_1: looking for AUTOBACKUP on day: 20230531
channel ORA_DISK_1: looking for AUTOBACKUP on day: 20230530
channel ORA_DISK_1: restoring spfile from AUTOBACKUP
/nfsfsxn/archlog/DB1/autobackup/2023_05_30/o1_mf_s_1138210401__08qlx
rrr_.bkp
channel ORA_DISK_1: SPFILE restore from AUTOBACKUP complete
Finished restore at 31-MAY-23

```

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 'OraCopyBKUPonFSxN_level_0';
```

```
List of Datafile Copies
```

```
=====
```

Key	File S	Completion Time	Ckp SCN	Ckp Time	Sparse
316	1 A	30-MAY-23	4120170	30-MAY-23	NO
	Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-1_4f1t506m				
	Tag: ORACOPYBKUPONFSXN_LEVEL_0				
322	3 A	30-MAY-23	4120175	30-MAY-23	NO
	Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-3_4g1t506m				
	Tag: ORACOPYBKUPONFSXN_LEVEL_0				
317	4 A	30-MAY-23	4120179	30-MAY-23	NO
	Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-4_4h1t5083				
	Tag: ORACOPYBKUPONFSXN_LEVEL_0				
221	5 A	26-MAY-23	2383520	12-MAY-23	NO
	Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-5_4q1t509n				
	Tag: ORACOPYBKUPONFSXN_LEVEL_0				
	Container ID: 2, PDB Name: PDB\$SEED				
216	6 A	26-MAY-23	2383520	12-MAY-23	NO
	Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-6_4m1t508t				
	Tag: ORACOPYBKUPONFSXN_LEVEL_0				
	Container ID: 2, PDB Name: PDB\$SEED				
323	7 A	30-MAY-23	4120207	30-MAY-23	NO
	Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-7_4u1t50a6				
	Tag: ORACOPYBKUPONFSXN_LEVEL_0				
227	8 A	26-MAY-23	2383520	12-MAY-23	NO
	Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-8_4t1t50a6				

```

Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 2, PDB Name: PDB$SEED

308      9      A 30-MAY-23      4120158      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
SYSTEM_FNO-9_4nlt509m
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 3, PDB Name: DB1_PDB1

307      10     A 30-MAY-23      4120166      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
SYS_AUX_FNO-10_4ilt5083
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 3, PDB Name: DB1_PDB1

313      11     A 30-MAY-23      4120154      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
UNDOTBS1_FNO-11_4l1t508t
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 3, PDB Name: DB1_PDB1

315      12     A 30-MAY-23      4120162      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-
12_4v1t50aa
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 3, PDB Name: DB1_PDB1

319      13     A 30-MAY-23      4120191      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
SYSTEM_FNO-13_4olt509m
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 4, PDB Name: DB1_PDB2

318      14     A 30-MAY-23      4120183      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
SYS_AUX_FNO-14_4j1t508s
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 4, PDB Name: DB1_PDB2

324      15     A 30-MAY-23      4120199      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
UNDOTBS1_FNO-15_4r1t50a6
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 4, PDB Name: DB1_PDB2

325      16     A 30-MAY-23      4120211      30-MAY-23      NO

```

```

Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-
16_501t50ad
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 4, PDB Name: DB1_PDB2

320    17    A 30-MAY-23      4120195      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
SYSTEM_FNO-17_4p1t509m
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 5, PDB Name: DB1_PDB3

321    18    A 30-MAY-23      4120187      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
SYS_AUX_FNO-18_4k1t508t
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 5, PDB Name: DB1_PDB3

326    19    A 30-MAY-23      4120203      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
UNDOTBS1_FNO-19_4s1t50a6
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 5, PDB Name: DB1_PDB3

327    20    A 30-MAY-23      4120216      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-
20_511t50ad
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 5, PDB Name: DB1_PDB3

298    21    A 30-MAY-23      4120166      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
21_3o1t4ut2
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 3, PDB Name: DB1_PDB1

302    22    A 30-MAY-23      4120154      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
22_3p1t4ut3
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 3, PDB Name: DB1_PDB1

297    23    A 30-MAY-23      4120158      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
23_3q1t4ut3
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 3, PDB Name: DB1_PDB1

```

306	24	A	30-MAY-23	4120162	30-MAY-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-						
24_3r1t4ut3						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
300	25	A	30-MAY-23	4120166	30-MAY-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-						
25_3s1t4v1a						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
305	26	A	30-MAY-23	4120154	30-MAY-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-						
26_451t4vt7						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
299	27	A	30-MAY-23	4120158	30-MAY-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-						
27_461t4vt7						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
310	28	A	30-MAY-23	4120162	30-MAY-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-						
28_471t4vt7						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
303	29	A	30-MAY-23	4120166	30-MAY-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-						
29_481t4vt7						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
309	30	A	30-MAY-23	4120154	30-MAY-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-						
30_491t5014						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
301	31	A	30-MAY-23	4120158	30-MAY-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-						
31_4a1t5015						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						

```

Container ID: 3, PDB Name: DB1_PDB1

312      32      A 30-MAY-23      4120162      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
32_4b1t501u
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 3, PDB Name: DB1_PDB1

314      33      A 30-MAY-23      4120162      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
33_4c1t501v
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 3, PDB Name: DB1_PDB1

304      34      A 30-MAY-23      4120158      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
34_4d1t5058
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 3, PDB Name: DB1_PDB1

311      35      A 30-MAY-23      4120154      30-MAY-23      NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
35_4e1t5059
Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 3, PDB Name: DB1_PDB1

```

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

```

RMAN> switch database to copy;

Starting implicit crosscheck backup at 31-MAY-23
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=11 device type=DISK
Crosschecked 33 objects
Finished implicit crosscheck backup at 31-MAY-23

Starting implicit crosscheck copy at 31-MAY-23
using channel ORA_DISK_1
Crosschecked 68 objects
Finished implicit crosscheck copy at 31-MAY-23

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

List of Cataloged Files

```

=====
File Name:

/nfsfsxn/archlog/DB1/autobackup/2023_05_30/o1_mf_s_1138210401__08qlx
rrr_.bkp

datafile 1 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-
1730530050_TS-SYSTEM_FNO-1_4f1t506m"

datafile 3 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-
1730530050_TS-SYSAUX_FNO-3_4g1t506m"

datafile 4 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-
1730530050_TS-UNDOTBS1_FNO-4_4h1t5083"

datafile 5 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-
1730530050_TS-SYSTEM_FNO-5_4q1t509n"

datafile 6 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-
1730530050_TS-SYSAUX_FNO-6_4m1t508t"

datafile 7 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-
1730530050_TS-USERS_FNO-7_4u1t50a6"

datafile 8 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-
1730530050_TS-UNDOTBS1_FNO-8_4t1t50a6"

datafile 9 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-
1730530050_TS-SYSTEM_FNO-9_4n1t509m"

datafile 10 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SYSAUX_FNO-10_4i1t5083"

datafile 11 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-UNDOTBS1_FNO-11_4l1t508t"

datafile 12 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-USERS_FNO-12_4v1t50aa"

datafile 13 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SYSTEM_FNO-13_4o1t509m"

datafile 14 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SYSAUX_FNO-14_4j1t508s"

datafile 15 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-UNDOTBS1_FNO-15_4r1t50a6"

datafile 16 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-USERS_FNO-16_501t50ad"

datafile 17 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SYSTEM_FNO-17_4p1t509m"

datafile 18 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SYSAUX_FNO-18_4k1t508t"

datafile 19 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-UNDOTBS1_FNO-19_4s1t50a6"

datafile 20 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-USERS_FNO-20_511t50ad"

datafile 21 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SOE_FNO-21_3o1t4ut2"

datafile 22 switched to datafile copy "/nfsfsxn/oracopy/data_D-


```
DB1_I-1730530050_TS-SOE_FNO-22_3p1t4ut3"
datafile 23 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SOE_FNO-23_3q1t4ut3"
datafile 24 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SOE_FNO-24_3r1t4ut3"
datafile 25 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SOE_FNO-25_3s1t4v1a"
datafile 26 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SOE_FNO-26_451t4vt7"
datafile 27 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SOE_FNO-27_461t4vt7"
datafile 28 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SOE_FNO-28_471t4vt7"
datafile 29 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SOE_FNO-29_481t4vt7"
datafile 30 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SOE_FNO-30_491t5014"
datafile 31 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SOE_FNO-31_4a1t5015"
datafile 32 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SOE_FNO-32_4b1t501u"
datafile 33 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SOE_FNO-33_4c1t501v"
datafile 34 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SOE_FNO-34_4d1t5058"
datafile 35 switched to datafile copy "/nfsfsxn/oracopy/data_D-
DB1_I-1730530050_TS-SOE_FNO-35_4e1t5059"
```

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

```
RMAN> run {
2> set until sequence=176;
3> recover database;
4> }

executing command: SET until clause

Starting recover at 31-MAY-23
using channel ORA_DISK_1

starting media recovery

archived log for thread 1 with sequence 142 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_142__02n3x2qb_.ar
```

c
archived log for thread 1 with sequence 143 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_143__02rotwyb_.ar
c
archived log for thread 1 with sequence 144 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_144__02x563wh_.ar
c
archived log for thread 1 with sequence 145 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_145__031kg2co_.ar
c
archived log for thread 1 with sequence 146 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_146__035xpcdt_.ar
c
archived log for thread 1 with sequence 147 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_147__03bds8qf_.ar
c
archived log for thread 1 with sequence 148 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_148__03gyt7rx_.ar
c
archived log for thread 1 with sequence 149 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_149__03mfxl7v_.ar
c
archived log for thread 1 with sequence 150 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_150__03qzz0ty_.ar
c
archived log for thread 1 with sequence 151 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_151__03wgxdry_.ar
c
archived log for thread 1 with sequence 152 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_152__040y85v3_.ar
c
archived log for thread 1 with sequence 153 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_153__04ox946w_.ar
c
archived log for thread 1 with sequence 154 is already on disk as

```
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_154__04rbv7n8_.ar
c
archived log for thread 1 with sequence 155 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_155__04tvlyvn_.ar
c
archived log for thread 1 with sequence 156 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_156__04xgfjtl_.ar
c
archived log for thread 1 with sequence 157 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_157__04zyg8hw_.ar
c
archived log for thread 1 with sequence 158 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_158__052gp9mt_.ar
c
archived log for thread 1 with sequence 159 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_159__0551wk7s_.ar
c
archived log for thread 1 with sequence 160 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_160__057146my_.ar
c
archived log for thread 1 with sequence 161 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_161__05b2dmwp_.ar
c
archived log for thread 1 with sequence 162 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_162__05drbj8n_.ar
c
archived log for thread 1 with sequence 163 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_163__05h81mlh_.ar
c
archived log for thread 1 with sequence 164 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_164__05krsqmh_.ar
c
archived log for thread 1 with sequence 165 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_165__05n378pw_.ar
```

```
c
archived log for thread 1 with sequence 166 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_166__05pmg741_.ar
c
archived log for thread 1 with sequence 167 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_167__05s3o01r_.ar
c
archived log for thread 1 with sequence 168 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_168__05vmwt34_.ar
c
archived log for thread 1 with sequence 169 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_169__05y45qdd_.ar
c
archived log for thread 1 with sequence 170 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_170__060kgh33_.ar
c
archived log for thread 1 with sequence 171 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_171__0631tvgv_.ar
c
archived log for thread 1 with sequence 172 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_172__065d94fq_.ar
c
archived log for thread 1 with sequence 173 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_173__067wnwy8_.ar
c
archived log for thread 1 with sequence 174 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_174__06b9zdh8_.ar
c
archived log for thread 1 with sequence 175 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_175__08c7jc2b_.ar
c
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_142__02n3x2q
b_.arc thread=1 sequence=142
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_143__02rotwy
```

b_.arc thread=1 sequence=143
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_144__02x563w
h_.arc thread=1 sequence=144
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_145__031kg2c
o_.arc thread=1 sequence=145
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_146__035xpcd
t_.arc thread=1 sequence=146
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_147__03bds8q
f_.arc thread=1 sequence=147
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_148__03gyt7r
x_.arc thread=1 sequence=148
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_149__03mfx17
v_.arc thread=1 sequence=149
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_150__03qzz0t
y_.arc thread=1 sequence=150
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_151__03wgxdr
y_.arc thread=1 sequence=151
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_152__040y85v
3_.arc thread=1 sequence=152
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_153__04ox946
w_.arc thread=1 sequence=153
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_154__04rbv7n
8_.arc thread=1 sequence=154
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_155__04tvlyv
n_.arc thread=1 sequence=155
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_156__04xgfjt
l_.arc thread=1 sequence=156
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_157__04zyg8h
w_.arc thread=1 sequence=157
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_158__052gp9m
t_.arc thread=1 sequence=158

archived log file

name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_159__0551wk7
s_.arc thread=1 sequence=159

archived log file

name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_160__057146m
y_.arc thread=1 sequence=160

archived log file

name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_161__05b2dmw
p_.arc thread=1 sequence=161

archived log file

name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_162__05drbj8
n_.arc thread=1 sequence=162

archived log file

name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_163__05h81m1
h_.arc thread=1 sequence=163

archived log file

name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_164__05krsqm
h_.arc thread=1 sequence=164

archived log file

name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_165__05n378p
w_.arc thread=1 sequence=165

archived log file

name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_166__05pmg74
l_.arc thread=1 sequence=166

archived log file

name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_167__05s3o01
r_.arc thread=1 sequence=167

archived log file

name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_168__05vmwt3
4_.arc thread=1 sequence=168

archived log file

name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_169__05y45qd
d_.arc thread=1 sequence=169

archived log file

name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_170__060kgh3
3_.arc thread=1 sequence=170

archived log file

name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_171__0631tv
g_.arc thread=1 sequence=171

archived log file

name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_172__065d94f
q_.arc thread=1 sequence=172

archived log file

name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_173__067wnwy
8_.arc thread=1 sequence=173

archived log file

```
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_174__06b9zdh
8_.arc thread=1 sequence=174
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_05_30/o1_mf_1_175__08c7jc2
b_.arc thread=1 sequence=175
media recovery complete, elapsed time: 00:48:34
Finished recover at 31-MAY-23
```



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
-----
DB1           MOUNTED
```

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

```
MEMBER
-----
+DATA/DB1/ONLINELOG/group_3.264.1136666437
+DATA/DB1/ONLINELOG/group_2.263.1136666437
+DATA/DB1/ONLINELOG/group_1.262.1136666437
```

```
SQL> alter database rename file
'+DATA/DB1/ONLINELOG/group_1.262.1136666437' to
'/nfsfsxn/oracopy/redo01.log';
```

Database altered.

```
SQL> alter database rename file
'+DATA/DB1/ONLINELOG/group_2.263.1136666437' to
'/nfsfsxn/oracopy/redo02.log';
```

Database altered.

```
SQL> alter database rename file
'+DATA/DB1/ONLINELOG/group_3.264.1136666437' to
'/nfsfsxn/oracopy/redo03.log';
```

Database altered.

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

Database altered.

19. Validate the database restored to new host that has the row we have inserted before primary database failure.


```
SQL> show pdbs
```

CON_ID	CON_NAME	OPEN MODE	RESTRICTED
2	PDB\$SEED	READ ONLY	NO
3	DB1_PDB1	READ WRITE	NO
4	DB1_PDB2	READ WRITE	NO
5	DB1_PDB3	READ WRITE	NO

```
SQL> alter session set container=db1_pdb1;
```

Session altered.

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

ID	DT
1	18-MAY-23 02.35.37.000000 PM
test oracle incremental merge switch to copy	
2	30-MAY-23 05.23.11.000000 PM
test recovery on a new EC2 instance host with image copy on FSxN	

20. Other post recovery tasks

Add FSxN NFS mount to fstab so that the NFS file system will be mounted when EC2 instance host rebooted.

As EC2 user, vi /etc/fstab and add following entry:

```
172.30.15.19:/ora_01_copy          /nfsfsxn          nfs
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=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 FSxN NFS file system to a new EC2 DB instance host.

Clone Oracle standby image copy for other use cases

Another benefit of using AWS FSx ONTAP for staging Oracle VLDB image copy is that it can be FlexCloned to serve many other purposes with minimal additional storage investment. In the following use case, we demonstrate how to snapshot and clone the staging NFS volume on FSx ONTAP for other Oracle use cases such as DEV, UAT, etc.

1. We begin with inserting a row into the same test table we have created before.

```
SQL> insert into test values (3, sysdate, 'test clone on a new EC2
instance host with image copy on FSxN');
```

```
1 row created.
```

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

```
          ID
-----
DT
-----
EVENT
-----
          1
18-MAY-23 02.35.37.000000 PM
test oracle incremental merge switch to copy

          2
30-MAY-23 05.23.11.000000 PM
test recovery on a new EC2 instance host with image copy on FSxN

          ID
-----
DT
-----
EVENT
-----
          3
05-JUN-23 03.19.46.000000 PM
test clone on a new EC2 instance host with image copy on FSxN

SQL>
```

2. Take a RMAN backup and merge to FSx ONTAP database image copy so that the transaction will be captured in the backup set on FSx NFS mount but not merged into copy until cloned database is recovered.

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

3. Login to FSx cluster via ssh as fsxadmin user to observe the snapshots created by scheduled backup policy - oracle and take an one-off snapshot so that it will include the transaction we committed in step 1.

```
FsxId06c3c8b2a7bd56458::> vol snapshot create -vserver svm_ora
-volume ora_01_copy -snapshot one-off.2023-06-05-1137 -foreground
true
```

```
FsxId06c3c8b2a7bd56458::> snapshot show
```

```
---Blocks---
```

```
Vserver Volume Snapshot Size
Total% Used%
```

```
-----
```

Vserver	Volume	Snapshot	Size
svm_ora	ora_01_copy		
		daily.2023-06-02_0010	3.59GB
2%	5%		
		daily.2023-06-03_0010	1.10GB
1%	1%		
		daily.2023-06-04_0010	608KB
0%	0%		
		daily.2023-06-05_0010	3.81GB
2%	5%		
		one-off.2023-06-05-1137	168KB
0%	0%		
	svm_ora_root		
		weekly.2023-05-28_0015	1.86MB
0%	78%		
		daily.2023-06-04_0010	152KB
0%	22%		
		weekly.2023-06-04_0015	1.24MB
0%	70%		
		daily.2023-06-05_0010	196KB
0%	27%		
		hourly.2023-06-05_1005	156KB
0%	22%		
		hourly.2023-06-05_1105	156KB
0%	22%		
		hourly.2023-06-05_1205	156KB
0%	22%		
		hourly.2023-06-05_1305	156KB
0%	22%		
		hourly.2023-06-05_1405	1.87MB
0%	78%		
		hourly.2023-06-05_1505	148KB
0%	22%		

```
15 entries were displayed.
```

4. Clone from the one-off snapshot to be used for standing up a new DB1 clone instance on an alternative EC2 Oracle host. You have the option to clone from any available daily snapshots for volume ora_01_copy.

```
FsxId06c3c8b2a7bd56458::> vol clone create -flexclone db1_20230605of
-type RW -parent-vserver svm_ora -parent-volume ora_01_copy
-junction-path /db1_20230605of -junction-active true -parent
-snapshot one-off.2023-06-05-1137
[Job 464] Job succeeded: Successful

FsxId06c3c8b2a7bd56458::>

FsxId06c3c8b2a7bd56458::> vol show db1*
Vserver   Volume           Aggregate      State          Type          Size
Available Used%
-----
-----
svm_ora   db1_20230605of
                aggr1          online        RW            200GB
116.6GB  38%

FsxId06c3c8b2a7bd56458::>
```

5. Turn off snapshot policy for the cloned volume as it inherits parent volume snapshot policy unless you want to protect the cloned volume, then leave it alone.

```
FsxId06c3c8b2a7bd56458::> vol modify -volume db1_20230605of
-snapshot-policy none

Warning: You are changing the Snapshot policy on volume
"db1_20230605of" to "none". Snapshot copies on this volume that do
not match any of the prefixes of the new Snapshot policy will not be
deleted. However, when the new Snapshot policy
                takes effect, depending on the new retention count, any
existing Snapshot copies that continue to use the same prefixes
might be deleted. See the 'volume modify' man page for more
information.
Do you want to continue? {y|n}: y
Volume modify successful on volume db1_20230605of of Vserver
svm_ora.

FsxId06c3c8b2a7bd56458::>
```

6. Login to a new EC2 Linux instance with Oracle software pre-installed with same version and patch level as your primary Oracle EC2 instance and mount the cloned volume.

```
[ec2-user@ip-172-30-15-124 ~]$ sudo mkdir /nfsfsxn
[ec2-user@ip-172-30-15-124 ~]$ sudo mount -t nfs
172.30.15.19:/db1_20230605of /nfsfsxn -o
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsiz=262144,wsiz=262144,noi
tr
```

7. Validate the database incremental backup sets, image copy, and available archived logs on FSx NFS mount.

```
[ec2-user@ip-172-30-15-124 ~]$ ls -ltr /nfsfsxn/oracopy
total 79450332
-rw-r----- 1 oracle 54331 482353152 Jun  1 19:02 data_D-DB1_I-
1730530050_TS-SYSAUX_FNO-6_891tkrhr
-rw-r----- 1 oracle 54331 419438592 Jun  1 19:03 data_D-DB1_I-
1730530050_TS-SYSTEM_FNO-5_8d1tkril
-rw-r----- 1 oracle 54331 241180672 Jun  1 19:03 data_D-DB1_I-
1730530050_TS-UNDOTBS1_FNO-8_8g1tkrj7
-rw-r----- 1 oracle 54331 912506880 Jun  1 20:21 8n1tkvv2_279_1_1
-rw-r----- 1 oracle 54331 925696 Jun  1 20:21 8q1tl05i_282_1_1
-rw-r----- 1 oracle 54331 1169014784 Jun  1 20:21 8p1tkvv2_281_1_1
-rw-r----- 1 oracle 54331 6455296 Jun  1 20:21 8r1tl05m_283_1_1
-rw-r----- 1 oracle 54331 139264 Jun  1 20:21 8t1tl05t_285_1_1
-rw-r----- 1 oracle 54331 3514368 Jun  1 20:21 8s1tl05t_284_1_1
-rw-r----- 1 oracle 54331 139264 Jun  1 20:21 8u1tl060_286_1_1
-rw-r----- 1 oracle 54331 425984 Jun  1 20:21 901tl062_288_1_1
-rw-r----- 1 oracle 54331 344064 Jun  1 20:21 911tl062_289_1_1
-rw-r----- 1 oracle 54331 245760 Jun  1 20:21 931tl063_291_1_1
-rw-r----- 1 oracle 54331 237568 Jun  1 20:21 941tl064_292_1_1
-rw-r----- 1 oracle 54331 57344 Jun  1 20:21 961tl065_294_1_1
-rw-r----- 1 oracle 54331 57344 Jun  1 20:21 971tl066_295_1_1
-rw-r----- 1 oracle 54331 57344 Jun  1 20:21 981tl067_296_1_1
-rw-r----- 1 oracle 54331 1040760832 Jun  1 20:23 8m1tkvv2_278_1_1
-rw-r----- 1 oracle 54331 932847616 Jun  1 20:24 8o1tkvv2_280_1_1
-rw-r----- 1 oracle 54331 1121984512 Jun  5 15:21 data_D-DB1_I-
1730530050_TS-SYSTEM_FNO-1_821tkrb8
-rw-r----- 1 oracle 54331 1027612672 Jun  5 15:21 data_D-DB1_I-
1730530050_TS-SYSAUX_FNO-3_831tkrd9
-rw-r----- 1 oracle 54331 429924352 Jun  5 15:21 data_D-DB1_I-
1730530050_TS-SYSTEM_FNO-9_8a1tkrhr
-rw-r----- 1 oracle 54331 707796992 Jun  5 15:21 data_D-DB1_I-
1730530050_TS-UNDOTBS1_FNO-4_851tkrgf
-rw-r----- 1 oracle 54331 534781952 Jun  5 15:21 data_D-DB1_I-
1730530050_TS-SYSAUX_FNO-14_871tkrhr
-rw-r----- 1 oracle 54331 534781952 Jun  5 15:21 data_D-DB1_I-
1730530050_TS-SYSAUX_FNO-18_881tkrhr
```

```

-rw-r----- 1 oracle 54331 429924352 Jun 5 15:21 data_D-DB1_I-
1730530050_TS-SYSTEM_FNO-13_8b1tkril
-rw-r----- 1 oracle 54331 429924352 Jun 5 15:21 data_D-DB1_I-
1730530050_TS-SYSTEM_FNO-17_8c1tkril
-rw-r----- 1 oracle 54331 246423552 Jun 5 15:21 data_D-DB1_I-
1730530050_TS-UNDOTBS1_FNO-15_8e1tkril
-rw-r----- 1 oracle 54331 246423552 Jun 5 15:21 data_D-DB1_I-
1730530050_TS-UNDOTBS1_FNO-19_8f1tkrj4
-rw-r----- 1 oracle 54331 5251072 Jun 5 15:21 data_D-DB1_I-
1730530050_TS-USERS_FNO-7_8h1tkrj9
-rw-r----- 1 oracle 54331 5251072 Jun 5 15:21 data_D-DB1_I-
1730530050_TS-USERS_FNO-16_8j1tkrja
-rw-r----- 1 oracle 54331 5251072 Jun 5 15:21 data_D-DB1_I-
1730530050_TS-USERS_FNO-20_8k1tkrjb
-rw-r----- 1 oracle 54331 5251072 Jun 5 15:21 data_D-DB1_I-
1730530050_TS-USERS_FNO-12_8i1tkrj9
-rw-r----- 1 oracle 54331 555753472 Jun 5 15:21 data_D-DB1_I-
1730530050_TS-SYSAUX_FNO-10_861tkrgo
-rw-r----- 1 oracle 54331 796925952 Jun 5 15:22 data_D-DB1_I-
1730530050_TS-UNDOTBS1_FNO-11_841tkrf2
-rw-r----- 1 oracle 54331 4294975488 Jun 5 15:22 data_D-DB1_I-
1730530050_TS-SOE_FNO-21_7j1tkqk6
-rw-r----- 1 oracle 54331 4294975488 Jun 5 15:22 data_D-DB1_I-
1730530050_TS-SOE_FNO-34_801tkram
-rw-r----- 1 oracle 54331 4294975488 Jun 5 15:22 data_D-DB1_I-
1730530050_TS-SOE_FNO-29_7r1tkr32
-rw-r----- 1 oracle 54331 4294975488 Jun 5 15:22 data_D-DB1_I-
1730530050_TS-SOE_FNO-25_7n1tkqrh
-rw-r----- 1 oracle 54331 4294975488 Jun 5 15:22 data_D-DB1_I-
1730530050_TS-SOE_FNO-31_7t1tkr3i
-rw-r----- 1 oracle 54331 4294975488 Jun 5 15:22 data_D-DB1_I-
1730530050_TS-SOE_FNO-33_7v1tkra6
-rw-r----- 1 oracle 54331 4294975488 Jun 5 15:22 data_D-DB1_I-
1730530050_TS-SOE_FNO-23_7l1tkqk6
-rw-r----- 1 oracle 54331 4294975488 Jun 5 15:22 data_D-DB1_I-
1730530050_TS-SOE_FNO-27_7p1tkqrq
-rw-r----- 1 oracle 54331 4294975488 Jun 5 15:22 data_D-DB1_I-
1730530050_TS-SOE_FNO-35_8l1tkrap
-rw-r----- 1 oracle 54331 4294975488 Jun 5 15:22 data_D-DB1_I-
1730530050_TS-SOE_FNO-32_7u1tkr42
-rw-r----- 1 oracle 54331 4294975488 Jun 5 15:22 data_D-DB1_I-
1730530050_TS-SOE_FNO-22_7k1tkqk6
-rw-r----- 1 oracle 54331 4294975488 Jun 5 15:22 data_D-DB1_I-
1730530050_TS-SOE_FNO-24_7m1tkqk6
-rw-r----- 1 oracle 54331 4294975488 Jun 5 15:22 data_D-DB1_I-
1730530050_TS-SOE_FNO-28_7q1tkqs1

```



```

-rw-r----- 1 oracle 54331 4294975488 Jun  5 15:22 data_D-DB1_I-
1730530050_TS-SOE_FNO-30_7s1tkr3a
-rw-r----- 1 oracle 54331 4294975488 Jun  5 15:22 data_D-DB1_I-
1730530050_TS-SOE_FNO-26_7o1tkqrj
-rw-r----- 1 oracle 54331 1241432064 Jun  5 15:30 9d1tv06n_301_1_1
-rw-r----- 1 oracle 54331 1019805696 Jun  5 15:31 9a1tv06m_298_1_1
-rw-r----- 1 oracle 54331      4612096 Jun  5 15:31 9e1tv01d_302_1_1
-rw-r----- 1 oracle 54331  967163904 Jun  5 15:31 9b1tv06n_299_1_1
-rw-r----- 1 oracle 54331  31563776 Jun  5 15:31 9g1tv01t_304_1_1
-rw-r----- 1 oracle 54331    319488 Jun  5 15:31 9h1tv01t_305_1_1
-rw-r----- 1 oracle 54331   335872 Jun  5 15:31 9i1tv0m0_306_1_1
-rw-r----- 1 oracle 54331   565248 Jun  5 15:31 9k1tv0m1_308_1_1
-rw-r----- 1 oracle 54331   581632 Jun  5 15:31 9l1tv0m5_309_1_1
-rw-r----- 1 oracle 54331  54345728 Jun  5 15:31 9f1tv01t_303_1_1
-rw-r----- 1 oracle 54331   368640 Jun  5 15:31 9n1tv0m5_311_1_1
-rw-r----- 1 oracle 54331   385024 Jun  5 15:31 9o1tv0m6_312_1_1
-rw-r----- 1 oracle 54331  985858048 Jun  5 15:31 9c1tv06n_300_1_1
-rw-r----- 1 oracle 54331    57344 Jun  5 15:31 9q1tv0m7_314_1_1
-rw-r----- 1 oracle 54331    57344 Jun  5 15:31 9r1tv0m8_315_1_1
-rw-r----- 1 oracle 54331    57344 Jun  5 15:31 9s1tv0m9_316_1_1
-rw-r--r-- 1 oracle 54331    12720 Jun  5 15:31 db1_ctl.sql
-rw-r----- 1 oracle 54331  11600384 Jun  5 15:48 bct_db1.ctf
[ec2-user@ip-172-30-15-124 ~]$

```

```

[oracle@ip-172-30-15-124 ~]$ ls -l
/nfsfsxn/archlog/DB1/archivelog/2023_06_05
total 2008864
-rw-r----- 1 oracle 54331    729088 Jun  5 14:38
o1_mf_1_190_17vwwvt9_.arc
-rw-r----- 1 oracle 54331 166651904 Jun  5 14:44
o1_mf_1_191_17vx6vmg_.arc
-rw-r----- 1 oracle 54331 167406080 Jun  5 14:47
o1_mf_1_192_17vxctms_.arc
-rw-r----- 1 oracle 54331 166868992 Jun  5 14:49
o1_mf_1_193_17vxjjps_.arc
-rw-r----- 1 oracle 54331 166087168 Jun  5 14:52
o1_mf_1_194_17vxnxrh_.arc
-rw-r----- 1 oracle 54331 175210496 Jun  5 14:54
o1_mf_1_195_17vxswv5_.arc
-rw-r----- 1 oracle 54331 167078400 Jun  5 14:57
o1_mf_1_196_17vxylwp_.arc
-rw-r----- 1 oracle 54331 169701888 Jun  5 14:59
o1_mf_1_197_17vy3cyw_.arc
-rw-r----- 1 oracle 54331 167845376 Jun  5 15:02
o1_mf_1_198_17vy8245_.arc
-rw-r----- 1 oracle 54331 170763776 Jun  5 15:05

```

```
o1_mf_1_199_17vydv4c_.arc
-rw-r----- 1 oracle 54331 193853440 Jun  5 15:07
o1_mf_1_200_17vykf23_.arc
-rw-r----- 1 oracle 54331 165523968 Jun  5 15:09
o1_mf_1_201_17vyp1dh_.arc
-rw-r----- 1 oracle 54331 161117184 Jun  5 15:12
o1_mf_1_202_17vyvrm5_.arc
-rw-r----- 1 oracle 54331  10098176 Jun  5 15:21
o1_mf_1_203_17vzdfwm_.arc
```

8. The recovery processes now are similar to previous use case of recovery to a new EC2 DB instance after a failure - set oracle environment (oratab, \$ORACLE_HOME, \$ORACLE_SID) to match with primary production instance, create an init file including db_recovery_file_dest_size and db_recovery_file_dest that point to flash recovery directory on FSx NFS mount. Then, launch RMAN to run recovery. Following are command steps and output.

```
[oracle@ip-172-30-15-124 dbs]$ rman target / nocatalog

Recovery Manager: Release 19.0.0.0.0 - Production on Wed Jun 7
14:44:33 2023
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

RMAN> set dbid = 1730530050;

executing command: SET DBID

RMAN> restore controlfile from autobackup;

Starting restore at 07-JUN-23
allocated channel: ORA_DISK_1
```

```

channel ORA_DISK_1: SID=2 device type=DISK

recovery area destination: /nfsfsxn/archlog/
database name (or database unique name) used for search: DB1
channel ORA_DISK_1: AUTOBACKUP
/nfsfsxn/archlog/DB1/autobackup/2023_06_05/o1_mf_s_1138721482_17vzyb
vq_.bkp found in the recovery area
channel ORA_DISK_1: looking for AUTOBACKUP on day: 20230607
channel ORA_DISK_1: looking for AUTOBACKUP on day: 20230606
channel ORA_DISK_1: looking for AUTOBACKUP on day: 20230605
channel ORA_DISK_1: restoring control file from AUTOBACKUP
/nfsfsxn/archlog/DB1/autobackup/2023_06_05/o1_mf_s_1138721482_17vzyb
vq_.bkp
channel ORA_DISK_1: control file restore from AUTOBACKUP complete
output file name=/nfsfsxn/oracopy/db1ctl
Finished restore at 07-JUN-23

```

```

RMAN> alter database mount;

```

```

released channel: ORA_DISK_1
Statement processed

```

```

RMAN> list incarnation;

```

List of Database Incarnations

DB Key	Inc Key	DB Name	DB ID	STATUS	Reset SCN	Reset Time
1	1	DB1	1730530050	PARENT	1	17-APR-19
2	2	DB1	1730530050	CURRENT	1920977	12-MAY-23

```

RMAN> list copy of database tag 'OraCopyBKUPonFSxN_level_0';

```

List of Datafile Copies

=====

Key	File S	Completion Time	Ckp SCN	Ckp Time	Sparse
362	1 A	05-JUN-23	8319160	01-JUN-23	NO
	Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-1_821tkrb8				
	Tag: ORACOPYBKUPONFSXN_LEVEL_0				
363	3 A	05-JUN-23	8319165	01-JUN-23	NO

```

      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
SYSAUX_FNO-3_831tkrd9
      Tag: ORACOPYBKUPONFSXN_LEVEL_0

365      4      A 05-JUN-23      8319171      01-JUN-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
UNDOTBS1_FNO-4_851tkrgf
      Tag: ORACOPYBKUPONFSXN_LEVEL_0

355      5      A 01-JUN-23      2383520      12-MAY-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
SYSTEM_FNO-5_8dltkri1
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 2, PDB Name: PDB$SEED

349      6      A 01-JUN-23      2383520      12-MAY-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
SYSAUX_FNO-6_891tkrhr
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 2, PDB Name: PDB$SEED

372      7      A 05-JUN-23      8319201      01-JUN-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-
7_8h1tkrj9
      Tag: ORACOPYBKUPONFSXN_LEVEL_0

361      8      A 01-JUN-23      2383520      12-MAY-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
UNDOTBS1_FNO-8_8g1tkrj7
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 2, PDB Name: PDB$SEED

364      9      A 05-JUN-23      8318717      01-JUN-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
SYSTEM_FNO-9_8altkrhr
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 3, PDB Name: DB1_PDB1

376      10     A 05-JUN-23      8318714      01-JUN-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
SYSAUX_FNO-10_861tkrgo
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 3, PDB Name: DB1_PDB1

377      11     A 05-JUN-23      8318720      01-JUN-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-

```

UNDOTBS1_FNO-11_841tkrf2
 Tag: ORACOPYBKUPONFSXN_LEVEL_0
 Container ID: 3, PDB Name: DB1_PDB1

375 12 A 05-JUN-23 8318719 01-JUN-23 NO
 Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-
 12_8i1tkrj9
 Tag: ORACOPYBKUPONFSXN_LEVEL_0
 Container ID: 3, PDB Name: DB1_PDB1

368 13 A 05-JUN-23 8319184 01-JUN-23 NO
 Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
 SYSTEM_FNO-13_8b1tkril
 Tag: ORACOPYBKUPONFSXN_LEVEL_0
 Container ID: 4, PDB Name: DB1_PDB2

366 14 A 05-JUN-23 8319175 01-JUN-23 NO
 Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
 SYSAUX_FNO-14_871tkrhr
 Tag: ORACOPYBKUPONFSXN_LEVEL_0
 Container ID: 4, PDB Name: DB1_PDB2

370 15 A 05-JUN-23 8319193 01-JUN-23 NO
 Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
 UNDOTBS1_FNO-15_8e1tkril
 Tag: ORACOPYBKUPONFSXN_LEVEL_0
 Container ID: 4, PDB Name: DB1_PDB2

373 16 A 05-JUN-23 8319206 01-JUN-23 NO
 Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-
 16_8j1tkrja
 Tag: ORACOPYBKUPONFSXN_LEVEL_0
 Container ID: 4, PDB Name: DB1_PDB2

369 17 A 05-JUN-23 8319188 01-JUN-23 NO
 Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
 SYSTEM_FNO-17_8c1tkril
 Tag: ORACOPYBKUPONFSXN_LEVEL_0
 Container ID: 5, PDB Name: DB1_PDB3

367 18 A 05-JUN-23 8319180 01-JUN-23 NO
 Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-
 SYSAUX_FNO-18_881tkrhr
 Tag: ORACOPYBKUPONFSXN_LEVEL_0
 Container ID: 5, PDB Name: DB1_PDB3

371	19	A	05-JUN-23	8319197	01-JUN-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-19_8f1tkrj4						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 5, PDB Name: DB1_PDB3						
374	20	A	05-JUN-23	8319210	01-JUN-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-20_8k1tkrjb						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 5, PDB Name: DB1_PDB3						
378	21	A	05-JUN-23	8318720	01-JUN-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-21_7j1tkqk6						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
388	22	A	05-JUN-23	8318714	01-JUN-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-22_7k1tkqk6						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
384	23	A	05-JUN-23	8318717	01-JUN-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-23_7l1tkqk6						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
389	24	A	05-JUN-23	8318719	01-JUN-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-24_7m1tkqk6						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
381	25	A	05-JUN-23	8318720	01-JUN-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-25_7n1tkqrh						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						
Container ID: 3, PDB Name: DB1_PDB1						
392	26	A	05-JUN-23	8318714	01-JUN-23	NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-26_7o1tkqrj						
Tag: ORACOPYBKUPONFSXN_LEVEL_0						

Container ID: 3, PDB Name: DB1_PDB1

```
385      27      A 05-JUN-23      8318717      01-JUN-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
27_7p1tkqrq
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 3, PDB Name: DB1_PDB1

390      28      A 05-JUN-23      8318719      01-JUN-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
28_7q1tkqsl
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 3, PDB Name: DB1_PDB1

380      29      A 05-JUN-23      8318720      01-JUN-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
29_7r1tkr32
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 3, PDB Name: DB1_PDB1

391      30      A 05-JUN-23      8318714      01-JUN-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
30_7s1tkr3a
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 3, PDB Name: DB1_PDB1

382      31      A 05-JUN-23      8318717      01-JUN-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
31_7t1tkr3i
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 3, PDB Name: DB1_PDB1

387      32      A 05-JUN-23      8318719      01-JUN-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
32_7ultkr42
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 3, PDB Name: DB1_PDB1

383      33      A 05-JUN-23      8318719      01-JUN-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
33_7v1tkra6
      Tag: ORACOPYBKUPONFSXN_LEVEL_0
      Container ID: 3, PDB Name: DB1_PDB1

379      34      A 05-JUN-23      8318717      01-JUN-23      NO
      Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
```

34_801tkram

Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 3, PDB Name: DB1_PDB1

386 35 A 05-JUN-23 8318714 01-JUN-23 NO
Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
35_811tkrap

Tag: ORACOPYBKUPONFSXN_LEVEL_0
Container ID: 3, PDB Name: DB1_PDB1

RMAN> switch database to copy;

datafile 1 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-1_821tkrb8"
datafile 3 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-3_831tkrd9"
datafile 4 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-4_851tkrgf"
datafile 5 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-5_8d1tkril"
datafile 6 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-6_891tkrhr"
datafile 7 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-7_8h1tkrj9"
datafile 8 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-8_8g1tkrj7"
datafile 9 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-9_8a1tkrhr"
datafile 10 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-10_861tkrgo"
datafile 11 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-11_841tkrf2"
datafile 12 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-12_8i1tkrj9"
datafile 13 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-13_8b1tkril"
datafile 14 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-14_871tkrhr"
datafile 15 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-15_8e1tkril"
datafile 16 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-16_8j1tkrja"
datafile 17 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-17_8c1tkril"
datafile 18 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-18_881tkrhr"


```
datafile 19 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-19_8f1tkrj4"  
datafile 20 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-20_8k1tkrjb"  
datafile 21 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-21_7j1tkqk6"  
datafile 22 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-22_7k1tkqk6"  
datafile 23 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-23_7l1tkqk6"  
datafile 24 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-24_7m1tkqk6"  
datafile 25 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-25_7n1tkqrh"  
datafile 26 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-26_7o1tkqrj"  
datafile 27 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-27_7p1tkqrq"  
datafile 28 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-28_7q1tkqs1"  
datafile 29 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-29_7r1tkr32"  
datafile 30 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-30_7s1tkr3a"  
datafile 31 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-31_7t1tkr3i"  
datafile 32 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-32_7u1tkr42"  
datafile 33 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-33_7v1tkra6"  
datafile 34 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-34_801tkram"  
datafile 35 switched to datafile copy "/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-35_811tkrap"
```

```
RMAN> run {  
2> set until sequence 204;  
3> recover database;  
4> }
```

executing command: SET until clause

Starting recover at 07-JUN-23
using channel ORA_DISK_1

starting media recovery

archived log for thread 1 with sequence 190 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_190_17vwvvt9_.arc
archived log for thread 1 with sequence 191 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_191_17vx6vmg_.arc
archived log for thread 1 with sequence 192 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_192_17vxctms_.arc
archived log for thread 1 with sequence 193 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_193_17vxjjps_.arc
archived log for thread 1 with sequence 194 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_194_17vxnxrh_.arc
archived log for thread 1 with sequence 195 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_195_17vxswv5_.arc
archived log for thread 1 with sequence 196 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_196_17vxylwp_.arc
archived log for thread 1 with sequence 197 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_197_17vy3cyw_.arc
archived log for thread 1 with sequence 198 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_198_17vy8245_.arc
archived log for thread 1 with sequence 199 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_199_17vydv4c_.arc
archived log for thread 1 with sequence 200 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_200_17vykf23_.arc
archived log for thread 1 with sequence 201 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_201_17vyp1dh_.arc
archived log for thread 1 with sequence 202 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_202_17vyvrm5_.arc
archived log for thread 1 with sequence 203 is already on disk as
file
/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_203_17vzdfwm_.arc
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_190_17vwvvt9
_.arc thread=1 sequence=190
archived log file

```
name=/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_191_17vx6vmg
_.arc thread=1 sequence=191
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_192_17vxctms
_.arc thread=1 sequence=192
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_193_17vxjjps
_.arc thread=1 sequence=193
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_194_17vxnxrh
_.arc thread=1 sequence=194
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_195_17vxswv5
_.arc thread=1 sequence=195
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_196_17vxyllwp
_.arc thread=1 sequence=196
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_197_17vy3cyw
_.arc thread=1 sequence=197
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_198_17vy8245
_.arc thread=1 sequence=198
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_199_17vydv4c
_.arc thread=1 sequence=199
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_200_17vykf23
_.arc thread=1 sequence=200
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_201_17vyp1dh
_.arc thread=1 sequence=201
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_202_17vyvrm5
_.arc thread=1 sequence=202
archived log file
name=/nfsfsxn/archlog/DB1/archivelog/2023_06_05/o1_mf_1_203_17vzdfwm
_.arc thread=1 sequence=203
media recovery complete, elapsed time: 00:19:30
Finished recover at 07-JUN-23

RMAN> exit

Recovery Manager complete.
[oracle@ip-172-30-15-124 dbs]$ sqlplus / as sysdba
```

SQL*Plus: Release 19.0.0.0.0 - Production on Wed Jun 7 15:58:12 2023
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 member from v\$logfile;

MEMBER

+DATA/DB1/ONLINELOG/group_3.264.1136666437
+DATA/DB1/ONLINELOG/group_2.263.1136666437
+DATA/DB1/ONLINELOG/group_1.262.1136666437

SQL> alter database rename file
'+DATA/DB1/ONLINELOG/group_1.262.1136666437' to
'/nfsfsxn/oracopy/redo01.log';

Database altered.

SQL> alter database rename file
'+DATA/DB1/ONLINELOG/group_2.263.1136666437' to
'/nfsfsxn/oracopy/redo02.log';

Database altered.

SQL> alter database rename file
'+DATA/DB1/ONLINELOG/group_3.264.1136666437' to
'/nfsfsxn/oracopy/redo03.log';

Database altered.

SQL> alter database noarchivelog;

Database altered.

SQL> alter database open resetlogs;

Database altered.

SQL> set lin 200;

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

NAME

```
-----
-----
/nfsfsxn/oracopy/DB1/FB864A929AEB79B9E053630F1EAC7046/datafile/o1_mf
_temp_l81bhgz6g_.tmp
/nfsfsxn/oracopy/DB1/FB867DA8C68C816EE053630F1EAC2BCF/datafile/o1_mf
_temp_l81bj16t_.tmp
/nfsfsxn/oracopy/DB1/FB867EA89ECF81C0E053630F1EACB901/datafile/o1_mf
_temp_l81bj135_.tmp
/nfsfsxn/oracopy/DB1/FB867F8A4D4F821CE053630F1EAC69CC/datafile/o1_mf
_temp_l81bj13g_.tmp
/nfsfsxn/oracopy/DB1/datafile/o1_mf_temp_l81bhwjg_.tmp
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-21_7jltkqk6
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-22_7kltkqk6
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-23_7lltkqk6
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-24_7mltkqk6
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-25_7nltkqrh
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-26_7oltkqrj
```

NAME

```
-----
-----
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-27_7pltkqrq
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-28_7qltkqs1
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-29_7rltkr32
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-30_7sltkr3a
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-31_7tltkr3i
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-32_7ultkr42
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-33_7vltkra6
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-34_80ltkram
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-35_81ltkrap
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-10_86ltkrgr
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-14_87ltkrhr
```

NAME

```
-----
-----
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-18_88ltkrhr
```

```

/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-3_831tkrd9
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-6_891tkrhr
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-13_8b1tkril
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-17_8c1tkril
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-1_821tkrb8
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-5_8d1tkril
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-9_8a1tkrhr
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-11_841tkrf2
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-15_8e1tkril
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-19_8f1tkrj4

```

NAME

```

-----
-----
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-4_851tkrgf
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-8_8g1tkrj7
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-12_8i1tkrj9
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-16_8j1tkrja
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-20_8k1tkrjb
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-7_8h1tkrj9
/nfsfsxn/oracopy/db1.ctl
/nfsfsxn/oracopy/redo01.log
/nfsfsxn/oracopy/redo02.log
/nfsfsxn/oracopy/redo03.log

```

43 rows selected.

SQL> show pdbs;

CON_ID	CON_NAME	OPEN MODE	RESTRICTED
2	PDB\$SEED	READ ONLY	NO
3	DB1_PDB1	READ WRITE	NO
4	DB1_PDB2	READ WRITE	NO
5	DB1_PDB3	READ WRITE	NO

SQL> alter session set container=db1_pdb1;

Session altered.

SQL> select * from test;

```

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

```

```

-----
1 18-MAY-23 02.35.37.000000 PM
test oracle incremental merge switch to copy
2 30-MAY-23 05.23.11.000000 PM
test recovery on a new EC2 instance host with image copy on FSxN
3 05-JUN-23 03.19.46.000000 PM
test clone on a new EC2 instance host with image copy on FSxN

SQL>

```

9. Rename the cloned database instance and change database ID with Oracle nid utility. The database instance state needs to be in mount to execute the command.

```

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

NAME          OPEN_MODE          LOG_MODE
-----
DB1           READ WRITE        NOARCHIVELOG

SQL> shutdown immediate;
Database closed.
Database dismounted.
ORACLE instance shut down.

SQL> startup mount;
ORACLE instance started.

Total System Global Area 1.0737E+10 bytes
Fixed Size                 9174800 bytes
Variable Size             1577058304 bytes
Database Buffers          9126805504 bytes
Redo Buffers              24379392 bytes
Database mounted.

SQL> exit
Disconnected from Oracle Database 19c Enterprise Edition Release
19.0.0.0.0 - Production
Version 19.18.0.0.0
[oracle@ip-172-30-15-124 dbs]$ nid target=/ dbname=db1tst

DBNEWID: Release 19.0.0.0.0 - Production on Wed Jun 7 16:15:14 2023

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

Connected to database DB1 (DBID=1730530050)

```

Connected to server version 19.18.0

Control Files in database:

/nfsfsxn/oracopy/db1.ctl

Change database ID and database name DB1 to DB1TST? (Y/[N]) => Y

Proceeding with operation

Changing database ID from 1730530050 to 3054879890

Changing database name from DB1 to DB1TST

Control File /nfsfsxn/oracopy/db1.ctl - modified

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-1_821tkrb - dbid changed, wrote new name

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-3_831tkrd - dbid changed, wrote new name

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-4_851tkrg - dbid changed, wrote new name

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-5_8d1tkri - dbid changed, wrote new name

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-6_891tkrh - dbid changed, wrote new name

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-7_8h1tkrj - dbid changed, wrote new name

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-8_8g1tkrj - dbid changed, wrote new name

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-9_8a1tkrh - dbid changed, wrote new name

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-10_861tkrg - dbid changed, wrote new name

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-11_841tkrf - dbid changed, wrote new name

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-12_8i1tkrj - dbid changed, wrote new name

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-13_8b1tkri - dbid changed, wrote new name

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-14_871tkrh - dbid changed, wrote new name

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-15_8e1tkri - dbid changed, wrote new name

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-16_8j1tkrj - dbid changed, wrote new name

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-17_8c1tkri - dbid changed, wrote new name

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-18_881tkrh - dbid changed, wrote new name

Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-

UNDOTBS1_FNO-19_8f1tkrj - dbid changed, wrote new name
Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-
20_8k1tkrj - dbid changed, wrote new name
Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
21_7j1tkqk - dbid changed, wrote new name
Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
22_7k1tkqk - dbid changed, wrote new name
Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
23_7l1tkqk - dbid changed, wrote new name
Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
24_7m1tkqk - dbid changed, wrote new name
Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
25_7n1tkqr - dbid changed, wrote new name
Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
26_7o1tkqr - dbid changed, wrote new name
Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
27_7p1tkqr - dbid changed, wrote new name
Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
28_7q1tkqs - dbid changed, wrote new name
Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
29_7r1tkr3 - dbid changed, wrote new name
Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
30_7s1tkr3 - dbid changed, wrote new name
Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
31_7t1tkr3 - dbid changed, wrote new name
Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
32_7ultkr4 - dbid changed, wrote new name
Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
33_7v1tkra - dbid changed, wrote new name
Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
34_801tkra - dbid changed, wrote new name
Datafile /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-
35_811tkra - dbid changed, wrote new name
Datafile /nfsfsxn/oracopy/DB1/datafile/o1_mf_temp_l81bhwjg_.tm -
dbid changed, wrote new name
Datafile
/nfsfsxn/oracopy/DB1/FB864A929AEB79B9E053630F1EAC7046/datafile/o1_mf
_temp_l81bh6g_.tm - dbid changed, wrote new name
Datafile
/nfsfsxn/oracopy/DB1/FB867DA8C68C816EE053630F1EAC2BCF/datafile/o1_mf
_temp_l81bj16t_.tm - dbid changed, wrote new name
Datafile
/nfsfsxn/oracopy/DB1/FB867EA89ECF81C0E053630F1EACB901/datafile/o1_mf
_temp_l81bj135_.tm - dbid changed, wrote new name
Datafile
/nfsfsxn/oracopy/DB1/FB867F8A4D4F821CE053630F1EAC69CC/datafile/o1_mf

```
_temp_l81bj13g_.tm - dbid changed, wrote new name
Control File /nfsfsxn/oracopy/db1.ctl - dbid changed, wrote new
name
Instance shut down

Database name changed to DB1TST.
Modify parameter file and generate a new password file before
restarting.
Database ID for database DB1TST changed to 3054879890.
All previous backups and archived redo logs for this database are
unusable.
Database is not aware of previous backups and archived logs in
Recovery Area.
Database has been shutdown, open database with RESETLOGS option.
Succesfully changed database name and ID.
DBNEWID - Completed succesfully.
```

10. Change Oracle database environment configuration to new database name or instance ID in oratab, init file, and create necessary admin directories that match with new instance ID. Then, start the instance with resetlogs option.

```
SQL> startup mount;
ORACLE instance started.
```

```
Total System Global Area 1.0737E+10 bytes
Fixed Size                  9174800 bytes
Variable Size               1577058304 bytes
Database Buffers           9126805504 bytes
Redo Buffers                 24379392 bytes
Database mounted.
```

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

```
Database altered.
```

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

NAME	OPEN_MODE	LOG_MODE
DB1TST	READ WRITE	NOARCHIVELOG

```
SQL> show pdbs
```

CON_ID	CON_NAME	OPEN MODE	RESTRICTED
2	PDB\$SEED	READ ONLY	NO
3	DB1_PDB1	MOUNTED	
4	DB1_PDB2	MOUNTED	
5	DB1_PDB3	MOUNTED	

```
SQL> alter pluggable database all open;
```

```
Pluggable database altered.
```

```
SQL> show pdbs
```

CON_ID	CON_NAME	OPEN MODE	RESTRICTED
2	PDB\$SEED	READ ONLY	NO
3	DB1_PDB1	READ WRITE	NO
4	DB1_PDB2	READ WRITE	NO
5	DB1_PDB3	READ WRITE	NO

```
SQL>
```

This completes the clone of a new Oracle instance from staging database copy on FSx NFS mount for DEV, UAT, or any other use cases. Multiple Oracle instances can be cloned off the same staging image copy.



If you run into error RMAN-06571: datafile 1 does not have recoverable copy when switching the database to copy, check database incarnation that matches with primary production DB. If needed, reset the incarnation to match with primary with RMAN command `reset database to incarnation n;`

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>

- Amazon FSx for NetApp ONTAP

<https://aws.amazon.com/fsx/netapp-ontap/>

- Amazon EC2

https://aws.amazon.com/pm/ec2/?trk=36c6da98-7b20-48fa-8225-4784bced9843&sc_channel=ps&s_kwcid=ALi4422!3!467723097970!e!!g!!aws%20ec2&ef_id=Cj0KCQiA54KfBhCKARIsAJzSrdqwQrghn6l71jiWzSeaT9Uh1-vY-VfhJixF-xnv5rWwn2S7RqZOTQ0aAh7eEALw_wcB:G:s&s_kwcid=ALi4422!3!467723097970!e!!g!!aws%20ec2

TR-4974: Oracle 19c in Standalone Restart on AWS FSx/EC2 with NFS/ASM

Allen Cao, Niyaz Mohamed, NetApp

This solution provides overview and details for Oracle database deployment and protection in AWS FSx ONTAP storage and EC2 compute instance with NFS protocol and Oracle database configured in standalone ReStart using asm as volume manager.

Purpose

ASM (Automatic Storage Management) is a popular Oracle storage volume manager that is employed in many Oracle installations. It is also Oracle's recommended storage management solution. It provides an alternative to conventional volume managers and file systems. Since Oracle version 11g, ASM has been packaged with grid infrastructure rather than a database. As a result, in order to utilize Oracle ASM for storage management without RAC, you must install Oracle grid infrastructure in a standalone server, also known as Oracle Restart. Doing so certainly adds more complexity in an otherwise simpler Oracle database deployment. However, as the name implies, when Oracle is deployed in Restart mode, any failed Oracle services are restarted after a host reboot without user intervention, which provides a certain degree of high availability or HA functionality.

Oracle ASM is generally deployed in FC, iSCSI storage protocols and luns as raw storage devices. However, ASM on NFS protocol and NFS file system is also supported configuration by Oracle. In this documentation, we demonstrate how to deploy an Oracle 19c database with the NFS protocol and Oracle ASM in an Amazon FSx for ONTAP storage environment with EC2 compute instances. We also demonstrate how to use the

NetApp SnapCenter service through the NetApp BlueXP console to backup, restore, and clone your Oracle database for dev/test or other use cases for storage-efficient database operation in the AWS public cloud.

This solution addresses the following use cases:

- Oracle database deployment in Amazon FSx for ONTAP storage and EC2 compute instances with NFS/ASM
- Testing and validating an Oracle workload in the public AWS cloud with NFS/ASM
- Testing and validating Oracle database Restart functionalities deployed in AWS

Audience

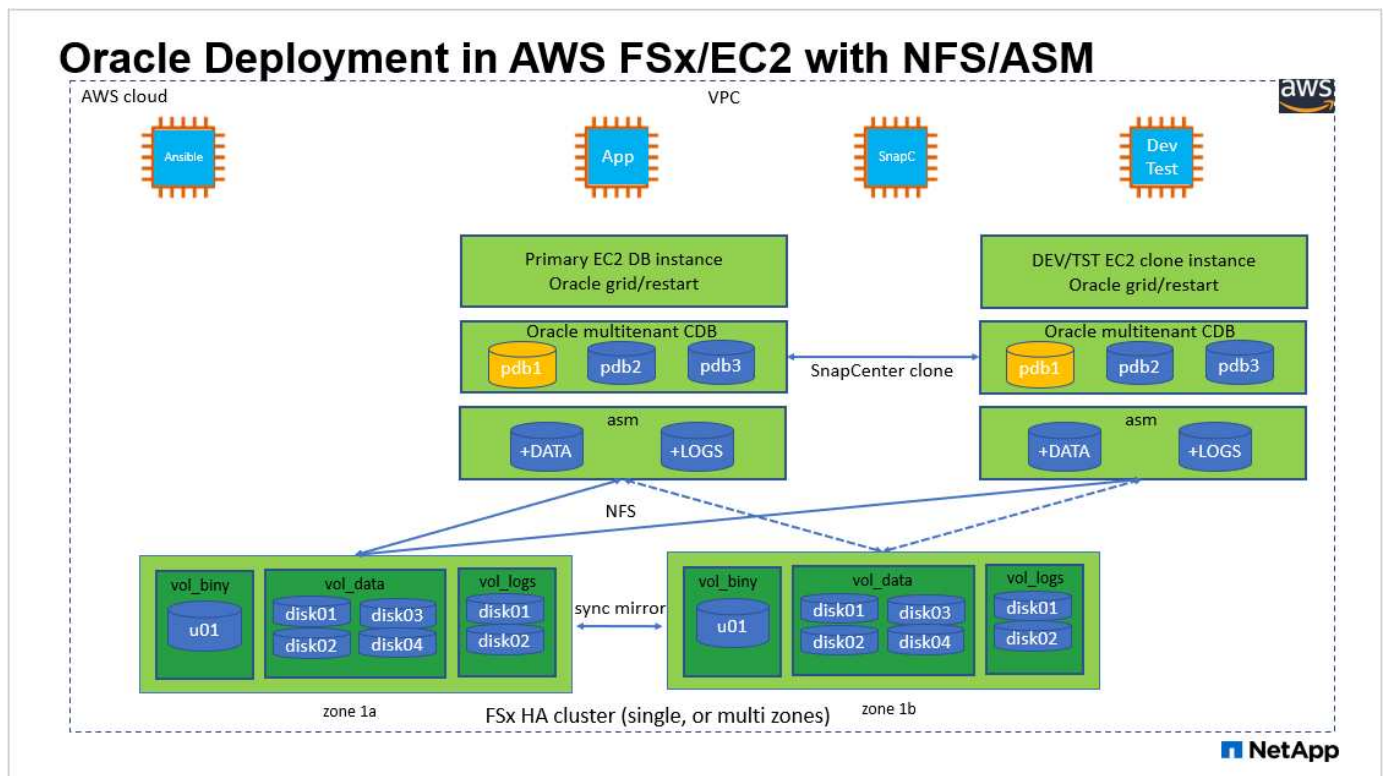
This solution is intended for the following people:

- A DBA who would like to deploy Oracle in an AWS public cloud with NFS/ASM.
- A database solution architect who would like to test Oracle workloads in the AWS public cloud.
- The storage administrator who would like to deploy and manage an Oracle database deployed to AWS FSx storage.
- The application owner who would like to stand up an Oracle database in AWS FSx/EC2.

Solution test and validation environment

The testing and validation of this solution was performed in an AWS FSx and EC2 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		
FSx ONTAP storage	Current version offered by AWS	One FSx HA cluster in the same VPC and availability zone
EC2 instance for compute	t2.xlarge/4vCPU/16G	Two EC2 T2 xlarge EC2 instances, one as primary DB server and the other as a clone DB server
Software		
RedHat Linux	RHEL-8.6.0_HVM-20220503-x86_64-2-Hourly2-GP2	Deployed RedHat subscription for testing
Oracle Grid Infrastructure	Version 19.18	Applied RU patch p34762026_190000_Linux-x86-64.zip
Oracle Database	Version 19.18	Applied RU patch p34765931_190000_Linux-x86-64.zip
Oracle OPatch	Version 12.2.0.1.36	Latest patch p6880880_190000_Linux-x86-64.zip
SnapCenter Service	Version	v2.3.1.2324

Key factors for deployment consideration

- **EC2 compute instances.** In these tests and validations, we used an AWS EC2 t2.xlarge instance type for the Oracle database compute instance. NetApp recommends using an M5 type EC2 instance as the compute instance for Oracle in production deployment because it is optimized for database workloads. You need to size the EC2 instance appropriately for the number of vCPUs and the amount of RAM based on actual workload requirements.
- **FSx storage HA clusters single- or multi-zone deployment.** In these tests and validations, we deployed an FSx HA cluster in a single AWS availability zone. For production deployment, NetApp recommends deploying an FSx HA pair in two different availability zones. An FSx HA cluster is always provisioned in a HA pair that is sync mirrored in a pair of active-passive file systems to provide storage-level redundancy. Multi-zone deployment further enhances high availability in the event of failure in a single AWS zone.
- **FSx storage cluster sizing.** An Amazon FSx for ONTAP storage file system provides up to 160,000 raw SSD IOPS, up to 4GBps throughput, and a maximum of 192TiB capacity. However, you can size the cluster in terms of provisioned IOPS, throughput, and the storage limit (minimum 1,024 GiB) based on your actual requirements at the time of deployment. The capacity can be adjusted dynamically on the fly without affecting application availability.
- **Oracle data and logs layout.** In our tests and validations, we deployed two ASM disk groups for data and logs respectively. Within the +DATA asm disk group, we provisioned four disks in a data NFS file system mount point. Within the +LOGS asm disk group, we provisioned two disks in a logs NFS file system mount point. For large database deployment, ASM disk groups can be built to span multiple FSx file systems with ASM NFS disks distributed through multiple NFS mount points anchored on FSx file systems. This particular setup is designed to meet database throughput over 4GBps throughput and 160,000 raw SSD IOPS requirement.
- **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 FSx file systems deployment for large database, dNFS multi-path should be properly configured.

- **Oracle ASM redundancy level to use for each Oracle ASM disk group that you create.** Because FSx already mirrors the storage on the FSx cluster level, you should **ONLY** use External Redundancy, which means that the option does not allow Oracle ASM to mirror the contents of the disk group. This is particularly important as NFS for Oracle database data storage requires HARD NFS mount option which is **NOT** desirable for mirroring ASM contents on the Oracle level.
- **Database backup.** NetApp provides a SaaS version of SnapCenter software service for database backup, restore, and clone in the cloud that is available through the NetApp BlueXP console UI. NetApp recommends implementing such a service to achieve fast (under a minute) SnapShot backup, quick (few minutes) database restore, and database cloning.

Solution deployment

The following section provides step-by-step deployment procedures.

Prerequisites for deployment

Deployment requires the following prerequisites.

1. An AWS account has been set up, and the necessary VPC and network segments have been created within your AWS account.
2. From the AWS EC2 console, you must deploy two EC2 Linux instances, one as the primary Oracle DB server and an optional alternative clone target DB server. See the architecture diagram in the previous section for more details about the environment setup. Also review the [User Guide for Linux instances](#) for more information.
3. From the AWS EC2 console, deploy Amazon FSx for ONTAP storage HA clusters to host the Oracle database volumes. If you are not familiar with the deployment of FSx storage, see the documentation [Creating FSx for ONTAP file systems](#) for step-by-step instructions.
4. Steps 2 and 3 can be performed using the following Terraform automation toolkit, which creates an EC2 instance named `ora_01` and an FSx file system named `fsx_01`. Review the instruction carefully and change the variables to suit your environment before execution.

```
git clone https://github.com/NetApp-  
Automation/na_aws_fsx_ec2_deploy.git
```



Ensure that you have allocated at least 50G in EC2 instance root volume in order to have sufficient space to stage Oracle installation files.

EC2 instance kernel configuration

With the prerequisites provisioned, log into the EC2 instance as `ec2-user` and `sudo` to root user to configure the Linux kernel for Oracle installation.

1. Create a staging directory `/tmp/archive` folder and set the `777` permission.

```
mkdir /tmp/archive  
  
chmod 777 /tmp/archive
```

2. Download and stage the Oracle binary installation files and other required rpm files to the `/tmp/archive` directory.

See the following list of installation files to be staged in `/tmp/archive` on the EC2 instance.

```
[ec2-user@ip-172-30-15-58 ~]$ ls -l /tmp/archive  
total 10537316  
-rw-rw-r--. 1 ec2-user ec2-user      19112 Mar 21 15:57 compat-  
libcap1-1.10-7.el7.x86_64.rpm  
-rw-rw-r--  1 ec2-user ec2-user 3059705302 Mar 21 22:01  
LINUX.X64_193000_db_home.zip  
-rw-rw-r--  1 ec2-user ec2-user 2889184573 Mar 21 21:09  
LINUX.X64_193000_grid_home.zip  
-rw-rw-r--. 1 ec2-user ec2-user      589145 Mar 21 15:56  
netapp_linux_unified_host_utilities-7-1.x86_64.rpm  
-rw-rw-r--. 1 ec2-user ec2-user      31828 Mar 21 15:55 oracle-  
database-preinstall-19c-1.0-2.el8.x86_64.rpm  
-rw-rw-r--  1 ec2-user ec2-user 2872741741 Mar 21 22:31  
p34762026_190000_Linux-x86-64.zip  
-rw-rw-r--  1 ec2-user ec2-user 1843577895 Mar 21 22:32  
p34765931_190000_Linux-x86-64.zip  
-rw-rw-r--  1 ec2-user ec2-user  124347218 Mar 21 22:33  
p6880880_190000_Linux-x86-64.zip  
-rw-r--r--  1 ec2-user ec2-user    257136 Mar 22 16:25  
policycoreutils-python-utils-2.9-9.el8.noarch.rpm
```

3. Install Oracle 19c preinstall RPM, which satisfies most kernel configuration requirements.

```
yum install /tmp/archive/oracle-database-preinstall-19c-1.0-  
2.el8.x86_64.rpm
```

4. Download and install the missing `compat-libcap1` in Linux 8.


```
yum install /tmp/archive/compat-libcap1-1.10-7.el7.x86_64.rpm
```

5. From NetApp, download and install NetApp host utilities.

```
yum install /tmp/archive/netapp_linux_unified_host_utilities-7-1.x86_64.rpm
```

6. Install `policycoreutils-python-utils`, which is not available in the EC2 instance.

```
yum install /tmp/archive/policycoreutils-python-utils-2.9-9.el8.noarch.rpm
```

7. Install open JDK version 1.8.

```
yum install java-1.8.0-openjdk.x86_64
```

8. Install `nfs-utils`.

```
yum install nfs-utils
```

9. Disable transparent hugepages in the current system.

```
echo never > /sys/kernel/mm/transparent_hugepage/enabled  
echo never > /sys/kernel/mm/transparent_hugepage/defrag
```

Add the following lines in `/etc/rc.local` to disable `transparent_hugepage` after reboot:

```
# Disable transparent hugepages  
if test -f /sys/kernel/mm/transparent_hugepage/enabled;  
then  
    echo never > /sys/kernel/mm/transparent_hugepage/enabled  
fi  
if test -f /sys/kernel/mm/transparent_hugepage/defrag;  
then  
    echo never > /sys/kernel/mm/transparent_hugepage/defrag  
fi
```

10. Disable selinux by changing `SELINUX=enforcing` to `SELINUX=disabled`. You must reboot the host to make the change effective.

```
vi /etc/sysconfig/selinux
```

11. Add the following lines to `limit.conf` to set the file descriptor limit and stack size without quotes "`"`.

```
vi /etc/security/limits.conf
**                hard    nofile           65536"
**                soft    stack           10240"
```

12. Add swap space to EC2 instance by following this instruction: [How do I allocate memory to work as swap space in an Amazon EC2 instance by using a swap file?](#) The exact amount of space to add depends on the size of RAM up to 16G.
13. Add the ASM group to be used for the asm sysasm group

```
groupadd asm
```

14. Modify the oracle user to add ASM as a secondary group (the oracle user should have been created after Oracle preinstall RPM installation).

```
usermod -a -G asm oracle
```

15. Reboot the EC2 instance.

Provision and export NFS volumes to be mounted to EC2 instance host

Provision three volumes from the command line by login to FSx cluster via ssh as fsxadmin user with FSx cluster management IP to host the Oracle database binary, data, and logs files.

1. Log into the FSx cluster through SSH as the fsxadmin user.

```
ssh fsxadmin@172.30.15.53
```

2. Execute the following command to create a volume for the Oracle binary.

```
vol create -volume ora_01_biny -aggregate aggr1 -size 50G -state  
online -type RW -junction-path /ora_01_biny -snapshot-policy none  
-tiering-policy snapshot-only
```

3. Execute the following command to create a volume for Oracle data.

```
vol create -volume ora_01_data -aggregate aggr1 -size 100G -state  
online -type RW -junction-path /ora_01_data -snapshot-policy none  
-tiering-policy snapshot-only
```

4. Execute the following command to create a volume for Oracle logs.

```
vol create -volume ora_01_logs -aggregate aggr1 -size 100G -state  
online -type RW -junction-path /ora_01_logs -snapshot-policy none  
-tiering-policy snapshot-only
```

5. Validate the DB volumes created.

```
vol show
```

This is expected to return:

```

FsxId02ad7bf3476b741df::> vol show
(vol show)
FsxId06c3c8b2a7bd56458::> vol show
Vserver      Volume      Aggregate   State      Type      Size
Available  Used%
-----  -----  -----  -----  -----  -----
svm_ora     ora_01_biny  aggr1      online    RW        50GB
47.50GB     0%
svm_ora     ora_01_data  aggr1      online    RW        100GB
95.00GB     0%
svm_ora     ora_01_logs  aggr1      online    RW        100GB
95.00GB     0%
svm_ora     svm_ora_root aggr1      online    RW        1GB
972.1MB     0%
4 entries were displayed.

```

Database storage configuration

Now, import and set up the FSx storage for the Oracle grid infrastructure and database installation on the EC2 instance host.

1. Log into the EC2 instance via SSH as the ec2-user with your SSH key and EC2 instance IP address.

```
ssh -i ora_01.pem ec2-user@172.30.15.58
```

2. Create /u01 directory to mount Oracle binary file system

```
sudo mkdir /u01
```

3. Mount the binary volume to /u01, changed to your FSx NFS lif IP address. If you deployed FSx cluster via NetApp automation toolkit, FSx virtual storage server NFS lif IP address will be listed in the output at the end of resources provision execution. Otherwise, it can be retrieved from AWS FSx console UI.

```
sudo mount -t nfs 172.30.15.19:/ora_01_biny /u01 -o  
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsiz=65536
```

4. Change /u01 mount point ownership to the Oracle user and it's associated primary group.

```
sudo chown oracle:oinstall /u01
```

5. Create /oradata directory to mount Oracle data file system

```
sudo mkdir /oradata
```

6. Mount the data volume to /oradata, changed to your FSx NFS lif IP address

```
sudo mount -t nfs 172.30.15.19:/ora_01_data /oradata -o  
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsiz=65536
```

7. Change /oradata mount point ownership to the Oracle user and it's associated primary group.

```
sudo chown oracle:oinstall /oradata
```

8. Create /orlogs directory to mount Oracle logs file system

```
sudo mkdir /orlogs
```

9. Mount the log volume to /oratalogs, changed to your FSx NFS lif IP address

```
sudo mount -t nfs 172.30.15.19:/ora_01_logs /oratalogs -o  
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsiz=65536,wsiz=65536
```

10. Change /oratalogs mount point ownership to the Oracle user and it's associated primary group.

```
sudo chown oracle:oinstall /oratalogs
```

11. Add a mount point to /etc/fstab.

```
sudo vi /etc/fstab
```

Add the following line.

```
172.30.15.19:/ora_01_biny      /u01          nfs  
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsiz=65536,wsiz=65536  0  
0  
172.30.15.19:/ora_01_data     /oradata      nfs  
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsiz=65536,wsiz=65536  0  
0  
172.30.15.19:/ora_01_logs     /oratalogs    nfs  
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsiz=65536,wsiz=65536  0  
0
```

12. sudo to oracle user, create asm folders to store asm disk files

```
sudo su  
su - oracle  
mkdir /oradata/asm  
mkdir /oratalogs/asm
```

13. As the oracle user, create asm data disk files, change the count to match to the disk size with block size.

```
dd if=/dev/zero of=/oradata/asm/nfs_data_disk01 bs=1M count=20480
oflag=direct
dd if=/dev/zero of=/oradata/asm/nfs_data_disk02 bs=1M count=20480
oflag=direct
dd if=/dev/zero of=/oradata/asm/nfs_data_disk03 bs=1M count=20480
oflag=direct
dd if=/dev/zero of=/oradata/asm/nfs_data_disk04 bs=1M count=20480
oflag=direct
```

14. As the root user, change data disk file permission to 640

```
chmod 640 /oradata/asm/*
```

15. AS the oracle user, create asm logs disk files, change to count to match to the disk size with block size.

```
dd if=/dev/zero of=/oralogs/asm/nfs_logs_disk01 bs=1M count=40960
oflag=direct
dd if=/dev/zero of=/oralogs/asm/nfs_logs_disk02 bs=1M count=40960
oflag=direct
```

16. As the root user, change logs disk file permission to 640

```
chmod 640 /oralogs/asm/*
```

17. Reboot the EC2 instance host.

Oracle grid infrastructure installation

1. Log into the EC2 instance as the ec2-user via SSH and enable password authentication by uncommenting `PasswordAuthentication yes` and then commenting out `PasswordAuthentication no`.

```
sudo vi /etc/ssh/sshd_config
```

2. Restart the sshd service.

```
sudo systemctl restart sshd
```

3. Reset the Oracle user password.

```
sudo passwd oracle
```

4. Log in as the Oracle Restart software owner user (oracle). Create an Oracle directory as follows:

```
mkdir -p /u01/app/oracle  
mkdir -p /u01/app/oraInventory
```

5. Change the directory permission setting.

```
chmod -R 775 /u01/app
```

6. Create a grid home directory and change to it.

```
mkdir -p /u01/app/oracle/product/19.0.0/grid  
cd /u01/app/oracle/product/19.0.0/grid
```

7. Unzip the grid installation files.

```
unzip -q /tmp/archive/LINUX.X64_193000_grid_home.zip
```

8. From grid home, delete the OPatch directory.

```
rm -rf OPatch
```

9. From grid home, copy `p6880880_190000_Linux-x86-64.zip` to the `grid_home`, and then unzip it.


```
cp /tmp/archive/p6880880_190000_Linux-x86-64.zip .
unzip p6880880_190000_Linux-x86-64.zip
```

10. From grid home, revise `cv/admin/cvu_config`, uncomment and replace `CV_ASSUME_DISTID=OEL5` with `CV_ASSUME_DISTID=OL7`.

```
vi cv/admin/cvu_config
```

11. Prepare a `gridsetup.rsp` file for silent installation and place the `rsp` file in the `/tmp/archive` directory. The `rsp` file should cover sections A, B, and G with the following information:

```
INVENTORY_LOCATION=/u01/app/oraInventory
oracle.install.option=HA_CONFIG
ORACLE_BASE=/u01/app/oracle
oracle.install.asm.OSDBA=dba
oracle.install.asm.OSOPER=oper
oracle.install.asm.OSASM=asm
oracle.install.asm.SYSASMPassword="SetPWD"
oracle.install.asm.diskGroup.name=DATA
oracle.install.asm.diskGroup.redundancy=EXTERNAL
oracle.install.asm.diskGroup.AUSize=4
oracle.install.asm.diskGroup.disks=/oradata/asm/*,/orlogs/asm/*
oracle.install.asm.diskGroup.diskDiscoveryString=/oradata/asm/nfs_data_
data_disk01,/oradata/asm/nfs_data_disk02,/oradata/asm/nfs_data_disk03,
/oradata/asm/nfs_data_disk04
oracle.install.asm.monitorPassword="SetPWD"
oracle.install.asm.configureAFD=false
```

12. Log into the EC2 instance as the root user.

13. Install `cvuqdisk-1.0.10-1.rpm`.

```
rpm -ivh /u01/app/oracle/product/19.0.0/grid/cv/rpm/cvuqdisk-1.0.10-
1.rpm
```

14. Log into the EC2 instance as the Oracle user and extract the patch in the `/tmp/archive` folder.

```
unzip p34762026_190000_Linux-x86-64.zip
```

15. From grid home `/u01/app/oracle/product/19.0.0/grid` and as the oracle user, launch `gridSetup.sh` for grid infrastructure installation.

```
./gridSetup.sh -applyRU /tmp/archive/34762026/ -silent  
-responseFile /tmp/archive/gridsetup.rsp
```

Ignore the warnings about wrong groups for grid infrastructure. We are using a single Oracle user to manage Oracle Restart, so this is expected.

16. As root user, execute the following script(s):

```
/u01/app/oraInventory/orainstRoot.sh  
  
/u01/app/oracle/product/19.0.0/grid/root.sh
```

17. As the Oracle user, execute the following command to complete the configuration:

```
/u01/app/oracle/product/19.0.0/grid/gridSetup.sh -executeConfigTools  
-responseFile /tmp/archive/gridsetup.rsp -silent
```

18. As the Oracle user, create the LOGS disk group.

```
bin/asmca -silent -sysAsmPassword 'yourPWD' -asmsnmpPassword  
'yourPWD' -createDiskGroup -diskGroupName LOGS -disk  
'/oralogs/asm/nfs_logs_disk*' -redundancy EXTERNAL -au_size 4
```

19. As the Oracle user, validate grid services after installation configuration.

```

bin/crsctl stat res -t
+
Name                Target  State      Server
State details
Local Resources
ora.DATA.dg         ONLINE ONLINE     ip-172-30-15-58
STABLE
ora.LISTENER.lsnr   ONLINE ONLINE     ip-172-30-15-58
STABLE
ora.LOGS.dg         ONLINE ONLINE     ip-172-30-15-58
STABLE
ora.asm             ONLINE ONLINE     ip-172-30-15-58
Started,STABLE
ora.ons             OFFLINE OFFLINE    ip-172-30-15-58
STABLE
Cluster Resources
ora.cssd            ONLINE ONLINE     ip-172-30-15-58
STABLE
ora.diskmon         OFFLINE OFFLINE
STABLE
ora.driver.afd      ONLINE ONLINE     ip-172-30-15-58
STABLE
ora.evmd            ONLINE ONLINE     ip-172-30-15-58
STABLE

```

Oracle database installation

1. Log in as the Oracle user and unset `$ORACLE_HOME` and `$ORACLE_SID` if it is set.

```
unset ORACLE_HOME
unset ORACLE_SID
```

2. Create the Oracle DB home directory and change to it.

```
mkdir /u01/app/oracle/product/19.0.0/db1
cd /u01/app/oracle/product/19.0.0/db1
```

3. Unzip the Oracle DB installation files.

```
unzip -q /tmp/archive/LINUX.X64_193000_db_home.zip
```

4. From the DB home, delete the OPatch directory.

```
rm -rf OPatch
```

5. From DB home, copy `p6880880_190000_Linux-x86-64.zip` to `grid_home`, and then unzip it.

```
cp /tmp/archive/p6880880_190000_Linux-x86-64.zip .
unzip p6880880_190000_Linux-x86-64.zip
```

6. From DB home, revise `cv/admin/cvu_config`, and uncomment and replace `CV_ASSUME_DISTID=OEL5` with `CV_ASSUME_DISTID=OL7`.

```
vi cv/admin/cvu_config
```

7. From the `/tmp/archive` directory, unpack the DB 19.18 RU patch.

```
unzip p34765931_190000_Linux-x86-64.zip
```

8. Prepare the DB silent install `rsp` file in `/tmp/archive/dbinstall.rsp` directory with the following values:

```
oracle.install.option=INSTALL_DB_SWONLY
UNIX_GROUP_NAME=oinstall
INVENTORY_LOCATION=/u01/app/oraInventory
ORACLE_HOME=/u01/app/oracle/product/19.0.0/db1
ORACLE_BASE=/u01/app/oracle
oracle.install.db.InstallEdition=EE
oracle.install.db.OSDBA_GROUP=dba
oracle.install.db.OSOPER_GROUP=oper
oracle.install.db.OSBACKUPDBA_GROUP=oper
oracle.install.db.OSDGDBA_GROUP=dba
oracle.install.db.OSKMDBA_GROUP=dba
oracle.install.db.OSRACDBA_GROUP=dba
oracle.install.db.rootconfig.executeRootScript=false
```

9. From db1 home /u01/app/oracle/product/19.0.0/db1, execute silent software-only DB installation.

```
./runInstaller -applyRU /tmp/archive/34765931/ -silent
-ignorePrereqFailure -responseFile /tmp/archive/dbinstall.rsp
```

10. As root user, run the `root.sh` script after software-only installation.

```
/u01/app/oracle/product/19.0.0/db1/root.sh
```

11. As Oracle user, create the `dbca.rsp` file with the following entries:

```
gdbName=db1.demo.netapp.com
sid=db1
createAsContainerDatabase=true
numberOfPDBs=3
pdbName=db1_pdb
useLocalUndoForPDBs=true
pdbAdminPassword="yourPWD"
templateName=General_Purpose.dbc
sysPassword="yourPWD"
systemPassword="yourPWD"
dbsnmpPassword="yourPWD"
storageType=ASM
diskGroupName=DATA
characterSet=AL32UTF8
nationalCharacterSet=AL16UTF16
listeners=LISTENER
databaseType=MULTIPURPOSE
automaticMemoryManagement=false
totalMemory=8192
```



Set the total memory based on available memory in EC2 instance host. Oracle allocates 75% of `totalMemory` to DB instance SGA or buffer cache.

12. As Oracle user, launch DB creation with dbca.

```
bin/dbca -silent -createDatabase -responseFile /tmp/archive/dbca.rsp
```

output:

Prepare for db operation

7% complete

Registering database with Oracle Restart

11% complete

Copying database files

33% complete

Creating and starting Oracle instance

35% complete

38% complete

42% complete

45% complete

48% complete

Completing Database Creation

53% complete

55% complete

56% complete

Creating Pluggable Databases

60% complete

64% complete

69% complete

78% complete

Executing Post Configuration Actions

100% complete

Database creation complete. For details check the logfiles at:

/u01/app/oracle/cfgtoollogs/dbca/db1.

Database Information:

Global Database Name:db1.demo.netapp.com

System Identifier(SID):db1

Look at the log file "/u01/app/oracle/cfgtoollogs/dbca/db1/db1.log"
for further details.

13. As Oracle user, validate Oracle Restart HA services after DB creation.

```

[oracle@ip-172-30-15-58 db1]$ ../grid/bin/crsctl stat res -t
-----
-----
Name          Target  State          Server          State
details
-----
-----
Local Resources
-----
-----
ora.DATA.dg
          ONLINE  ONLINE        ip-172-30-15-58  STABLE
ora.LISTENER.lsnr
          ONLINE  ONLINE        ip-172-30-15-58  STABLE
ora.LOGS.dg
          ONLINE  ONLINE        ip-172-30-15-58  STABLE
ora.asm
          ONLINE  ONLINE        ip-172-30-15-58
Started,STABLE
ora.ons
          OFFLINE OFFLINE        ip-172-30-15-58  STABLE
-----
-----
Cluster Resources
-----
-----
ora.cssd
   1      ONLINE  ONLINE        ip-172-30-15-58  STABLE
ora.dbf.db
   1      ONLINE  ONLINE        ip-172-30-15-58
Open,HOME=/u01/app/o
racle/product/19.0.0
/db1,STABLE
ora.diskmon
   1      OFFLINE OFFLINE        STABLE
ora.evmd
   1      ONLINE  ONLINE        ip-172-30-15-58  STABLE
-----
-----
[oracle@ip-172-30-15-58 db1]$

```

14. Set the Oracle user `.bash_profile`.


```
vi ~/.bash_profile
```

15. Add following entries:

```
export ORACLE_HOME=/u01/app/oracle/product/19.0.0/db1
export ORACLE_SID=db1
export PATH=$PATH:$ORACLE_HOME/bin
alias asm='export
ORACLE_HOME=/u01/app/oracle/product/19.0.0/grid;export
ORACLE_SID=+ASM;export PATH=$PATH:$ORACLE_HOME/bin'
```

16. Validate the CDB/PDB created.

```
. ~/.bash_profile

sqlplus / as sysdba

SQL> select name, open_mode from v$database;

NAME          OPEN_MODE

DB1           READ WRITE

SQL> select name from v$datafile;

NAME

+DATA/DB1/DATAFILE/system.256.1132176177
+DATA/DB1/DATAFILE/sysaux.257.1132176221
+DATA/DB1/DATAFILE/undotbs1.258.1132176247
+DATA/DB1/86B637B62FE07A65E053F706E80A27CA/DATAFILE/system.265.11321
77009
+DATA/DB1/86B637B62FE07A65E053F706E80A27CA/DATAFILE/sysaux.266.11321
77009
+DATA/DB1/DATAFILE/users.259.1132176247
+DATA/DB1/86B637B62FE07A65E053F706E80A27CA/DATAFILE/undotbs1.267.113
2177009
+DATA/DB1/F7852758DCD6B800E0533A0F1EAC1DC6/DATAFILE/system.271.11321
77853
+DATA/DB1/F7852758DCD6B800E0533A0F1EAC1DC6/DATAFILE/sysaux.272.11321
77853
+DATA/DB1/F7852758DCD6B800E0533A0F1EAC1DC6/DATAFILE/undotbs1.270.113
2177853
+DATA/DB1/F7852758DCD6B800E0533A0F1EAC1DC6/DATAFILE/users.274.113217
```

```
7871
```

```
NAME
```

```
+DATA/DB1/F785288BBCD1BA78E0533A0F1EACCD6F/DATAFILE/system.276.11321
77871
+DATA/DB1/F785288BBCD1BA78E0533A0F1EACCD6F/DATAFILE/sysaux.277.11321
77871
+DATA/DB1/F785288BBCD1BA78E0533A0F1EACCD6F/DATAFILE/undotbs1.275.113
2177871
+DATA/DB1/F785288BBCD1BA78E0533A0F1EACCD6F/DATAFILE/users.279.113217
7889
+DATA/DB1/F78529A14DD8BB18E0533A0F1EACB8ED/DATAFILE/system.281.11321
77889
+DATA/DB1/F78529A14DD8BB18E0533A0F1EACB8ED/DATAFILE/sysaux.282.11321
77889
+DATA/DB1/F78529A14DD8BB18E0533A0F1EACB8ED/DATAFILE/undotbs1.280.113
2177889
+DATA/DB1/F78529A14DD8BB18E0533A0F1EACB8ED/DATAFILE/users.284.113217
7907
```

```
19 rows selected.
```

```
SQL> show pdbs
```

CON_ID	CON_NAME	OPEN MODE	RESTRICTED
2	PDB\$SEED	READ ONLY	NO
3	DB1_PDB1	READ WRITE	NO
4	DB1_PDB2	READ WRITE	NO
5	DB1_PDB3	READ WRITE	NO

```
SQL>
```

17. As oracle user, change to Oracle database home directory /u01/app/oracle/product/19.0.0/db1 and Enable dNFS

```
cd /u01/app/oracle/product/19.0.0/db1

mkdir rdbms/lib/odm

cp lib/libnfsodm19.so rdbms/lib/odm/
```

18. Configure oranfstab file in ORACLE_HOME

```
vi $ORACLE_HOME/dbs/oranfstab

add following entries:

server: fsx_01
local: 172.30.15.58 path: 172.30.15.19
nfs_version: nfsv3
export: /ora_01_biny mount: /u01
export: /ora_01_data mount: /oradata
export: /ora_01_logs mount: /orlogs
```

19. As oracle user, login to database from sqlplus and set the DB recovery size and location to the +LOGS disk group.

```
. ~/.bash_profile

sqlplus / as sysdba

alter system set db_recovery_file_dest_size = 80G scope=both;

alter system set db_recovery_file_dest = '+LOGS' scope=both;
```

20. Enable archive log mode and reboot Oracle DB instance

```
shutdown immediate;

startup mount;

alter database archivelog;

alter database open;

alter system switch logfile;
```

21. Validate DB log mode and dNFS after instance reboot

```

SQL> select name, log_mode from v$database;

NAME          LOG_MODE
-----
DB1           ARCHIVELOG

SQL> select svrname, dirname from v$dnfs_servers;

SVRNAME
-----
-----
DIRNAME
-----
-----
fsx_01
/ora_01_data

fsx_01
/ora_01_biny

fsx_01
/ora_01_logs

```

22. Validate Oracle ASM

```

[oracle@ip-172-30-15-58 db1]$ asm
[oracle@ip-172-30-15-58 db1]$ sqlplus / as sysasm

SQL*Plus: Release 19.0.0.0.0 - Production on Tue May 9 20:39:39 2023
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> set lin 200
SQL> col path form a30
SQL> select name, path, header_status, mount_status, state from
v$asm_disk;

NAME          PATH

```

```

HEADER_STATU MOUNT_S STATE
-----
-----
DATA_0002          /oradata/asm/nfs_data_disk01  MEMBER
  CACHED  NORMAL
DATA_0000          /oradata/asm/nfs_data_disk02  MEMBER
  CACHED  NORMAL
DATA_0001          /oradata/asm/nfs_data_disk03  MEMBER
  CACHED  NORMAL
DATA_0003          /oradata/asm/nfs_data_disk04  MEMBER
  CACHED  NORMAL
LOGS_0000          /orlogs/asm/nfs_logs_disk01   MEMBER
  CACHED  NORMAL
LOGS_0001          /orlogs/asm/nfs_logs_disk02   MEMBER
  CACHED  NORMAL

```

6 rows selected.

```

SQL> select name, state, ALLOCATION_UNIT_SIZE, TOTAL_MB, FREE_MB
from v$asm_diskgroup;

```

```

NAME                STATE           ALLOCATION_UNIT_SIZE
TOTAL_MB    FREE_MB
-----
DATA                MOUNTED                4194304
81920          73536
LOGS                MOUNTED                4194304
81920          81640

```

This completes Oracle 19c version 19.18 Restart deployment on an Amazon FSx for ONTAP and EC2 compute instance with NFS/ASM. If desired, NetApp recommends relocating the Oracle control file and online log files to the +LOGS disk group.

Automated deployment option

NetApp will release a fully automated solution deployment toolkit with Ansible to facilitate the implementation of this solution. Please check back for the availability of the toolkit. After it is released, a link will be posted here.

Oracle Database backup, restore, and clone with SnapCenter Service

At this moment, Oracle database with NFS and ASM storage option is only supported by traditional SnapCenter Server UI tool See [Hybrid Cloud Database Solutions with SnapCenter](#) for details on Oracle database backup, restore, and clone with NetApp SnapCenter UI tool.

Where to find additional information

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

- Installing Oracle Grid Infrastructure for a Standalone Server with a New Database Installation

<https://docs.oracle.com/en/database/oracle/oracle-database/19/ladbi/installing-oracle-grid-infrastructure-for-a-standalone-server-with-a-new-database-installation.html#GUID-0B1CEE8C-C893-46AA-8A6A-7B5FAAEC72B3>

- Installing and Configuring Oracle Database Using Response Files

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

- Amazon FSx for NetApp ONTAP

<https://aws.amazon.com/fsx/netapp-ontap/>

- Amazon EC2

https://aws.amazon.com/pm/ec2/?trk=36c6da98-7b20-48fa-8225-4784bced9843&sc_channel=ps&s_kwcid=ALi4422!3!467723097970!e!!g!!aws%20ec2&ef_id=Cj0KCQiA54KfBhCKARIsAJzSrdqwQrghn6l71jiWzSeaT9Uh1-vY-VfhJixF-xnv5rWwn2S7RqZOTQ0aAh7eEALw_wcB:G:s&s_kwcid=ALi4422!3!467723097970!e!!g!!aws%20ec2

TR-4965: Oracle Database Deployment and Protection in AWS FSx/EC2 with iSCSI/ASM

Allen Cao, Niyaz Mohamed, NetApp

This solution provides overview and details for Oracle database deployment and protection in AWS FSx ONTAP storage and EC2 compute instance with iSCSI protocol and Oracle database configured in standalone ReStart using asm as volume manager.

Purpose

ASM (Automatic Storage Management) is a popular Oracle storage volume manager employed in many Oracle installations. It is also Oracle's recommended storage management solution. It provides an alternative to conventional volume managers and file systems. Since Oracle version 11g, ASM packaged with grid infrastructure rather than a database. As a result, in order to utilize Oracle ASM for storage management without RAC, you must install Oracle grid infrastructure in a standalone server, also known as Oracle Restart. Doing so certainly adds more complexity in Oracle database deployment. However, as the name implies, when Oracle deployed in Restart mode, failed Oracle services restarted automatically by grid infrastructure or after a host reboot without user intervention, which provides a certain degree of high availability or HA functionality.

In this documentation, we demonstrate how to deploy an Oracle database with the iSCSI protocol and Oracle ASM in an Amazon FSx for ONTAP storage environment with EC2 compute instances. We also demonstrate how to use the NetApp SnapCenter service through the NetApp BlueXP console to backup, restore, and clone your Oracle database for dev/test or other use cases for storage-efficient database operation in the AWS public cloud.

This solution addresses the following use cases:

- Oracle database deployment in Amazon FSx for ONTAP storage and EC2 compute instances with iSCSI/ASM
- Testing and validating an Oracle workload in the public AWS cloud with iSCSI/ASM
- Testing and validating Oracle database Restart functionalities deployed in AWS

Audience


This solution is intended for the following people:

- A DBA who would like to deploy Oracle in an AWS public cloud with iSCSI/ASM.
- A database solution architect who would like to test Oracle workloads in the AWS public cloud.
- The storage administrator who would like to deploy and manage an Oracle database deployed to AWS FSx storage.
- The application owner who would like to stand up an Oracle database in AWS FSx/EC2.

Solution test and validation environment

The testing and validation of this solution was performed in an AWS FSx and EC2 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		
FSx ONTAP storage	Current version offered by AWS	One FSx HA cluster in the same VPC and availability zone

EC2 instance for compute	t2.xlarge/4vCPU/16G	Two EC2 T2 xlarge EC2 instances, one as primary DB server and the other as a clone DB server
Software		
RedHat Linux	RHEL-8.6.0_HVM-20220503-x86_64-2-Hourly2-GP2	Deployed RedHat subscription for testing
Oracle Grid Infrastructure	Version 19.18	Applied RU patch p34762026_190000_Linux-x86-64.zip
Oracle Database	Version 19.18	Applied RU patch p34765931_190000_Linux-x86-64.zip
Oracle OPatch	Version 12.2.0.1.36	Latest patch p6880880_190000_Linux-x86-64.zip
SnapCenter Service	Version	v2.3.1.2324

Key factors for deployment consideration

- **EC2 compute instances.** In these tests and validations, we used an AWS EC2 t2.xlarge instance type for the Oracle database compute instance. NetApp recommends using an M5 type EC2 instance as the compute instance for Oracle in production deployment because it is optimized for database workloads. You need to size the EC2 instance appropriately for the number of vCPUs and the amount of RAM based on actual workload requirements.
- **FSx storage HA clusters single- or multi-zone deployment.** In these tests and validations, we deployed an FSx HA cluster in a single AWS availability zone. For production deployment, NetApp recommends deploying an FSx HA pair in two different availability zones. An FSx HA cluster is always provisioned in a HA pair that is sync mirrored in a pair of active-passive file systems to provide storage-level redundancy. Multi-zone deployment further enhances high availability in the event of failure in a single AWS zone.
- **FSx storage cluster sizing.** An Amazon FSx for ONTAP storage file system provides up to 160,000 raw SSD IOPS, up to 4GBps throughput, and a maximum of 192TiB capacity. However, you can size the cluster in terms of provisioned IOPS, throughput, and the storage limit (minimum 1,024 GiB) based on your actual requirements at the time of deployment. The capacity can be adjusted dynamically on the fly without affecting application availability.
- **Oracle data and logs layout.** In our tests and validations, we deployed two ASM disk groups for data and logs respectively. Within the +DATA asm disk group, we provisioned four LUNs in a data volume. Within the +LOGS asm disk group, we provisioned two LUNs in a logs volume. In general, multiple LUNs laid out within an Amazon FSx for ONTAP volume provides better performance.
- **iSCSI configuration.** The EC2 instance database server connects to FSx storage with the iSCSI protocol. EC2 instances generally deploy with a single network interface or ENI. The single NIC interface carries both iSCSI and application traffic. It is important to gauge the Oracle database peak I/O throughput requirement by carefully analyzing the Oracle AWR report in order to choose a right EC2 compute instance that meets both application and iSCSI traffic-throughput requirements. NetApp also recommends allocating four iSCSI connections to both FSx iSCSI endpoints with multipath properly configured.
- **Oracle ASM redundancy level to use for each Oracle ASM disk group that you create.** Because FSx already mirrors the storage on the FSx cluster level, you should use External Redundancy, which means that the option does not allow Oracle ASM to mirror the contents of the disk group.

- **Database backup.** NetApp provides a SaaS version of SnapCenter software service for database backup, restore, and clone in the cloud that is available through the NetApp BlueXP console UI. NetApp recommends implementing such a service to achieve fast (under a minute) SnapShot backup, quick (few minutes) database restore, and database cloning.

Solution deployment

The following section provides step-by-step deployment procedures.

Prerequisites for deployment

Deployment requires the following prerequisites.

1. An AWS account has been set up, and the necessary VPC and network segments have been created within your AWS account.
2. From the AWS EC2 console, you must deploy two EC2 Linux instances, one as the primary Oracle DB server and an optional alternative clone target DB server. See the architecture diagram in the previous section for more details about the environment setup. Also review the [User Guide for Linux instances](#) for more information.
3. From the AWS EC2 console, deploy Amazon FSx for ONTAP storage HA clusters to host the Oracle database volumes. If you are not familiar with the deployment of FSx storage, see the documentation [Creating FSx for ONTAP file systems](#) for step-by-step instructions.
4. Steps 2 and 3 can be performed using the following Terraform automation toolkit, which creates an EC2 instance named `ora_01` and an FSx file system named `fsx_01`. Review the instruction carefully and change the variables to suit your environment before execution.

```
git clone https://github.com/NetApp-  
Automation/na_aws_fsx_ec2_deploy.git
```



Ensure that you have allocated at least 50G in EC2 instance root volume in order to have sufficient space to stage Oracle installation files.

EC2 instance kernel configuration

With the prerequisites provisioned, log into the EC2 instance as `ec2-user` and `sudo` to root user to configure the Linux kernel for Oracle installation.

1. Create a staging directory `/tmp/archive` folder and set the `777` permission.

```
mkdir /tmp/archive  
  
chmod 777 /tmp/archive
```

2. Download and stage the Oracle binary installation files and other required rpm files to the `/tmp/archive` directory.

See the following list of installation files to be staged in `/tmp/archive` on the EC2 instance.

```
[ec2-user@ip-172-30-15-58 ~]$ ls -l /tmp/archive  
total 10537316  
-rw-rw-r--. 1 ec2-user ec2-user      19112 Mar 21 15:57 compat-  
libcap1-1.10-7.el7.x86_64.rpm  
-rw-rw-r--  1 ec2-user ec2-user 3059705302 Mar 21 22:01  
LINUX.X64_193000_db_home.zip  
-rw-rw-r--  1 ec2-user ec2-user 2889184573 Mar 21 21:09  
LINUX.X64_193000_grid_home.zip  
-rw-rw-r--. 1 ec2-user ec2-user      589145 Mar 21 15:56  
netapp_linux_unified_host_utilities-7-1.x86_64.rpm  
-rw-rw-r--. 1 ec2-user ec2-user      31828 Mar 21 15:55 oracle-  
database-preinstall-19c-1.0-2.el8.x86_64.rpm  
-rw-rw-r--  1 ec2-user ec2-user 2872741741 Mar 21 22:31  
p34762026_190000_Linux-x86-64.zip  
-rw-rw-r--  1 ec2-user ec2-user 1843577895 Mar 21 22:32  
p34765931_190000_Linux-x86-64.zip  
-rw-rw-r--  1 ec2-user ec2-user  124347218 Mar 21 22:33  
p6880880_190000_Linux-x86-64.zip  
-rw-r--r--  1 ec2-user ec2-user    257136 Mar 22 16:25  
policycoreutils-python-utils-2.9-9.el8.noarch.rpm
```

3. Install Oracle 19c preinstall RPM, which satisfies most kernel configuration requirements.

```
yum install /tmp/archive/oracle-database-preinstall-19c-1.0-  
2.el8.x86_64.rpm
```

4. Download and install the missing `compat-libcap1` in Linux 8.

```
yum install /tmp/archive/compat-libcap1-1.10-7.el7.x86_64.rpm
```

5. From NetApp, download and install NetApp host utilities.

```
yum install /tmp/archive/netapp_linux_unified_host_utilities-7-1.x86_64.rpm
```

6. Install `policycoreutils-python-utils`, which is not available in the EC2 instance.

```
yum install /tmp/archive/policycoreutils-python-utils-2.9-9.el8.noarch.rpm
```

7. Install open JDK version 1.8.

```
yum install java-1.8.0-openjdk.x86_64
```

8. Install iSCSI initiator utils.

```
yum install iscsi-initiator-utils
```

9. Install `sg3_utils`.

```
yum install sg3_utils
```

10. Install `device-mapper-multipath`.

```
yum install device-mapper-multipath
```

11. Disable transparent hugepages in the current system.

```
echo never > /sys/kernel/mm/transparent_hugepage/enabled  
echo never > /sys/kernel/mm/transparent_hugepage/defrag
```

Add the following lines in `/etc/rc.local` to disable `transparent_hugepage` after reboot:

```
# Disable transparent hugepages
    if test -f /sys/kernel/mm/transparent_hugepage/enabled;
then
    echo never > /sys/kernel/mm/transparent_hugepage/enabled
fi
    if test -f /sys/kernel/mm/transparent_hugepage/defrag;
then
    echo never > /sys/kernel/mm/transparent_hugepage/defrag
fi
```

12. Disable selinux by changing SELINUX=enforcing to SELINUX=disabled. You must reboot the host to make the change effective.

```
vi /etc/sysconfig/selinux
```

13. Add the following lines to `limit.conf` to set the file descriptor limit and stack size without quotes "`"`.

```
vi /etc/security/limits.conf
**          hard    nofile      65536"
**          soft    stack       10240"
```

14. Add swap space to EC2 instance by following this instruction: [How do I allocate memory to work as swap space in an Amazon EC2 instance by using a swap file?](#) The exact amount of space to add depends on the size of RAM up to 16G.
15. Change `node.session.timeo.replacement_timeout` in the `iscsi.conf` configuration file from 120 to 5 seconds.

```
vi /etc/iscsi/iscsid.conf
```

16. Enable and start the iSCSI service on the EC2 instance.

```
systemctl enable iscsid
systemctl start iscsid
```

17. Retrieve the iSCSI initiator address to be used for database LUN mapping.

```
cat /etc/iscsi/initiatorname.iscsi
```

18. Add the ASM group to be used for the asm sysasm group.

```
groupadd asm
```

19. Modify the oracle user to add ASM as a secondary group (the oracle user should have been created after Oracle preinstall RPM installation).

```
usermod -a -G asm oracle
```

20. Stop and disable Linux firewall if it is active.

```
systemctl stop firewalld  
systemctl disable firewalld
```

21. Reboot the EC2 instance.

Provision and map database volumes and LUNs to the EC2 instance host

Provision three volumes from the command line by login to FSx cluster via ssh as fsxadmin user with FSx cluster management IP to host the Oracle database binary, data, and logs files.

1. Log into the FSx cluster through SSH as the fsxadmin user.

```
ssh fsxadmin@172.30.15.53
```

2. Execute the following command to create a volume for the Oracle binary.

```
vol create -volume ora_01_biny -aggregate aggr1 -size 50G -state  
online -type RW -snapshot-policy none -tiering-policy snapshot-only
```

3. Execute the following command to create a volume for Oracle data.

```
vol create -volume ora_01_data -aggregate aggr1 -size 100G -state  
online -type RW -snapshot-policy none -tiering-policy snapshot-only
```

4. Execute the following command to create a volume for Oracle logs.

```
vol create -volume ora_01_logs -aggregate aggr1 -size 100G -state  
online -type RW -snapshot-policy none -tiering-policy snapshot-only
```

5. Create a binary LUN within the database binary volume.

```
lun create -path /vol/ora_01_biny/ora_01_biny_01 -size 40G -ostype  
linux
```

6. Create data LUNs within the database data volume.

```
lun create -path /vol/ora_01_data/ora_01_data_01 -size 20G -ostype  
linux
```

```
lun create -path /vol/ora_01_data/ora_01_data_02 -size 20G -ostype  
linux
```

```
lun create -path /vol/ora_01_data/ora_01_data_03 -size 20G -ostype  
linux
```

```
lun create -path /vol/ora_01_data/ora_01_data_04 -size 20G -ostype  
linux
```

7. Create log LUNs within the database logs volume.

```
lun create -path /vol/ora_01_logs/ora_01_logs_01 -size 40G -ostype linux  
lun create -path /vol/ora_01_logs/ora_01_logs_02 -size 40G -ostype linux
```

8. Create an igroup for the EC2 instance with the initiator retrieved from step 14 of the EC2 kernel configuration above.

```
igroup create -igroup ora_01 -protocol iscsi -ostype linux  
-initiator iqn.1994-05.com.redhat:f65fed7641c2
```

9. Map the LUNs to the igroup created above. Increment the LUN ID sequentially for each additional LUN within a volume.

```
lun map -path /vol/ora_01_biny/ora_01_biny_01 -igroup ora_01  
-vserver svm_ora -lun-id 0  
lun map -path /vol/ora_01_data/ora_01_data_01 -igroup ora_01  
-vserver svm_ora -lun-id 1  
lun map -path /vol/ora_01_data/ora_01_data_02 -igroup ora_01  
-vserver svm_ora -lun-id 2  
lun map -path /vol/ora_01_data/ora_01_data_03 -igroup ora_01  
-vserver svm_ora -lun-id 3  
lun map -path /vol/ora_01_data/ora_01_data_04 -igroup ora_01  
-vserver svm_ora -lun-id 4  
lun map -path /vol/ora_01_logs/ora_01_logs_01 -igroup ora_01  
-vserver svm_ora -lun-id 5  
lun map -path /vol/ora_01_logs/ora_01_logs_02 -igroup ora_01  
-vserver svm_ora -lun-id 6
```

10. Validate the LUN mapping.

```
mapping show
```

This is expected to return:

```
FsxId02ad7bf3476b741df::> mapping show
```

```
(lun mapping show)
```

Vserver Protocol	Path	Igroup	LUN ID
svm_ora iscsi	/vol/ora_01_biny/ora_01_biny_01	ora_01	0
svm_ora iscsi	/vol/ora_01_data/ora_01_data_01	ora_01	1
svm_ora iscsi	/vol/ora_01_data/ora_01_data_02	ora_01	2
svm_ora iscsi	/vol/ora_01_data/ora_01_data_03	ora_01	3
svm_ora iscsi	/vol/ora_01_data/ora_01_data_04	ora_01	4
svm_ora iscsi	/vol/ora_01_logs/ora_01_logs_01	ora_01	5
svm_ora iscsi	/vol/ora_01_logs/ora_01_logs_02	ora_01	6

Database storage configuration

Now, import and set up the FSx storage for the Oracle grid infrastructure and database installation on the EC2 instance host.

1. Log into the EC2 instance via SSH as the ec2-user with your SSH key and EC2 instance IP address.

```
ssh -i ora_01.pem ec2-user@172.30.15.58
```

2. Discover the FSx iSCSI endpoints using either SVM iSCSI IP address. Then change to your environment-specific portal address.

```
sudo iscsiadm iscsiadm --mode discovery --op update --type  
sendtargets --portal 172.30.15.51
```

3. Establish iSCSI sessions by logging into each target.

```
sudo iscsiadm --mode node -l all
```

The expected output from the command is:

```
[ec2-user@ip-172-30-15-58 ~]$ sudo iscsiadm --mode node -l all  
Logging in to [iface: default, target: iqn.1992-  
08.com.netapp:sn.1f795e65c74911edb785affbf0a2b26e:vs.3, portal:  
172.30.15.51,3260]  
Logging in to [iface: default, target: iqn.1992-  
08.com.netapp:sn.1f795e65c74911edb785affbf0a2b26e:vs.3, portal:  
172.30.15.13,3260]  
Login to [iface: default, target: iqn.1992-  
08.com.netapp:sn.1f795e65c74911edb785affbf0a2b26e:vs.3, portal:  
172.30.15.51,3260] successful.  
Login to [iface: default, target: iqn.1992-  
08.com.netapp:sn.1f795e65c74911edb785affbf0a2b26e:vs.3, portal:  
172.30.15.13,3260] successful.
```

4. View and validate a list of active iSCSI sessions.

```
sudo iscsiadm --mode session
```

Return the iSCSI sessions.

```
[ec2-user@ip-172-30-15-58 ~]$ sudo iscsiadm --mode session
tcp: [1] 172.30.15.51:3260,1028 iqn.1992-
08.com.netapp:sn.1f795e65c74911edb785affbf0a2b26e:vs.3 (non-flash)
tcp: [2] 172.30.15.13:3260,1029 iqn.1992-
08.com.netapp:sn.1f795e65c74911edb785affbf0a2b26e:vs.3 (non-flash)
```

5. Verify that the LUNs were imported into the host.

```
sudo sanlun lun show
```

This will return a list of Oracle LUNs from FSx.

```

[ec2-user@ip-172-30-15-58 ~]$ sudo sanlun lun show
controller(7mode/E-Series)/                               device
host                lun
vservers(cDOT/FlashRay)  lun-pathname
filename             adapter  protocol  size  product

svm_ora              /vol/ora_01_logs/ora_01_logs_02
/dev/sdn             host3    iSCSI    40g   cDOT
svm_ora              /vol/ora_01_logs/ora_01_logs_01
/dev/sdm             host3    iSCSI    40g   cDOT
svm_ora              /vol/ora_01_data/ora_01_data_03
/dev/sdk             host3    iSCSI    20g   cDOT
svm_ora              /vol/ora_01_data/ora_01_data_04
/dev/sdl             host3    iSCSI    20g   cDOT
svm_ora              /vol/ora_01_data/ora_01_data_01
/dev/sdi             host3    iSCSI    20g   cDOT
svm_ora              /vol/ora_01_data/ora_01_data_02
/dev/sdj             host3    iSCSI    20g   cDOT
svm_ora              /vol/ora_01_biny/ora_01_biny_01
/dev/sdh             host3    iSCSI    40g   cDOT
svm_ora              /vol/ora_01_logs/ora_01_logs_02
/dev/sdg             host2    iSCSI    40g   cDOT
svm_ora              /vol/ora_01_logs/ora_01_logs_01
/dev/sdf             host2    iSCSI    40g   cDOT
svm_ora              /vol/ora_01_data/ora_01_data_04
/dev/sde             host2    iSCSI    20g   cDOT
svm_ora              /vol/ora_01_data/ora_01_data_02
/dev/sdc             host2    iSCSI    20g   cDOT
svm_ora              /vol/ora_01_data/ora_01_data_03
/dev/sdd             host2    iSCSI    20g   cDOT
svm_ora              /vol/ora_01_data/ora_01_data_01
/dev/sdb             host2    iSCSI    20g   cDOT
svm_ora              /vol/ora_01_biny/ora_01_biny_01
/dev/sda             host2    iSCSI    40g   cDOT

```

6. Configure the `multipath.conf` file with following default and blacklist entries.

```
sudo vi /etc/multipath.conf

defaults {
    find_multipaths yes
    user_friendly_names yes
}

blacklist {
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9]*"
    devnode "^hd[a-z]"
    devnode "^cciss.*"
}
```

7. Start the multipath service.

```
sudo systemctl start multipathd
```

Now multipath devices appear in the `/dev/mapper` directory.

```
[ec2-user@ip-172-30-15-58 ~]$ ls -l /dev/mapper
total 0
lrwxrwxrwx 1 root root          7 Mar 21 20:13
3600a09806c574235472455534e68512d -> ../dm-0
lrwxrwxrwx 1 root root          7 Mar 21 20:13
3600a09806c574235472455534e685141 -> ../dm-1
lrwxrwxrwx 1 root root          7 Mar 21 20:13
3600a09806c574235472455534e685142 -> ../dm-2
lrwxrwxrwx 1 root root          7 Mar 21 20:13
3600a09806c574235472455534e685143 -> ../dm-3
lrwxrwxrwx 1 root root          7 Mar 21 20:13
3600a09806c574235472455534e685144 -> ../dm-4
lrwxrwxrwx 1 root root          7 Mar 21 20:13
3600a09806c574235472455534e685145 -> ../dm-5
lrwxrwxrwx 1 root root          7 Mar 21 20:13
3600a09806c574235472455534e685146 -> ../dm-6
crw----- 1 root root    10, 236 Mar 21 18:19 control
```

8. Log into the FSx cluster as the `fsxadmin` user via SSH to retrieve the serial-hex number for each LUN start with `6c574xxx...`, the HEX number start with `3600a0980`, which is AWS vendor ID.

```
lun show -fields serial-hex
```

and return as follow:

```
FsxId02ad7bf3476b741df::> lun show -fields serial-hex
vserver path                               serial-hex
-----
svm_ora /vol/ora_01_biny/ora_01_biny_01 6c574235472455534e68512d
svm_ora /vol/ora_01_data/ora_01_data_01 6c574235472455534e685141
svm_ora /vol/ora_01_data/ora_01_data_02 6c574235472455534e685142
svm_ora /vol/ora_01_data/ora_01_data_03 6c574235472455534e685143
svm_ora /vol/ora_01_data/ora_01_data_04 6c574235472455534e685144
svm_ora /vol/ora_01_logs/ora_01_logs_01 6c574235472455534e685145
svm_ora /vol/ora_01_logs/ora_01_logs_02 6c574235472455534e685146
7 entries were displayed.
```

9. Update the `/dev/multipath.conf` file to add a user-friendly name for the multipath device.

```
sudo vi /etc/multipath.conf
```

with following entries:

```

multipaths {
    multipath {
        wwid          3600a09806c574235472455534e68512d
        alias         ora_01_biny_01
    }
    multipath {
        wwid          3600a09806c574235472455534e685141
        alias         ora_01_data_01
    }
    multipath {
        wwid          3600a09806c574235472455534e685142
        alias         ora_01_data_02
    }
    multipath {
        wwid          3600a09806c574235472455534e685143
        alias         ora_01_data_03
    }
    multipath {
        wwid          3600a09806c574235472455534e685144
        alias         ora_01_data_04
    }
    multipath {
        wwid          3600a09806c574235472455534e685145
        alias         ora_01_logs_01
    }
    multipath {
        wwid          3600a09806c574235472455534e685146
        alias         ora_01_logs_02
    }
}

```

10. Reboot the multipath service to verify that the devices under `/dev/mapper` have changed to LUN names versus serial-hex IDs.

```
sudo systemctl restart multipathd
```

Check `/dev/mapper` to return as following:

```
[ec2-user@ip-172-30-15-58 ~]$ ls -l /dev/mapper
total 0
crw----- 1 root root 10, 236 Mar 21 18:19 control
lrwxrwxrwx 1 root root      7 Mar 21 20:41 ora_01_biny_01 -> ../dm-
0
lrwxrwxrwx 1 root root      7 Mar 21 20:41 ora_01_data_01 -> ../dm-
1
lrwxrwxrwx 1 root root      7 Mar 21 20:41 ora_01_data_02 -> ../dm-
2
lrwxrwxrwx 1 root root      7 Mar 21 20:41 ora_01_data_03 -> ../dm-
3
lrwxrwxrwx 1 root root      7 Mar 21 20:41 ora_01_data_04 -> ../dm-
4
lrwxrwxrwx 1 root root      7 Mar 21 20:41 ora_01_logs_01 -> ../dm-
5
lrwxrwxrwx 1 root root      7 Mar 21 20:41 ora_01_logs_02 -> ../dm-
6
```

11. Partition the binary LUN with a single primary partition.

```
sudo fdisk /dev/mapper/ora_01_biny_01
```

12. Format the partitioned binary LUN with an XFS file system.

```
sudo mkfs.xfs /dev/mapper/ora_01_biny_01p1
```

13. Mount the binary LUN to /u01.

```
sudo mount -t xfs /dev/mapper/ora_01_biny_01p1 /u01
```

14. Change /u01 mount point ownership to the Oracle user and its associated primary group.

```
sudo chown oracle:oinstall /u01
```

15. Find the UUID of the binary LUN.

```
sudo blkid /dev/mapper/ora_01_biny_01p1
```

16. Add a mount point to /etc/fstab.

```
sudo vi /etc/fstab
```

Add the following line.

```
UUID=d89fb1c9-4f89-4de4-b4d9-17754036d11d    /u01    xfs
defaults,nofail 0                2
```



It is important to mount the binary with only the UUID and with the nofail option to avoid possible root-lock issues during EC2-instance reboot.

17. As the root user, add the udev rule for Oracle devices.

```
vi /etc/udev/rules.d/99-oracle-asmdevices.rules
```

Include following entries:

```
ENV{DM_NAME}=="ora*", GROUP=="oinstall", OWNER=="oracle",
MODE=="660"
```

18. As the root user, reload the udev rules.

```
udevadm control --reload-rules
```

19. As the root user, trigger the udev rules.

```
udevadm trigger
```

20. As the root user, reload multipathd.

```
systemctl restart multipathd
```

21. Reboot the EC2 instance host.

Oracle grid infrastructure installation

1. Log into the EC2 instance as the ec2-user via SSH and enable password authentication by uncommenting `PasswordAuthentication yes` and then commenting out `PasswordAuthentication no`.

```
sudo vi /etc/ssh/sshd_config
```

2. Restart the sshd service.

```
sudo systemctl restart sshd
```

3. Reset the Oracle user password.

```
sudo passwd oracle
```

4. Log in as the Oracle Restart software owner user (oracle). Create an Oracle directory as follows:

```
mkdir -p /u01/app/oracle  
mkdir -p /u01/app/oraInventory
```

5. Change the directory permission setting.

```
chmod -R 775 /u01/app
```

6. Create a grid home directory and change to it.

```
mkdir -p /u01/app/oracle/product/19.0.0/grid  
cd /u01/app/oracle/product/19.0.0/grid
```

7. Unzip the grid installation files.

```
unzip -q /tmp/archive/LINUX.X64_193000_grid_home.zip
```

8. From grid home, delete the OPatch directory.

```
rm -rf OPatch
```

9. From grid home, unzip `p6880880_190000_Linux-x86-64.zip`.

```
unzip -q /tmp/archive/p6880880_190000_Linux-x86-64.zip
```

10. From grid home, revise `cv/admin/cvu_config`, uncomment and replace `CV_ASSUME_DISTID=OEL5` with `CV_ASSUME_DISTID=OL7`.

```
vi cv/admin/cvu_config
```

11. Prepare a `gridsetup.rsp` file for silent installation and place the `rsp` file in the `/tmp/archive` directory. The `rsp` file should cover sections A, B, and G with the following information:

```
INVENTORY_LOCATION=/u01/app/oraInventory
oracle.install.option=HA_CONFIG
ORACLE_BASE=/u01/app/oracle
oracle.install.asm.OSDBA=dba
oracle.install.asm.OSOPER=oper
oracle.install.asm.OSASM=asm
oracle.install.asm.SYSASMPassword="SetPWD"
oracle.install.asm.diskGroup.name=DATA
oracle.install.asm.diskGroup.redundancy=EXTERNAL
oracle.install.asm.diskGroup.AUSize=4
oracle.install.asm.diskGroup.disks=/dev/mapper/ora_01_data_01,/dev/mapper/ora_01_data_02,/dev/mapper/ora_01_data_03,/dev/mapper/ora_01_data_04
oracle.install.asm.diskGroup.diskDiscoveryString=/dev/mapper/*
oracle.install.asm.monitorPassword="SetPWD"
oracle.install.asm.configureAFD=true
```

12. Log into the EC2 instance as the root user and set `ORACLE_HOME` and `ORACLE_BASE`.

```
export ORACLE_HOME=/u01/app/oracle/product/19.0.0/grid
export ORACLE_BASE=/tmp
cd /u01/app/oracle/product/19.0.0/grid/bin
```

13. Provision disk devices for use with the Oracle ASM filter driver.

```
./asmcmd afd_label DATA01 /dev/mapper/ora_01_data_01 --init  
./asmcmd afd_label DATA02 /dev/mapper/ora_01_data_02 --init  
./asmcmd afd_label DATA03 /dev/mapper/ora_01_data_03 --init  
./asmcmd afd_label DATA04 /dev/mapper/ora_01_data_04 --init  
./asmcmd afd_label LOGS01 /dev/mapper/ora_01_logs_01 --init  
./asmcmd afd_label LOGS02 /dev/mapper/ora_01_logs_02 --init
```

14. Install `cvuqdisk-1.0.10-1.rpm`.

```
rpm -ivh /u01/app/oracle/product/19.0.0/grid/cv/rpm/cvuqdisk-1.0.10-1.rpm
```

15. Unset `$ORACLE_BASE`.

```
unset ORACLE_BASE
```

16. Log into the EC2 instance as the Oracle user and extract the patch in the `/tmp/archive` folder.

```
unzip /tmp/archive/p34762026_190000_Linux-x86-64.zip -d /tmp/archive
```

17. From grid home `/u01/app/oracle/product/19.0.0/grid` and as the oracle user, launch `gridSetup.sh` for grid infrastructure installation.

```
./gridSetup.sh -applyRU /tmp/archive/34762026/ -silent  
-responseFile /tmp/archive/gridsetup.rsp
```

Ignore the warnings about wrong groups for grid infrastructure. We are using a single Oracle user to manage Oracle Restart, so this is expected.

18. As root user, execute the following script(s):

```
/u01/app/oraInventory/orainstRoot.sh  
  
/u01/app/oracle/product/19.0.0/grid/root.sh
```

19. As root user, reload the multipathd.

```
systemctl restart multipathd
```

20. As the Oracle user, execute the following command to complete the configuration:

```
/u01/app/oracle/product/19.0.0/grid/gridSetup.sh -executeConfigTools  
-responseFile /tmp/archive/gridsetup.rsp -silent
```

21. As the Oracle user, create the LOGS disk group.

```
bin/asmca -silent -sysAsmPassword 'yourPWD' -asmsnmpPassword  
'yourPWD' -createDiskGroup -diskGroupName LOGS -disk 'AFD:LOGS*'  
-redundancy EXTERNAL -au_size 4
```

22. As the Oracle user, validate grid services after installation configuration.

```
bin/crsctl stat res -t  
+  
Name                          Target  State        Server  
State details  
Local Resources  
ora.DATA.dg                    ONLINE  ONLINE      ip-172-30-15-58  
STABLE  
ora.LISTENER.lsnr              ONLINE  ONLINE      ip-172-30-15-58  
STABLE  
ora.LOGS.dg                    ONLINE  ONLINE      ip-172-30-15-58  
STABLE  
ora.asm                        ONLINE  ONLINE      ip-172-30-15-58  
Started,STABLE  
ora.ons                        OFFLINE  OFFLINE     ip-172-30-15-58  
STABLE  
Cluster Resources  
ora.cssd                       ONLINE  ONLINE      ip-172-30-15-58  
STABLE  
ora.diskmon                    OFFLINE  OFFLINE  
STABLE  
ora.driver.afd                 ONLINE  ONLINE      ip-172-30-15-58  
STABLE  
ora.evmd                       ONLINE  ONLINE      ip-172-30-15-58  
STABLE
```

23. Valiate ASM filter driver status.

```

[oracle@ip-172-30-15-58 grid]$ export
ORACLE_HOME=/u01/app/oracle/product/19.0.0/grid
[oracle@ip-172-30-15-58 grid]$ export ORACLE_SID=+ASM
[oracle@ip-172-30-15-58 grid]$ export PATH=$PATH:$ORACLE_HOME/bin
[oracle@ip-172-30-15-58 grid]$ asmcmd
ASMCMDB> lsdg
State      Type      Rebal  Sector  Logical_Sector  Block      AU
Total_MB  Free_MB  Req_mir_free_MB  Usable_file_MB  Offline_disks
Voting_files  Name
MOUNTED  EXTERN  N      512     512     4096    1048576
81920    81847      0      81847      0
N  DATA/
MOUNTED  EXTERN  N      512     512     4096    1048576
81920    81853      0      81853      0
N  LOGS/
ASMCMDB> afd_state
ASMCMDB-9526: The AFD state is 'LOADED' and filtering is 'ENABLED' on
host 'ip-172-30-15-58.ec2.internal'

```

Oracle database installation

1. Log in as the Oracle user and unset `$ORACLE_HOME` and `$ORACLE_SID` if it is set.

```
unset ORACLE_HOME
unset ORACLE_SID
```

2. Create the Oracle DB home directory and change to it.

```
mkdir /u01/app/oracle/product/19.0.0/db1
cd /u01/app/oracle/product/19.0.0/db1
```

3. Unzip the Oracle DB installation files.

```
unzip -q /tmp/archive/LINUX.X64_193000_db_home.zip
```

4. From the DB home, delete the OPatch directory.

```
rm -rf OPatch
```

5. From DB home, unzip `p6880880_190000_Linux-x86-64.zip`.

```
unzip -q /tmp/archive/p6880880_190000_Linux-x86-64.zip
```

6. From DB home, revise `cv/admin/cvu_config`, and uncomment and replace `CV_ASSUME_DISTID=OEL5` with `CV_ASSUME_DISTID=OL7`.

```
vi cv/admin/cvu_config
```

7. From the `/tmp/archive` directory, unpack the DB 19.18 RU patch.

```
unzip p34765931_190000_Linux-x86-64.zip
```

8. Prepare the DB silent install `rsp` file in `/tmp/archive/dbinstall.rsp` directory with the following values:

```
oracle.install.option=INSTALL_DB_SWONLY
UNIX_GROUP_NAME=oinstall
INVENTORY_LOCATION=/u01/app/oraInventory
ORACLE_HOME=/u01/app/oracle/product/19.0.0/db1
ORACLE_BASE=/u01/app/oracle
oracle.install.db.InstallEdition=EE
oracle.install.db.OSDBA_GROUP=dba
oracle.install.db.OSOPER_GROUP=oper
oracle.install.db.OSBACKUPDBA_GROUP=oper
oracle.install.db.OSDGDBA_GROUP=dba
oracle.install.db.OSKMDBA_GROUP=dba
oracle.install.db.OSRACDBA_GROUP=dba
oracle.install.db.rootconfig.executeRootScript=false
```

9. From db1 home /u01/app/oracle/product/19.0.0/db1, execute silent software-only DB installation.

```
./runInstaller -applyRU /tmp/archive/34765931/ -silent
-ignorePrereqFailure -responseFile /tmp/archive/dbinstall.rsp
```

10. As root user, run the `root.sh` script after software-only installation.

```
/u01/app/oracle/product/19.0.0/db1/root.sh
```

11. As Oracle user, create the `dbca.rsp` file with the following entries:

```
gdbName=db1.demo.netapp.com
sid=db1
createAsContainerDatabase=true
numberOfPDBs=3
pdbName=db1_pdb
useLocalUndoForPDBs=true
pdbAdminPassword="yourPWD"
templateName=General_Purpose.dbc
sysPassword="yourPWD"
systemPassword="yourPWD"
dbsnmpPassword="yourPWD"
datafileDestination=+DATA
recoveryAreaDestination=+LOGS
storageType=ASM
diskGroupName=DATA
characterSet=AL32UTF8
nationalCharacterSet=AL16UTF16
listeners=LISTENER
databaseType=MULTIPURPOSE
automaticMemoryManagement=false
totalMemory=8192
```

12. As Oracle user, launch DB creation with dbca.


```
bin/dbca -silent -createDatabase -responseFile /tmp/archive/dbca.rsp
```

output:

Prepare for db operation

7% complete

Registering database with Oracle Restart

11% complete

Copying database files

33% complete

Creating and starting Oracle instance

35% complete

38% complete

42% complete

45% complete

48% complete

Completing Database Creation

53% complete

55% complete

56% complete

Creating Pluggable Databases

60% complete

64% complete

69% complete

78% complete

Executing Post Configuration Actions

100% complete

Database creation complete. For details check the logfiles at:

/u01/app/oracle/cfgtoollogs/dbca/db1.

Database Information:

Global Database Name:db1.demo.netapp.com

System Identifier(SID):db1

Look at the log file "/u01/app/oracle/cfgtoollogs/dbca/db1/db1.log"
for further details.

13. As Oracle user, validate Oracle Restart HA services after DB creation.

```
[oracle@ip-172-30-15-58 db1]$ ../grid/bin/crsctl stat res -t
```

Name	Target	State	Server	State
details				
Local Resources				
ora.DATA.dg	ONLINE	ONLINE	ip-172-30-15-58	STABLE
ora.LISTENER.lsnr	ONLINE	ONLINE	ip-172-30-15-58	STABLE
ora.LOGS.dg	ONLINE	ONLINE	ip-172-30-15-58	STABLE
ora.asm	ONLINE	ONLINE	ip-172-30-15-58	Started,STABLE
ora.ons	OFFLINE	OFFLINE	ip-172-30-15-58	STABLE
Cluster Resources				
ora.cssd	ONLINE	ONLINE	ip-172-30-15-58	STABLE
ora.dbf.db	ONLINE	ONLINE	ip-172-30-15-58	Open,HOME=/u01/app/oracle/product/19.0.0/db1,STABLE
ora.diskmon	OFFLINE	OFFLINE		STABLE
ora.driver.afd	ONLINE	ONLINE	ip-172-30-15-58	STABLE
ora.evmd	ONLINE	ONLINE	ip-172-30-15-58	STABLE

14. Set the Oracle user `.bash_profile`.

```
vi ~/.bash_profile
```

15. Add following entries:

```
export ORACLE_HOME=/u01/app/oracle/product/19.0.0/db1
export ORACLE_SID=db1
export PATH=$PATH:$ORACLE_HOME/bin
alias asm='export
ORACLE_HOME=/u01/app/oracle/product/19.0.0/grid;export
ORACLE_SID=+ASM;export PATH=$PATH:$ORACLE_HOME/bin'
```

16. Validate the CDB/PDB created.

```
/home/oracle/.bash_profile

sqlplus / as sysdba
```

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

```
NAME          OPEN_MODE
```

```
DB1           READ WRITE
```

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

```
NAME
```

```
+DATA/DB1/DATAFILE/system.256.1132176177
```

```
+DATA/DB1/DATAFILE/sysaux.257.1132176221
```

```
+DATA/DB1/DATAFILE/undotbs1.258.1132176247
```

```
+DATA/DB1/86B637B62FE07A65E053F706E80A27CA/DATAFILE/system.265.1132177009
```

```
+DATA/DB1/86B637B62FE07A65E053F706E80A27CA/DATAFILE/sysaux.266.1132177009
```

```
+DATA/DB1/DATAFILE/users.259.1132176247
```

```
+DATA/DB1/86B637B62FE07A65E053F706E80A27CA/DATAFILE/undotbs1.267.1132177009
```

```
+DATA/DB1/F7852758DCD6B800E0533A0F1EAC1DC6/DATAFILE/system.271.1132177853
```

```
+DATA/DB1/F7852758DCD6B800E0533A0F1EAC1DC6/DATAFILE/sysaux.272.1132177853
```

```
+DATA/DB1/F7852758DCD6B800E0533A0F1EAC1DC6/DATAFILE/undotbs1.270.1132177853
```

```
+DATA/DB1/F7852758DCD6B800E0533A0F1EAC1DC6/DATAFILE/users.274.1132177871
```

```
NAME
```

```
+DATA/DB1/F785288BBCD1BA78E0533A0F1EACCD6F/DATAFILE/system.276.1132177871
```

```
+DATA/DB1/F785288BBCD1BA78E0533A0F1EACCD6F/DATAFILE/sysaux.277.1132177871
```

```
+DATA/DB1/F785288BBCD1BA78E0533A0F1EACCD6F/DATAFILE/undotbs1.275.1132177871
```

```
+DATA/DB1/F785288BBCD1BA78E0533A0F1EACCD6F/DATAFILE/users.279.1132177889
```

```
+DATA/DB1/F78529A14DD8BB18E0533A0F1EACB8ED/DATAFILE/system.281.1132177889
```

```
+DATA/DB1/F78529A14DD8BB18E0533A0F1EACB8ED/DATAFILE/sysaux.282.1132177889
```

```
+DATA/DB1/F78529A14DD8BB18E0533A0F1EACB8ED/DATAFILE/undotbs1.280.1132177889
```

```
+DATA/DB1/F78529A14DD8BB18E0533A0F1EACB8ED/DATAFILE/users.284.113217  
7907
```

```
19 rows selected.
```

```
SQL> show pdbs
```

CON_ID	CON_NAME	OPEN MODE	RESTRICTED
2	PDB\$SEED	READ ONLY	NO
3	DB1_PDB1	READ WRITE	NO
4	DB1_PDB2	READ WRITE	NO
5	DB1_PDB3	READ WRITE	NO

```
SQL>
```

17. Set the DB recovery destination size to the +LOGS disk group size.

```
alter system set db_recovery_file_dest_size = 80G scope=both;
```

18. Log into the database with sqlplus and enable archive log mode.

```
sqlplus /as sysdba.  
  
shutdown immediate;  
  
startup mount;  
  
alter database archivelog;  
  
alter database open;
```

This completes Oracle 19c version 19.18 Restart deployment on an Amazon FSx for ONTAP and EC2 compute instance. If desired, NetApp recommends relocating the Oracle control file and online log files to the +LOGS disk group.

Automated deployment option

Refer to [TR-4986: Simplified, Automated Oracle Deployment on Amazon FSx ONTAP with iSCSI](#) for details.

Oracle Database backup, restore, and clone with SnapCenter Service

See [SnapCenter Services for Oracle](#) for details on Oracle database backup, restore, and clone with NetApp BlueXP console.

Where to find additional information

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

- Installing Oracle Grid Infrastructure for a Standalone Server with a New Database Installation

<https://docs.oracle.com/en/database/oracle/oracle-database/19/ladbi/installing-oracle-grid-infrastructure-for-a-standalone-server-with-a-new-database-installation.html#GUID-0B1CEE8C-C893-46AA-8A6A-7B5FAAEC72B3>

- Installing and Configuring Oracle Database Using Response Files

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

- Amazon FSx for NetApp ONTAP

<https://aws.amazon.com/fsx/netapp-ontap/>

- Amazon EC2

https://aws.amazon.com/pm/ec2/?trk=36c6da98-7b20-48fa-8225-4784bced9843&sc_channel=ps&s_kwcid=AL!4422!3!467723097970!e!!g!!aws%20ec2&ef_id=Cj0KCQiA54KfBhCKARIsAJzSrdqwQrghn6l71jiWzSeaT9Uh1-vY-VfhJixF-xnv5rWwn2S7RqZOTQ0aAh7eEALw_wcB:G:s&s_kwcid=AL!4422!3!467723097970!e!!g!!aws%20ec2

Oracle Database Deployment on AWS EC2 and FSx Best Practices

WP-7357: Oracle Database Deployment on EC2 and FSx Best Practices Introduction

Allen Cao, Niyaz Mohamed, Jeffrey Steiner, NetApp

Many mission-critical enterprise Oracle 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. AWS EC2 compute instances and the AWS FSx 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.

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides secure, resizable compute capacity in the cloud. It is designed to make web-scale cloud computing easier for enterprises. The simple Amazon EC2 web-service interface allows you to obtain and configure capacity with minimal friction. It provides you with complete control of your computing resources and lets you run on Amazon's proven computing environment.

Amazon FSx for ONTAP is an AWS storage service that uses industry-leading NetApp ONTAP block and file storage, which exposes NFS, SMB, and iSCSI. With such a powerful storage engine, it has never been easier to relocate mission-critical Oracle database apps to AWS with sub-millisecond response times, multiple GBps of throughput, and 100,000+ IOPS per database instance. Better yet, the FSx storage service comes with

native replication capability that allows you to easily migrate your on-premises Oracle database to AWS or to replicate your mission critical Oracle database to a secondary AWS availability zone for HA or DR.

The goal of this documentation is to provide step-by-step processes, procedures, and best-practice guidance on how to deploy and configure an Oracle database with FSx storage and an EC2 instance that delivers performance similar to an on-premises system. NetApp also provides an automation toolkit that automates most of the tasks that are required for the deployment, configuration, and management of your Oracle database workload in the AWS public cloud.

To learn more about the solution and use case, take a look at following overview video:

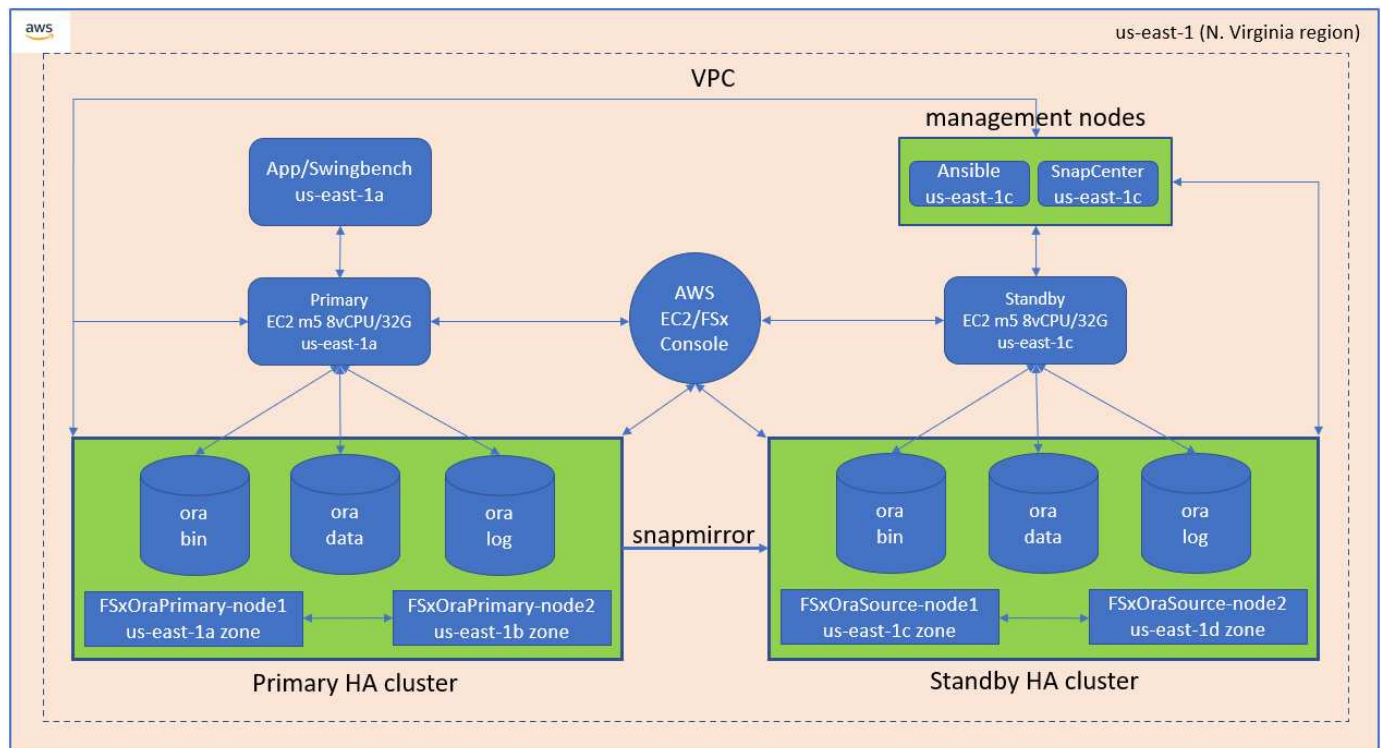
[Modernize your Oracle database with hybrid cloud in AWS and FSx ONTAP, Part1 - Use case and solution architecture](#)

Solution architecture

The following architecture diagram illustrates a highly available Oracle database deployment on an AWS EC2 instance with the FSx storage service. A similar deployment scheme but with the standby in a different region can be set up for disaster recovery.

Within the environment, the Oracle compute instance is deployed via an AWS EC2 instance console. There are multiple EC2 instance types available from the console. NetApp recommends deploying a database-oriented EC2 instance type such as an m5 Ami image with RedHat enterprise Linux 8 and up to 10Gbps of network bandwidth.

Oracle database storage on FSx volumes on the other hand is deployed with the AWS FSx console or CLI. The Oracle binary, data, or log volumes are subsequently presented and mounted on an EC2 instance Linux host. Each data or log volume can have multiple LUNs allocated depending on the underlying storage protocol employed.



An FSx storage cluster is designed with double redundancy, so that both the primary and standby storage clusters are deployed in two different availability zones. Database volumes are replicated from a primary FSx

cluster to a standby FSx cluster at a user-configurable interval for all Oracle binary, data, and log volumes.

This high availability Oracle environment is managed with an Ansible controller node and a SnapCenter backup server and UI tool. Oracle installation, configuration, and replication are automated using Ansible playbook-based tool kits. Any update to the Oracle EC2 instance kernel operating system or Oracle patching can be executed in parallel to keep the primary and standby in sync. In fact, the initial automation setup can be easily expanded to perform some repeating daily Oracle tasks if needed.

SnapCenter provides workflows for Oracle database point-in-time recovery or for database cloning at either the primary or standby zones if needed. Through the SnapCenter UI, you can configure Oracle database backup and replication to standby FSx storage for high availability or disaster recovery based on your RTO or RPO objectives.

The solution provides an alternative process that delivers capabilities similar to those available from Oracle RAC and Data Guard deployment.

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 Oracle database in an AWS public cloud on an EC2 instance with FSx storage.

VM performance


Selecting the right VM size is important for optimal performance of a relational database in a public cloud. For better performance, NetApp recommends using an EC2 M5 Series instance for Oracle deployment, which is optimized for database workloads. The same instance type is also used to power a RDS instance for Oracle by AWS.

- Choose the correct vCPU and RAM combination based on workload characteristics.
- Add swap space to a VM. The default EC2 instance deployment does not create a swap space, which is not optimal for a database.


Storage layout and settings

NetApp recommends the following storage layout:

- For NFS storage, the recommended volume layout is three volumes: one for the Oracle binary; one for Oracle data and a duplicate control file; and one for the Oracle active log, archived log, and control file.



- For iSCSI storage, the recommended volume layout is three volumes: one for the Oracle binary; one for Oracle data and a duplicate control file; and one for the Oracle active log, archived log, and control file. However, each data and log volume ideally should contain four LUNs. The LUNs are ideally balanced on the HA cluster nodes.



- For storage IOPS and throughput, you can choose the threshold for provisioned IOPS and throughput for the FSx storage cluster, and these parameters can be adjusted on the fly anytime the workload changes.
 - The auto IOPS setting is three IOPS per GiB of allocated storage capacity or user defined storage up to 80,000.
 - The throughput level is incremented as follow: 128, 256, 512, 1024, 2045 MBps.

Review the [Amazon FSx for NetApp ONTAP performance](#) documentation when sizing throughput and IOPS.

NFS configuration

Linux, the most common operating system, includes native NFS capabilities. Oracle offers the direct NFS (dNFS) client natively integrated into Oracle. Oracle has supported NFSv3 for over 20 years. dNFS is supported with NFSv3 with all versions of Oracle. NFSv4 is supported with all OS's that follow the NFSv4 standard. dNFS support for NFSv4 requires Oracle 12.1.0.2 or higher. NFSv4.1 requires specific OS support. Consult the NetApp Interoperability Matrix Tool (IMT) for supported OS's. dNFS support for NFSv4.1 requires Oracle version 19.3.0.0 or higher.

Automated Oracle deployment using the NetApp automation toolkit automatically configures dNFS on NFSv3.

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 the TCP slot tables to 128.

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

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

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,rsize=65536,wsiz=65536</code>
<ul style="list-style-type: none">• ORACLE_HOME• ORACLE_BASE	<code>rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=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.

High availability

As indicated in the solution architecture, HA is built on storage-level replication. Therefore, the startup and availability of Oracle is contingent on how quickly the compute and storage can be brought up and recovered. See the following key factors:

- Have a standby compute instance ready and synced up with the primary through Ansible parallel update to both hosts.
- Replicate the binary volume from the primary for standby purposes so that you do not need to install Oracle at the last minute and figure out what needs to be installed and patched.
- Replication frequency dictates how fast the Oracle database can be recovered to make service available. There is a trade off between the replication frequency and storage consumption.
- Leverage automation to make recovery and switch over to standby quick and free of human error. NetApp provides an automation toolkit for this purpose.

Step-by-Step Oracle Deployment Procedures on AWS EC2 and FSx

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

Deploy an EC2 Linux instance for Oracle via EC2 console

If you are new to AWS, you first need to set up an AWS environment. The documentation tab at the AWS website landing page provides EC2 instruction links on how to deploy a Linux EC2 instance that can be used to host your Oracle database via the AWS EC2 console. The following section is a summary of these steps. For details, see the linked AWS EC2-specific documentation.

Setting up your AWS EC2 environment

You must create an AWS account to provision the necessary resources to run your Oracle environment on the EC2 and FSx service. The following AWS documentation provides the necessary details:

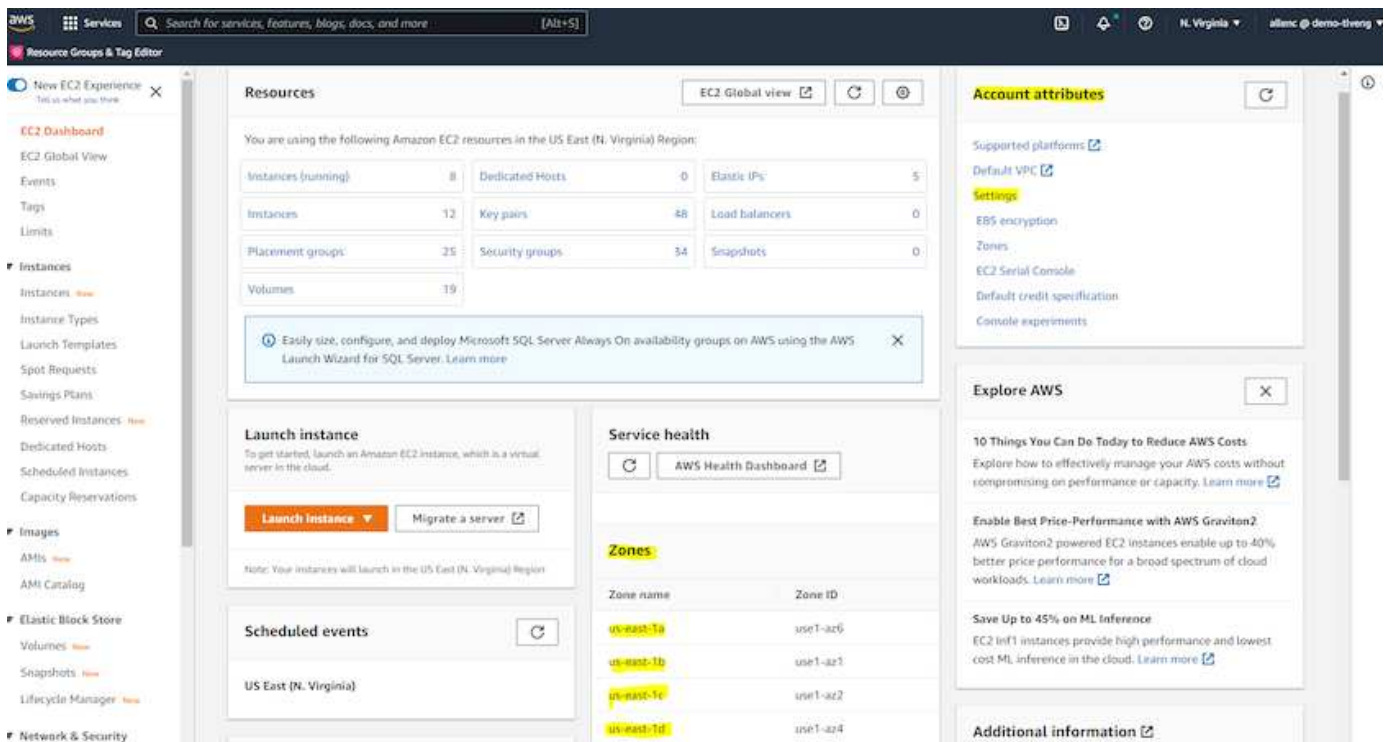
- [Set up to use Amazon EC2](#)

Key topics:

- Sign up for AWS.
- Create a key pair.
- Create a security group.

Enabling multiple availability zones in AWS account attributes

For an Oracle high availability configuration as demonstrated in the architecture diagram, you must enable at least four availability zones in a region. The multiple availability zones can also be situated in different regions to meet the required distances for disaster recovery.



The screenshot displays the AWS Management Console interface. The top navigation bar includes the AWS logo, a search bar, and the current region (N. Virginia) and user profile (allenc @ demo-12345). The main content area is divided into several sections:

- Resources:** A table showing EC2 resources in the US East (N. Virginia) Region. The table includes columns for resource type and count.
- Account attributes:** A panel on the right showing various account settings such as Supported platforms, Default VPC, Settings, EBS encryption, Zones, EC2 Serial Console, Default credit specification, and Console experiments.
- Launch instance:** A section with a "Launch instance" button and a "Migrate a server" button. A note indicates that instances will launch in the US East (N. Virginia) Region.
- Service health:** A section with an "AWS Health Dashboard" button.
- Zones:** A table listing the available availability zones in the region.
- Scheduled events:** A section showing a scheduled event for "US East (N. Virginia)".
- Explore AWS:** A section with several promotional cards, including "10 Things You Can Do Today to Reduce AWS Costs", "Enable Best Price-Performance with AWS Graviton2", and "Save Up to 45% on ML Inference".
- Additional information:** A section with a link to "Additional information".

Resource type	Count
Instances (running)	8
Dedicated Hosts	0
Elastic IPs	5
Instances	12
Key pairs	48
Load balancers	0
Placement groups	25
Security groups	34
Snapshots	0
Volumes	19

Zone name	Zone ID
us-east-1a	use1-az6
us-east-1b	use1-az3
us-east-1c	use1-az2
us-east-1d	use1-az4

Creating and connecting to an EC2 instance for hosting Oracle database

See the tutorial [Get started with Amazon EC2 Linux instances](#) for step-by-step deployment procedures and best practices.


Key topics:

- Overview.
- Prerequisites.
- Step 1: Launch an instance.
- Step 2: Connect to your instance.

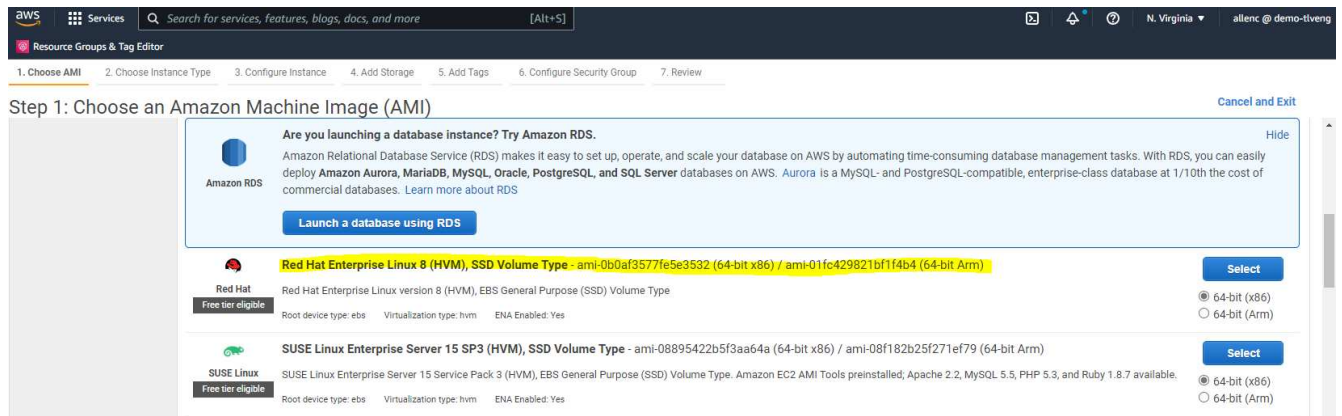
- Step 3: Clean up your instance.

The following screen shots demonstrate the deployment of an m5-type Linux instance with the EC2 console for running Oracle.

1. From the EC2 dashboard, click the yellow Launch Instance button to start the EC2 instance deployment workflow.



2. In Step 1, select "Red Hat Enterprise Linux 8 (HVM), SSD Volume Type - ami-0b0af3577fe5e3532 (64-bit x86) / ami-01fc429821bf1f4b4 (64-bit Arm)."



3. In Step 2, select an m5 instance type with the appropriate CPU and memory allocation based on your Oracle database workload. Click "Next: Configure Instance Details."

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia allenc @ demo-tlven

Resource Groups & Tag Editor

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 2: Choose an Instance Type

<input type="checkbox"/>	m4	m4.16xlarge	64	256	EBS only	Yes	25 Gigabit	Yes
<input type="checkbox"/>	m5	m5.large	2	8	EBS only	Yes	Up to 10 Gigabit	Yes
<input type="checkbox"/>	m5	m5.xlarge	4	16	EBS only	Yes	Up to 10 Gigabit	Yes
<input checked="" type="checkbox"/>	m5	m5.2xlarge	8	32	EBS only	Yes	Up to 10 Gigabit	Yes
<input type="checkbox"/>	m5	m5.4xlarge	16	64	EBS only	Yes	Up to 10 Gigabit	Yes
<input type="checkbox"/>	m5	m5.8xlarge	32	128	EBS only	Yes	10 Gigabit	Yes
<input type="checkbox"/>	m5	m5.12xlarge	48	192	EBS only	Yes	10 Gigabit	Yes
<input type="checkbox"/>	m5	m5.16xlarge	64	256	EBS only	Yes	20 Gigabit	Yes
<input type="checkbox"/>	m5	m5.24xlarge	96	384	EBS only	Yes	25 Gigabit	Yes
<input type="checkbox"/>	m5	m5.metal	96	384	EBS only	Yes	25 Gigabit	Yes

4. In Step 3, choose the VPC and subnet where the instance should be placed and enable public IP assignment. Click "Next: Add Storage."

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia allenc @ demo-tlven

Resource Groups & Tag Editor

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 3: Configure Instance Details

No default VPC found. Select another VPC, or create a new default VPC.

Configure the instance to suit your requirements. You can launch multiple instances from the same AMI, request Spot instances to take advantage of the lower pricing, assign an access management role to the instance, and more.

Number of instances Launch into Auto Scaling Group

Purchasing option Request Spot instances

Network Create new VPC
No default VPC found. Create a new default VPC.

Subnet Create new subnet
250 IP Addresses available

Auto-assign Public IP

Hostname type

DNS Hostname Enable IP name IPv4 (A record) DNS requests
 Enable resource-based IPv4 (A record) DNS requests
 Enable resource-based IPv6 (AAAA record) DNS requests

Placement group Add instance to placement group

Capacity Reservation

Domain join directory Create new directory

IAM role Create new IAM role

Cancel Previous **Review and Launch** Next: Add Storage

5. In Step 4, allocate enough space for the root disk. You may need the space to add a swap. By default, EC2 instance assign zero swap space, which is not optimal for running Oracle.

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia allenc @ demo-tlveng

Resource Groups & Tag Editor

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 4: Add Storage

Your instance will be launched with the following storage device settings. You can attach additional EBS volumes and instance store volumes to your instance, or edit the settings of the root volume. You can also attach additional EBS volumes after launching an instance, but not instance store volumes. [Learn more](#) about storage options in Amazon EC2.

Volume Type	Device	Snapshot	Size (GiB)	Volume Type	IOPS	Throughput (MB/s)	Delete on Termination	Encryption
Root	/dev/sda1	snap-03a3ad00558b4d17c	50	General Purpose SSD (gp2)	150 / 3000	N/A	<input checked="" type="checkbox"/>	Not Encrypted

[Add New Volume](#)

Free tier eligible customers can get up to 30 GB of EBS General Purpose (SSD) or Magnetic storage. [Learn more](#) about free usage tier eligibility and usage restrictions.

Shared file systems

You currently don't have any file systems on this instance. Select "Add file system" button below to add a file system.

[Add file system](#)

[Cancel](#) [Previous](#) [Review and Launch](#) [Next: Add Tags](#)

6. In Step 5, add a tag for instance identification if needed.

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia allenc @ demo-tlveng

Resource Groups & Tag Editor

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 5: Add Tags

A tag consists of a case-sensitive key-value pair. For example, you could define a tag with key = Name and value = Webservier.
A copy of a tag can be applied to volumes, instances or both.
Tags will be applied to all instances and volumes. [Learn more](#) about tagging your Amazon EC2 resources.

Key (128 characters maximum)	Value (256 characters maximum)	Instances	Volumes	Network Interfaces
<p><i>This resource currently has no tags</i></p> <p>Choose the Add tag button or click to add a Name tag.</p> <p>Make sure your IAM policy includes permissions to create tags.</p>				

[Add Tag](#) (Up to 50 tags maximum)

[Cancel](#) [Previous](#) [Review and Launch](#) [Next: Configure Security Group](#)

7. In Step 6, select an existing security group or create a new one with the desired inbound and outbound policy for the instance.

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia allenc @ demo-tveng

Resource Groups & Tag Editor

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 6: Configure Security Group

A security group is a set of firewall rules that control the traffic for your instance. On this page, you can add rules to allow specific traffic to reach your instance. For example, if you want to set up a web server and allow Internet traffic to reach your instance, add rules that allow unrestricted access to the HTTP and HTTPS ports. You can create a new security group or select from an existing one below. [Learn more](#) about Amazon EC2 security groups.

Assign a security group: Create a new security group Select an existing security group

Security Group ID	Name	Description	Actions
<input type="checkbox"/> sg-0d746a0908b897c48	AviOCCM03112021OCCM1635951256631-OCCMSecurityGroup-B3QFHUJRUVW	NetApp OCCM Instance External Security Group	Copy to new
<input type="checkbox"/> sg-07b0625cd544aee16	AVIOCCM0311OCCM1635943382952-OCCMSecurityGroup-1L8D4QX2SC945	NetApp OCCM Instance External Security Group	Copy to new
<input type="checkbox"/> sg-0618122caef6c50e9	AviOCCM1103OCCM1635944222133-OCCMSecurityGroup-DX5PHX6CKVKC	NetApp OCCM Instance External Security Group	Copy to new
<input type="checkbox"/> sg-0d653ea8c78987e660	AviOCCM1209OCCM1631452667252-OCCMSecurityGroup-TSKVZ1Q4SH48	NetApp OCCM Instance External Security Group	Copy to new
<input type="checkbox"/> sg-0aed9f8836b48c52d	AviOCCMFSXOCCM1638110371156-OCCMSecurityGroup-N0ENZJW3TVYB	NetApp OCCM Instance External Security Group	Copy to new
<input type="checkbox"/> sg-083a6ea5cba912375	connector01OCCM1631455604110-OCCMSecurityGroup-1790QV45PH3ZW	NetApp OCCM Instance External Security Group	Copy to new
<input checked="" type="checkbox"/> sg-08148ca915189ac87	default	default VPC security group	Copy to new
<input type="checkbox"/> sg-07f6c527620e3bb22	fsx02OCCM1633339531669-OCCMSecurityGroup-1XZYC5WM15NP7	NetApp OCCM Instance External Security Group	Copy to new
<input type="checkbox"/> sg-0f359d2ba38db749f	SG-Version10-0CE6MEs-NetAppExternalSecurityGroup-N8B50KGTK5U	ONTAP Cloud firewall rules for management and data interface	Copy to new

Inbound rules for sg-08148ca915189ac87 (Selected security groups: sg-08148ca915189ac87)

Type	Protocol	Port Range	Source	Description
All traffic	All	All	192.168.1.0/24	
All traffic	All	All	sg-08148ca915189ac87 (default)	

Cancel Previous **Review and Launch**

8. In Step 7, review the instance configuration summary, and click Launch to start instance deployment. You are prompted to create a key pair or select a key pair for access to the instance.

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia allenc @ demo-tveng

Resource Groups & Tag Editor

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 7: Review Instance Launch

Please review your instance launch details. You can go back to edit changes for each section. Click **Launch** to assign a key pair to your instance and complete the launch process.

AMI Details [Edit AMI](#)

Red Hat Enterprise Linux 8 (HVM), SSD Volume Type - ami-0b0af3577fe5e3532
 Free tier eligible Red Hat Enterprise Linux version 8 (HVM), EBS General Purpose (SSD) Volume Type
 Root Device Type: ebs Virtualization type: hvm

Instance Type [Edit instance type](#)

Instance Type	ECUs	vCPUs	Memory (GiB)	Instance Storage (GB)	EBS-Optimized Available	Network Performance
m5.2xlarge	-	8	32	EBS only	Yes	Up to 10 Gigabit

Security Groups [Edit security groups](#)

Security Group ID	Name	Description
sg-08148ca915189ac87	default	default VPC security group

All selected security groups inbound rules

Type	Protocol	Port Range	Source	Description
All traffic	All	All	192.168.1.0/24	
All traffic	All	All	sg-08148ca915189ac87 (default)	

Instance Details [Edit instance details](#)

Storage [Edit storage](#)

Cancel Previous **Launch**

Select an existing key pair or create a new key pair ✕

A key pair consists of a **public key** that AWS stores, and a **private key file** that you store. Together, they allow you to connect to your instance securely. For Windows AMIs, the private key file is required to obtain the password used to log into your instance. For Linux AMIs, the private key file allows you to securely SSH into your instance. Amazon EC2 supports ED25519 and RSA key pair types.

Note: The selected key pair will be added to the set of keys authorized for this instance. [Learn more about removing existing key pairs from a public AMI.](#)

Choose an existing key pair ▼

Select a key pair

accesststkey | RSA ▼

I acknowledge that I have access to the corresponding private key file, and that without this file, I won't be able to log into my instance.

Cancel
Launch Instances

9. Log into EC2 instance using an SSH key pair. Make changes to your key name and instance IP address as appropriate.

```
ssh -i ora-dblv2.pem ec2-user@54.80.114.77
```

You need to create two EC2 instances as primary and standby Oracle servers in their designated availability zone as demonstrated in the architecture diagram.

Provision FSx for ONTAP file systems for Oracle database storage

EC2 instance deployment allocates an EBS root volume for the OS. FSx for ONTAP file systems provides Oracle database storage volumes, including the Oracle binary, data, and log volumes. The FSx storage NFS volumes can be either provisioned from the AWS FSx console or from Oracle installation, and configuration automation that allocates the volumes as the user configures in a automation parameter file.

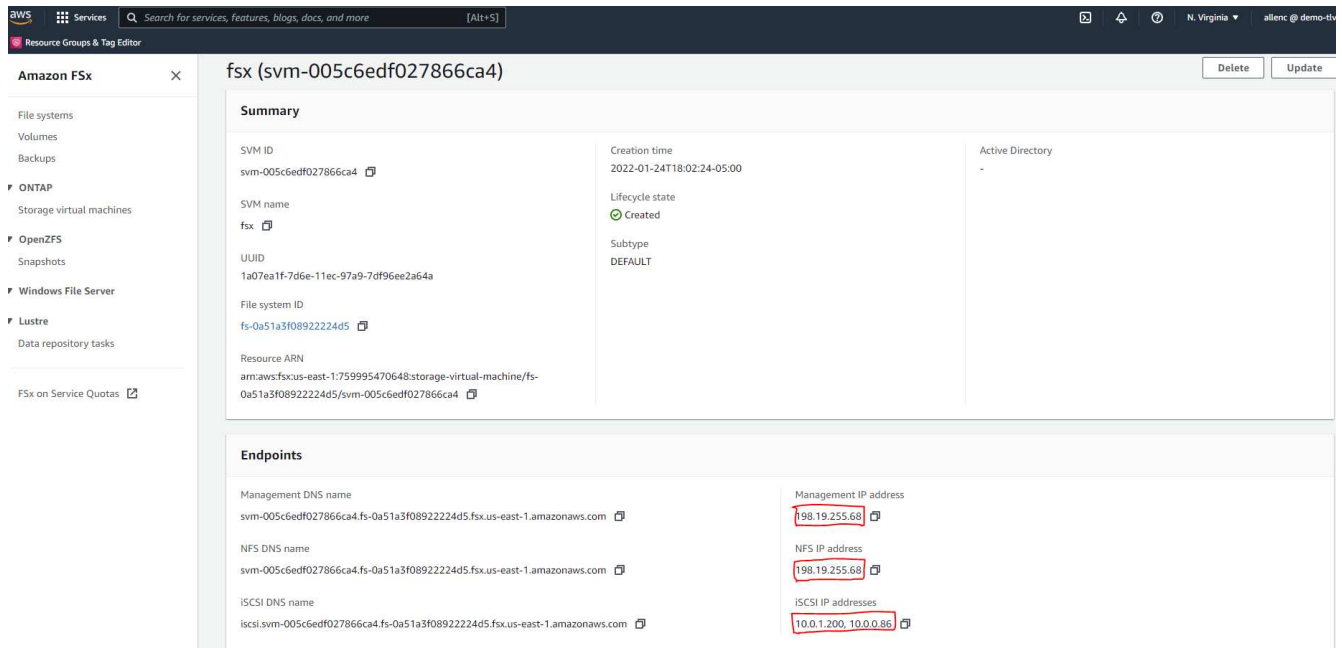
Creating FSx for ONTAP file systems

Referred to this documentation [Managing FSx for ONTAP file systems](#) for creating FSx for ONTAP file systems.

Key considerations:


- SSD storage capacity. Minimum 1024 GiB, maximum 192 TiB.
- Provisioned SSD IOPS. Based on workload requirements, a maximum of 80,000 SSD IOPS per file system.
- Throughput capacity.

- Set administrator fsxadmin/vsadmin password. Required for FSx configuration automation.
- Backup and maintenance. Disable automatic daily backups; database storage backup is executed through SnapCenter scheduling.
- Retrieve the SVM management IP address as well as protocol-specific access addresses from SVM details page. Required for FSx configuration automation.




See the following step-by-step procedures for setting up either a primary or standby HA FSx cluster.

1. From the FSx console, click Create File System to start the FSx provision workflow.



2. Select Amazon FSx for NetApp ONTAP. Then click Next.



3. Select Standard Create and, in File System Details, name your file system, Multi-AZ HA. Based on your database workload, choose either Automatic or User-Provisioned IOPS up to 80,000 SSD IOPS. FSx storage comes with up to 2TiB NVMe caching at the backend that can deliver even higher measured IOPS.

File system details

File system name - optional [Info](#)

Maximum of 256 Unicode letters, whitespace, and numbers, plus + - = . _ : /

Deployment type [Info](#)

Multi-AZ

Single-AZ

SSD storage capacity [Info](#)

Minimum 1024 GiB; Maximum 192 TiB.

Provisioned SSD IOPS

Amazon FSx provides 3 IOPS per GiB of storage capacity. You can also provision additional SSD IOPS as needed.

Automatic (3 IOPS per GiB of SSD storage)

User-provisioned

Maximum 80,000 IOPS

Throughput capacity [Info](#)

The sustained speed at which the file server hosting your file system can serve data. The file server can also burst to higher speeds for periods of time.

Recommended throughput capacity

128 MB/s

Specify throughput capacity

Throughput capacity

4. In the Network & Security section, select the VPC, security group, and subnets. These should be created before FSx deployment. Based on the role of the FSx cluster (primary or standby), place the FSx storage nodes in the appropriate zones.

Network & security

Virtual Private Cloud (VPC) [Info](#)

Specify the VPC from which your file system is accessible.

vpc-0474064fc537e5182

VPC Security Groups [Info](#)

Specify VPC Security Groups to associate with your file system's network interfaces.

Choose VPC security group(s)

sg-08148ca915189ac87 (default) X

Preferred subnet [Info](#)

Specify the preferred subnet for your file system.

subnet-08c952541f4ab282d (us-east-1a)

Standby subnet

subnet-0a84d6eeeb0f4e5c0 (us-east-1b)

VPC route tables

Specify the VPC route tables associated with your file system.

- VPC's default route table
- Select one or more VPC route tables

Endpoint IP address range

Specify the IP address range in which the endpoints to access your file system will be created

- No preference
- Select an IP address range

5. In the Security & Encryption section, accept the default, and enter the fsxadmin password.

Security & encryption

Encryption key [Info](#)

AWS Key Management Service (KMS) encryption key that protects your file system data at rest.

aws/fsx (default)

Description	Account	KMS key ID
Default master key that protects my FSx resources when no other key is defined	759995470648	5b31feff-6759-4306-a852-9c99a743982a

File system administrative password

Password for this file system's "fsxadmin" user, which you can use to access the ONTAP CLI or REST API.

- Don't specify a password
- Specify a password

Password

Confirm password

6. Enter the SVM name and the vsadmin password.

Default storage virtual machine configuration

Storage virtual machine name

SVM administrative password
Password for this SVM's "vsadmin" user, which you can use to access the ONTAP CLI or REST API.

Don't specify a password

Specify a password

Password

Confirm password

Active Directory
Joining an Active Directory enables access from Windows and MacOS clients over the SMB protocol.

Do not join an Active Directory

Join an Active Directory

7. Leave the volume configuration blank; you do not need to create a volume at this point.

Default volume configuration

Volume name

Maximum of 203 alphanumeric characters, plus _.

Junction path

The location within your file system where your volume will be mounted.

Volume size

Minimum 20 MiB; Maximum 104857600 MiB

Storage efficiency
Select whether you would like to enable ONTAP storage efficiencies on your volume: deduplication, compression, and compaction.

Enabled (recommended)
 Disabled

Capacity pool tiering policy
You can optionally enable automatic tiering of your data to lower-cost capacity pool storage.

► **Backup and maintenance - optional**

► **Tags - optional**

Cancel

8. Review the Summary page, and click Create File System to complete FSx file system provision.

aws Services Search for services, features, blogs, docs, and more [Alt+S]

Resource Groups & Tag Editor

Step 1 Select file system type

Step 2 Specify file system details

Step 3 Review and create

Create file system

Summary
Verify the following attributes before proceeding

Attribute	Value	Editable after creation
File system type	Amazon FSx for NetApp ONTAP	
File system name	aws_ora_prod	✓
Deployment type	Multi-AZ	
Storage type	SSD	
SSD storage capacity	1,024 GiB	✓
Minimum SSD IOPS	40000 IOPS	✓
Throughput capacity	512 MB/s	✓
Virtual Private Cloud (VPC)	vpc-0474064fc537e5182	
VPC Security Groups	sg-08148ca915189ac87	✓
Preferred subnet	subnet-08c952541f4ab282d	
Standby subnet	subnet-0a84d6eeeb0f4e5c0	
VPC route tables	VPC's default route table	
Endpoint IP address range	No preference	
KMS key ID	arn:aws:kms:us-east-1:759995470648:key/5b31feff-6759-4306-a852-9c99a743982a	
Daily automatic backup window	No preference	✓
Automatic backup	7 day(s)	✓

Provisioning of database volumes for Oracle database

See [Managing FSx for ONTAP volumes - creating a volume](#) for details.

Key considerations:

- Sizing the database volumes appropriately.
- Disabling capacity pool tiering policy for performance configuration.
- Enabling Oracle dNFS for NFS storage volumes.
- Setting up multipath for iSCSI storage volumes.

Create database volume from FSx console

From the AWS FSx console, you can create three volumes for Oracle database file storage: one for the Oracle binary, one for the Oracle data, and one for the Oracle log. Make sure that volume naming matches the Oracle host name (defined in the hosts file in the automation toolkit) for proper identification. In this example, we use db1 as the EC2 Oracle host name instead of a typical IP-address-based host name for an EC2 instance.

Create volume



File system

ONTAP | fs-0a51a3f08922224d5 | rdscustomfs007



Storage virtual machine

svm-005c6edf027866ca4 | fsx



Volume name

db1_bin

Maximum of 203 alphanumeric characters, plus _.

Junction path

/db1_bin

The location within your file system where your volume will be mounted.

Volume size

51200

Minimum 20 MiB; Maximum 104857600 MiB

Storage efficiency

Select whether you would like to enable ONTAP storage efficiencies on your volume: deduplication, compression, and compaction.

- Enabled (recommended)
- Disabled

Capacity pool tiering policy

You can optionally enable automatic tiering of your data to lower-cost capacity pool storage.

None



Cancel

Confirm

Create volume



File system

ONTAP | fs-0a51a3f08922224d5 | rdscustomfs007



Storage virtual machine

svm-005c6edf027866ca4 | fsx



Volume name

db1_data

Maximum of 203 alphanumeric characters, plus _ .

Junction path

/db1_data

The location within your file system where your volume will be mounted.

Volume size

512000

Minimum 20 MiB; Maximum 104857600 MiB

Storage efficiency

Select whether you would like to enable ONTAP storage efficiencies on your volume: deduplication, compression, and compaction.

Enabled (recommended)

Disabled

Capacity pool tiering policy

You can optionally enable automatic tiering of your data to lower-cost capacity pool storage.

None



Cancel

Confirm

Create volume
✕

File system

ONTAP | fs-0a51a3f08922224d5 | rdscustomfs007 ▼

Storage virtual machine

svm-005c6edf027866ca4 | fsx ▼

Volume name

db1_log

Maximum of 203 alphanumeric characters, plus _.

Junction path

/db1_log

The location within your file system where your volume will be mounted.

Volume size

256000

Minimum 20 MiB; Maximum 104857600 MiB

Storage efficiency

Select whether you would like to enable ONTAP storage efficiencies on your volume: deduplication, compression, and compaction.

Enabled (recommended)

 Disabled

Capacity pool tiering policy

You can optionally enable automatic tiering of your data to lower-cost capacity pool storage.

None ▼

Cancel
Confirm



Creating iSCSI LUNs is not currently supported by the FSx console. For iSCSI LUNs deployment for Oracle, the volumes and LUNs can be created by using automation for ONTAP with the NetApp Automation Toolkit.

Install and configure Oracle on an EC2 instance with FSx database volumes

The NetApp automation team provide an automation kit to run Oracle installation and configuration on EC2 instances according to best practices. The current version of the automation kit supports Oracle 19c on NFS with the default RU patch 19.8. The automation kit can be easily adapted for other RU patches if needed.

Prepare a Ansible controller to run automation

Follow the instruction in the section "[Creating and connecting to an EC2 instance for hosting Oracle database](#)" to provision a small EC2 Linux instance to run the Ansible controller. Rather than using RedHat, Amazon Linux t2.large with 2vCPU and 8G RAM should be sufficient.

Retrieve NetApp Oracle deployment automation toolkit

Log into the EC2 Ansible controller instance provisioned from step 1 as ec2-user and from the ec2-user home directory, execute the `git clone` command to clone a copy of the automation code.

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

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

Execute automated Oracle 19c deployment using automation toolkit

See these detailed instruction [CLI deployment Oracle 19c Database](#) to deploy Oracle 19c with CLI automation. There is a small change in command syntax for playbook execution because you are using an SSH key pair instead of a password for host access authentication. The following list is a high level summary:

1. By default, an EC2 instance uses an SSH key pair for access authentication. From Ansible controller automation root directories `/home/ec2-user/na_oracle19c_deploy`, and `/home/ec2-user/na_rds_fsx_oranfs_config`, make a copy of the SSH key `accesststkey.pem` for the Oracle host deployed in the step "[Creating and connecting to an EC2 instance for hosting Oracle database](#)."
2. Log into the EC2 instance DB host as `ec2-user`, and install the `python3` library.

```
sudo yum install python3
```

3. Create a 16G swap space from the root disk drive. By default, an EC2 instance creates zero swap space. Follow this AWS documentation: [How do I allocate memory to work as swap space in an Amazon EC2 instance by using a swap file?](#)
4. Return to the Ansible controller (`cd /home/ec2-user/na_rds_fsx_oranfs_config`), and execute the `preclone` playbook with the appropriate requirements and `linux_config` tags.

```
ansible-playbook -i hosts rds_preclone_config.yml -u ec2-user --private  
-key accesststkey.pem -e @vars/fsx_vars.yml -t requirements_config
```

```
ansible-playbook -i hosts rds_preclone_config.yml -u ec2-user --private  
-key accesststkey.pem -e @vars/fsx_vars.yml -t linux_config
```

5. Switch to the `/home/ec2-user/na_oracle19c_deploy-master` directory, read the `README` file, and populate the `global_vars.yml` file with the relevant global parameters.

6. Populate the `host_name.yml` file with the relevant parameters in the `host_vars` directory.
7. Execute the playbook for Linux, and press Enter when prompted for the `vsadmin` password.

```
ansible-playbook -i hosts all_playbook.yml -u ec2-user --private-key
accesststkey.pem -t linux_config -e @vars/vars.yml
```

8. Execute the playbook for Oracle, and press enter when prompted for the `vsadmin` password.

```
ansible-playbook -i hosts all_playbook.yml -u ec2-user --private-key
accesststkey.pem -t oracle_config -e @vars/vars.yml
```

Change the permission bit on the SSH key file to 400 if needed. Change the Oracle host (`ansible_host` in the `host_vars` file) IP address to your EC2 instance public address.

Setting up SnapMirror between primary and standby FSx HA cluster

For high availability and disaster recovery, you can set up SnapMirror replication between the primary and standby FSx storage cluster. Unlike other cloud storage services, FSx enables a user to control and manage storage replication at a desired frequency and replication throughput. It also enables users to test HA/DR without any effect on availability.

The following steps show how to set up replication between a primary and standby FSx storage cluster.

1. Setup primary and standby cluster peering. Log into the primary cluster as the `fsxadmin` user and execute the following command. This reciprocal create process executes the create command on both the primary cluster and the standby cluster. Replace `standby_cluster_name` with the appropriate name for your environment.

```
cluster peer create -peer-addr
standby_cluster_name,inter_cluster_ip_address -username fsxadmin
-initial-allowed-vserver-peers *
```

2. Set up vServer peering between the primary and standby cluster. Log into the primary cluster as the `vsadmin` user and execute the following command. Replace `primary_vserver_name`, `standby_vserver_name`, `standby_cluster_name` with the appropriate names for your environment.

```
vserver peer create -vserver primary_vserver_name -peer-vserver
standby_vserver_name -peer-cluster standby_cluster_name -applications
snapmirror
```

3. Verify that the cluster and vserver peerings are set up correctly.

```

FsxId00164454fac5591e6::> cluster peer show
Peer Cluster Name          Cluster Serial Number Availability Authentication
-----
FsxId0b6a95149d07aa82e    1-80-000011             Available         ok

FsxId00164454fac5591e6::> vserver peer show
Vserver      Peer      Peer      Peering      Remote
Vserver      Vserver   State     Peer Cluster Applications Vserver
-----
svm_FSxOraSource
svm_FSxOraTarget
peered       FsxId0b6a95149d07aa82e
snapmirror   svm_FSxOraTarget

FsxId00164454fac5591e6::>

```

4. Create target NFS volumes at the standby FSx cluster for each source volume at the primary FSx cluster. Replace the volume name as appropriate for your environment.

```

vol create -volume dr_db1_bin -aggregate aggr1 -size 50G -state online
-policy default -type DP

```

```

vol create -volume dr_db1_data -aggregate aggr1 -size 500G -state online
-policy default -type DP

```

```

vol create -volume dr_db1_log -aggregate aggr1 -size 250G -state online
-policy default -type DP

```

5. You can also create iSCSI volumes and LUNs for the Oracle binary, Oracle data, and the Oracle log if the iSCSI protocol is employed for data access. Leave approximately 10% free space in the volumes for snapshots.

```

vol create -volume dr_db1_bin -aggregate aggr1 -size 50G -state online
-policy default -unix-permissions ---rwxr-xr-x -type RW

```

```

lun create -path /vol/dr_db1_bin/dr_db1_bin_01 -size 45G -ostype linux

```

```

vol create -volume dr_db1_data -aggregate aggr1 -size 500G -state online
-policy default -unix-permissions ---rwxr-xr-x -type RW

```

```

lun create -path /vol/dr_db1_data/dr_db1_data_01 -size 100G -ostype
linux

```

```
lun create -path /vol/dr_db1_data/dr_db1_data_02 -size 100G -ostype linux
```

```
lun create -path /vol/dr_db1_data/dr_db1_data_03 -size 100G -ostype linux
```

```
lun create -path /vol/dr_db1_data/dr_db1_data_04 -size 100G -ostype linux
```

```
vol create -volume dr_db1_log -aggregate aggr1 -size 250G -state online -policy default -unix-permissions ---rwxr-xr-x -type RW
```

```
lun create -path /vol/dr_db1_log/dr_db1_log_01 -size 45G -ostype linux
```

```
lun create -path /vol/dr_db1_log/dr_db1_log_02 -size 45G -ostype linux
```

```
lun create -path /vol/dr_db1_log/dr_db1_log_03 -size 45G -ostype linux
```

```
lun create -path /vol/dr_db1_log/dr_db1_log_04 -size 45G -ostype linux
```

6. For iSCSI LUNs, create mapping for the Oracle host initiator for each LUN, using the binary LUN as an example. Replace the igroup with an appropriate name for your environment, and increment the lun-id for each additional LUN.

```
lun mapping create -path /vol/dr_db1_bin/dr_db1_bin_01 -igroup ip-10-0-1-136 -lun-id 0
```

```
lun mapping create -path /vol/dr_db1_data/dr_db1_data_01 -igroup ip-10-0-1-136 -lun-id 1
```

7. Create a SnapMirror relationship between the primary and standby database volumes. Replace the appropriate SVM name for your environment.s

```
snapmirror create -source-path svm_FSxOraSource:db1_bin -destination  
-path svm_FSxOraTarget:dr_db1_bin -vserver svm_FSxOraTarget -throttle  
unlimited -identity-preserve false -policy MirrorAllSnapshots -type DP
```

```
snapmirror create -source-path svm_FSxOraSource:db1_data -destination  
-path svm_FSxOraTarget:dr_db1_data -vserver svm_FSxOraTarget -throttle  
unlimited -identity-preserve false -policy MirrorAllSnapshots -type DP
```

```
snapmirror create -source-path svm_FSxOraSource:db1_log -destination  
-path svm_FSxOraTarget:dr_db1_log -vserver svm_FSxOraTarget -throttle  
unlimited -identity-preserve false -policy MirrorAllSnapshots -type DP
```

This SnapMirror setup can be automated with a NetApp Automation Toolkit for NFS database volumes. The toolkit is available for download from the NetApp public GitHub site.

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

Read the README instructions carefully before attempting setup and failover testing.



Replicating the Oracle binary from the primary to a standby cluster might have Oracle license implications. Contact your Oracle license representative for clarification. The alternative is to have Oracle installed and configured at the time of recovery and failover.


SnapCenter Deployment

SnapCenter installation

Follow [Installing the SnapCenter Server](#) to install SnapCenter server. This documentation covers how to install a standalone SnapCenter server. A SaaS version of SnapCenter is in beta review and could be available shortly. Check with your NetApp representative for availability if needed.

Configure SnapCenter plugin for EC2 Oracle host

1. After automated SnapCenter installation, log into SnapCenter as an administrative user for the Window host on which the SnapCenter server is installed.



- From the left-side menu, click Settings, and then Credential and New to add ec2-user credentials for SnapCenter plugin installation.

Credential Name	Authentication Mode	Details
244rdscustomdb	SQL	UserId:admin
42rdscustomdb	SQL	UserId:admin
admin	SQL	UserId:admin
administrator	Windows	UserId:administrator
ec2-user	Linux	UserId:ec2-user
onpremSQL	Windows	UserId:rdscustomval/administrator
rdsdb2	Windows	UserId:administrator
rdsdb244	Windows	UserId:administrator
rdssql	Windows	UserId:administrator
tst244	SQL	UserId:admin
tstcredfordemo	Windows	UserId:administrator

- Reset the ec2-user password and enable password SSH authentication by editing the `/etc/ssh/sshd_config` file on the EC2 instance host.
- Verify that the "Use sudo privileges" checkbox is selected. You just reset the ec2-user password in the previous step.

Credential
✕

Credential Name

Authentication Mode

Linux ▼

Username

i

Password

Use sudo privileges

i

Cancel

OK

5. Add the SnapCenter server name and the IP address to the EC2 instance host file for name resolution.

```

[ec2-user@ip-10-0-0-151 ~]$ sudo vi /etc/hosts
[ec2-user@ip-10-0-0-151 ~]$ cat /etc/hosts
127.0.0.1    localhost localhost.localdomain localhost4
localhost4.localhostdomain4
::1         localhost localhost.localdomain localhost6
localhost6.localhostdomain6
10.0.1.233  rdscustomvalsc.rdscustomval.com rdscustomvalsc
```

6. On the SnapCenter server Windows host, add the EC2 instance host IP address to the Windows host file C:\Windows\System32\drivers\etc\hosts.

```

10.0.0.151    ip-10-0-0-151.ec2.internal
```

7. In the left-side menu, select Hosts > Managed Hosts, and then click Add to add the EC2 instance host to SnapCenter.

NetApp SnapCenter®

Managed Hosts | Disks | Shares | Initiator Groups | iSCSI Session

Search by Name

Name	Type	System	Plug-in	Version	Overall Status
RDSAMAZ-VJ0DQK0	Windows	Stand-alone	Microsoft Windows Server, Microsoft SQL Server	4.5	Host down
rdscustommssql1.rdscustomval.com	Windows	Stand-alone	Microsoft Windows Server, Microsoft SQL Server	4.5	Running

Dashboard | Resources | Monitor | Reports | Hosts | Storage Systems | Settings | Alerts

Check Oracle Database, and, before you submit, click More Options.

rdscustomval\administrator | SnapCenterAdmin | Sign Out

Add Host

Host Type: Linux

Host Name: 10.0.0.151

Credentials: ec2-user

Select Plug-ins to Install SnapCenter Plug-ins Package 4.5 P2 for Linux

Oracle Database

SAP HANA

[More Options](#): Port, Install Path, Custom Plug-Ins...

Submit | Cancel

Check Skip Preinstall Checks. Confirm Skipping Preinstall Checks, and then click Submit After Save.

More Options ✕

Port i

Installation Path i

Skip preinstall checks

Custom Plug-ins _____

Choose a File

No plug-ins found.

You are prompted with Confirm Fingerprint, and then click Confirm and Submit.

Confirm Fingerprint ✕

Authenticity of the host cannot be determined i

Host name	Fingerprint	Valid
ip-10-0-0-151.ec2.internal	ssh-rsa 2048 97:6F:3C:7D:38:42:F6:54:B7:AF:E3:61:61:BA:2E:6F	

After successful plugin configuration, the managed host's overall status show as Running.

Managed Hosts							
Disks		Shares		Initiator Groups		iSCSI Session	
Search by Name							
<input type="checkbox"/>	Name	Type	System	Plug-in	Version	Overall Status	
<input type="checkbox"/>	ip-10-0-0-151.ec2.internal	Linux	Stand-alone	UNIX, Oracle Database	4.5	● Running	

Configure backup policy for Oracle database

Refer to this section [Setup database backup policy in SnapCenter](#) for details on configuring the Oracle database backup policy.

Generally you need create a policy for the full snapshot Oracle database backup and a policy for the Oracle archive-log-only snapshot backup.



You can enable Oracle archive log pruning in the backup policy to control log-archive space. Check "Update SnapMirror after creating a local Snapshot copy" in "Select secondary replication option" as you need to replicate to a standby location for HA or DR.

Configure Oracle database backup and scheduling

Database backup in SnapCenter is user configurable and can be set up either individually or as a group in a resource group. The backup interval depends on the RTO and RPO objectives. NetApp recommends that you run a full database backup every few hours and archive the log backup at a higher frequency such as 10-15 mins for quick recovery.

Refer to the Oracle section of [Implement backup policy to protect database](#) for a detailed step-by-step processes for implementing the backup policy created in the section [Configure backup policy for Oracle database](#) and for backup job scheduling.

The following image provides an example of the resources groups that are set up to back up an Oracle database.

Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Overall Status
ORCL	Single Instance	ip-10-0-0-151.ec2.internal	orcl_full_backup orcl_log_backup	Oracle full backup Oracle log backup	03/24/2022 8:40:08 PM	Backup succeeded

EC2 and FSx Oracle database management

In addition to the AWS EC2 and FSx management console, the Ansible control node and the SnapCenter UI tool are deployed for database management in this Oracle environment.

An Ansible control node can be used to manage Oracle environment configuration, with parallel updates that keep primary and standby instances in sync for kernel or patch updates. Failover, resync, and failback can be automated with the NetApp Automation Toolkit to archive fast application recovery and availability with Ansible. Some repeatable database management tasks can be executed using a playbook to reduce human errors.

The SnapCenter UI tool can perform database snapshot backup, point-in-time recovery, database cloning, and so on with the SnapCenter plugin for Oracle databases. For more information about Oracle plugin features, see the [SnapCenter Plug-in for Oracle Database overview](#).

The following sections provide details on how key functions of Oracle database management are fulfilled with the SnapCenter UI:

- Database snapshot backups
- Database point-in-time restore
- Database clone creation

Database cloning creates a replica of a primary database on a separate EC2 host for data recovery in the event of logical data error or corruption, and clones can also be used for application testing, debugging, patch validation, and so on.

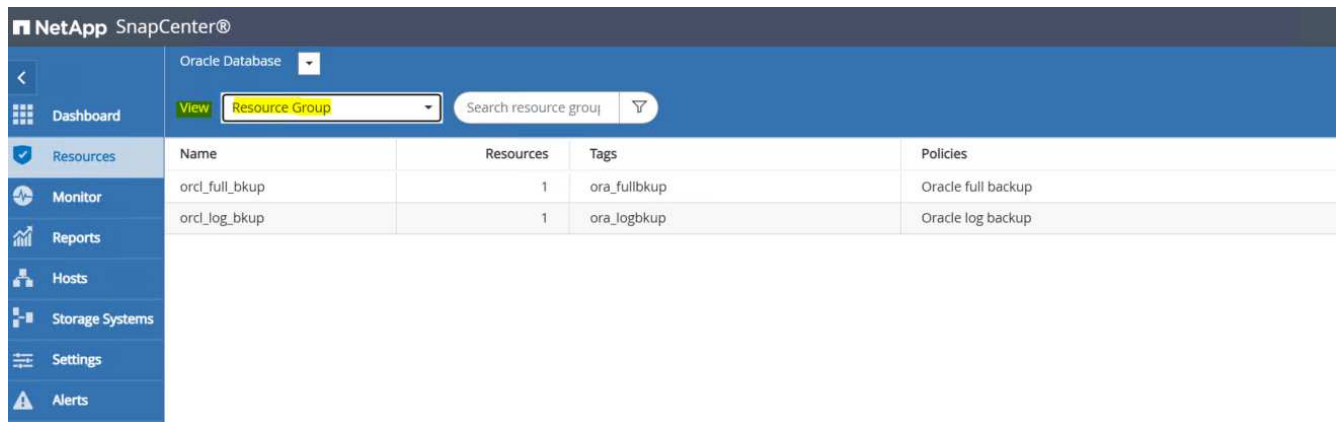
Taking a snapshot

An EC2/FSx Oracle database is regularly backed up at intervals configured by the user. A user can also take a one-off snapshot backup at any time. This applies to both full-database snapshot backups as well as archive-log-only snapshot backups.


Taking a full database snapshot

A full database snapshot includes all Oracle files, including data files, control files, and archive log files.


1. Log into the SnapCenter UI and click Resources in the left-side menu. From the View dropdown, change to the Resource Group view.




2. Click the full backup resource name, and then click the Backup Now icon to initiate an add-hoc backup.



3. Click Backup and then confirm the backup to start a full database backup.



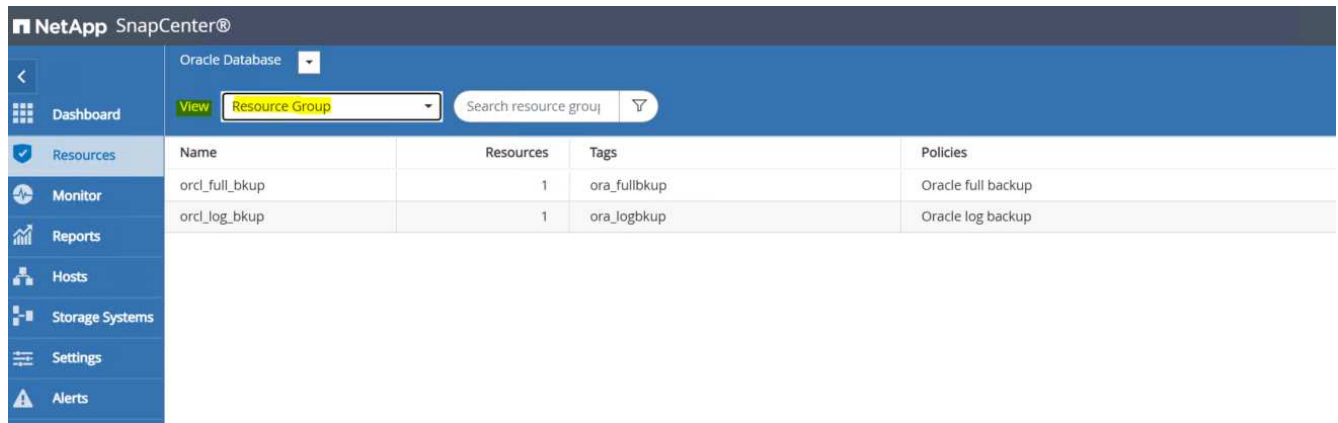
From the Resource view for the database, open the database Managed Backup Copies page to verify that the one-off backup completed successfully. A full database backup creates two snapshots: one for the data volume and one for the log volume.



Taking an archive log snapshot

An archive log snapshot is only taken for the Oracle archive log volume.

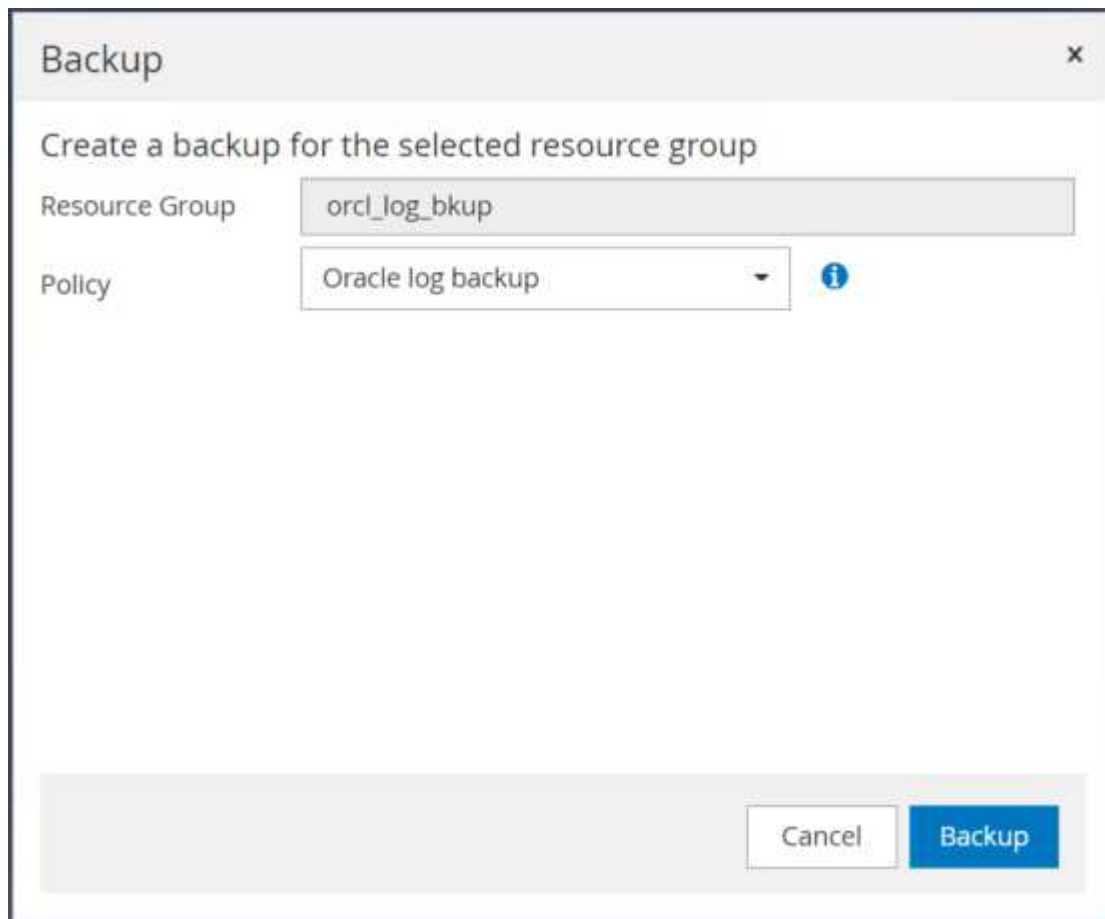
1. Log into the SnapCenter UI and click the Resources tab in the left-side menu bar. From the View dropdown, change to the Resource Group view.



- Click the log backup resource name, and then click the Backup Now icon to initiate an add-hoc backup for archive logs.



- Click Backup and then confirm the backup to start an archive log backup.



From the Resource view for the database, open the database Managed Backup Copies page to verify that the one-off archive log backup completed successfully. An archive log backup creates one snapshot for the log volume.




Restoring to a point in time

SnapCenter-based restore to a point in time is executed on the same EC2 instance host. Complete the following steps to perform the restore:


1. From the SnapCenter Resources tab > Database view, click the database name to open the database backup.




2. Select the database backup copy and the desired point in time to be restored. Also mark down the corresponding SCN number for the point in time. The point-in-time restore can be performed using either the time or the SCN.




3. Highlight the log volume snapshot and click the Mount button to mount the volume.



4. Choose the primary EC2 instance to mount the log volume.



- Verify that the mount job completes successfully. Also check on the EC2 instance host to see that log volume mounted and also the mount point path.



- Copy the archive logs from the mounted log volume to the current archive log directory.

```
[ec2-user@ip-10-0-0-151 ~]$ cp /var/opt/snapcenter/sco/backup_mount/ip-10-0-0-151_03-25-2022_11.15.01.1503_1/ORCL/1/db/ORCL_A/arch/*.arc /ora_nfs_log/db/ORCL_A/arch/
```

- Return to the SnapCenter Resource tab > database backup page, highlight the data snapshot copy, and click the Restore button to start the database restore workflow.

Manage Copies

80 Backups

0 Clones

Local copies

Summary Card

80 Backups

5 Data Backups

75 Log Backups

0 Clones

Primary Backup(s)

Backup Name	Count	Type	End Date	Verified	Mounted	RMAN Cataloged	SCN
ip-10-0-0-151_03-25-2022_12.10.01.1097_1	1	Log	03/25/2022 12:10:09 PM	Not Applicable	False	Not Cataloged	1782417
ip-10-0-0-151_03-25-2022_11.55.01.0500_1	1	Log	03/25/2022 11:55:09 AM	Not Applicable	False	Not Cataloged	1781160
ip-10-0-0-151_03-25-2022_11.40.01.0323_1	1	Log	03/25/2022 11:40:09 AM	Not Applicable	False	Not Cataloged	1780268
ip-10-0-0-151_03-25-2022_11.25.01.0430_1	1	Log	03/25/2022 11:25:09 AM	Not Applicable	False	Not Cataloged	1779368
ip-10-0-0-151_03-25-2022_11.15.01.1503_1	1	Log	03/25/2022 11:15:17 AM	Not Applicable	True	Not Cataloged	1778546
ip-10-0-0-151_03-25-2022_11.15.01.1503_0	1	Data	03/25/2022 11:15:11 AM	Unverified	False	Not Cataloged	1778504
ip-10-0-0-151_03-25-2022_11.10.01.1834_1	1	Log	03/25/2022 11:10:09 AM	Not Applicable	False	Not Cataloged	1778184

8. Check "All Datafiles" and "Change database state if needed for restore and recovery", and click Next.

Restore ORCL

- 1 Restore Scope
- 2 Recovery Scope
- 3 PreOps
- 4 PostOps
- 5 Notification
- 6 Summary

Restore Scope ⓘ

All Datafiles

Tablespaces

Control files

Database State

Change database state if needed for restore and recovery

Restore Mode ⓘ

Force in place restore

If this check box is not selected and if any of the in place restore criteria is not met, restore will be performed using the connect and copy method. The connect and copy restore method might take time based on the files being restored.

Previous
Next

9. Choose a desired recovery scope using either SCN or time. Rather than copying the mounted archive logs to the current log directory as demonstrated in step 6, the mounted archive log path can be listed in

"Specify external archive log files locations" for recovery.

The screenshot shows a window titled "Restore ORCL" with a close button (x) in the top right corner. On the left is a vertical navigation pane with six steps: 1 Restore Scope, 2 Recovery Scope (highlighted in blue), 3 PreOps, 4 PostOps, 5 Notification, and 6 Summary. The main content area is titled "Choose Recovery Scope" and contains the following options:

- All Logs (with an information icon)
- Until SCN (System Change Number) (with an information icon)
 - SCN: (with an information icon)
- Date and Time
- No recovery

Below these options is a section titled "Specify external archive log files locations" with a plus (+), minus (-), and information (i) icon. Underneath is a large, empty text input area.

At the bottom right of the window are two buttons: "Previous" and "Next".

10. Specify an optional prescript to run if necessary.

Restore ORCL x

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Specify optional scripts to run before performing a restore job ⓘ

Prescript full path

Arguments

Script timeout

11. Specify an optional afterscript to run if necessary. Check the open database after recovery.

Restore ORCL x

1 Restore Scope
2 Recovery Scope
3 PreOps
4 PostOps
5 Notification
6 Summary

Specify optional scripts to run after performing a restore job ⓘ

Postscript full path

Arguments

Open the database or container database in READ-WRITE mode after recovery

12. Provide an SMTP server and email address if a job notification is needed.

Restore ORCL x

- 1 Restore Scope
- 2 Recovery Scope
- 3 PreOps
- 4 PostOps
- 5 Notification**
- 6 Summary

Provide email settings ⓘ

Email preference:

From:

To:

Subject:

Attach job report

13. Restore the job summary. Click finish to launch the restore job.


Restore ORCL x

- 1 Restore Scope
- 2 Recovery Scope
- 3 PreOps
- 4 PostOps
- 5 Notification
- 6 Summary**

Summary

Backup name	ip-10-0-0-151_03-25-2022_11.15.01.1503_0
Backup date	03/25/2022 11:15:11 AM
Restore scope	All DataFiles
Recovery scope	Until SCN 1778546
Auxiliary destination	
Options	Change database state if necessary , Open the database or container database in READ-WRITE mode after recovery
Prescript full path	None
Prescript arguments	
Postscript full path	None
Postscript arguments	
Send email	No

14. Validate the restore from SnapCenter.



15. Validate the restore from the EC2 instance host.


```

-bash-4.2$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Fri Mar 25 15:44:08 2022
Version 19.8.0.0.0

Copyright (c) 1982, 2020, 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> select name, RESETLOGS_CHANGE#, RESETLOGS_TIME, open_mode from v$database;

NAME          RESETLOGS_CHANGE# RESETLOGS_TIME OPEN_MODE
-----
ORCL          1778547 25-MAR-22 READ WRITE

SQL>


```

16. To unmount the restore log volume, reverse the steps in step 4.

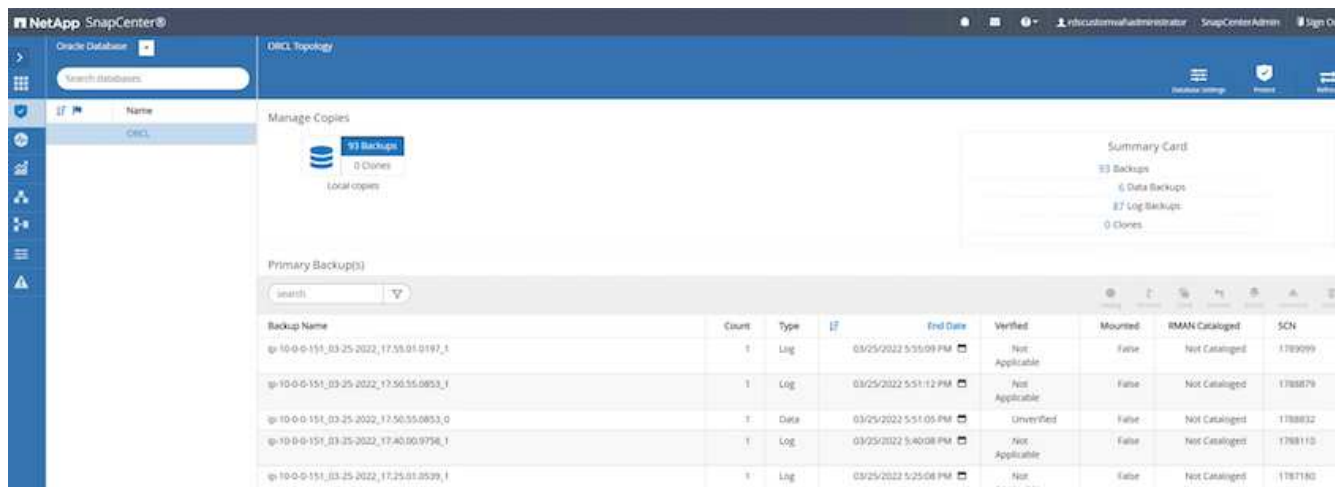
Creating a database clone

The following section demonstrates how to use the SnapCenter clone workflow to create a database clone from a primary database to a standby EC2 instance.

1. Take a full snapshot backup of the primary database from SnapCenter using the full backup resource group.



2. From the SnapCenter Resource tab > Database view, open the Database Backup Management page for the primary database that the replica is to be created from.



3. Mount the log volume snapshot taken in step 4 to the standby EC2 instance host.

ORCL Topology

Database Settings Protect Refresh

Manage Copies

95 Backups
0 Clones
Local copies

Summary Card

95 Backups
6 Data Backups
89 Log Backups
0 Clones

Primary Backup(s)

search

Backup Name	Count	Type	End Date	Verified	Mounted	RMAN Cataloged	SCN
ip-10-0-0-151_03-25-2022_18.55.01.0309_1	1	Log	03/25/2022 6:55:09 PM	Not Applicable	False	Not Cataloged	1892563
ip-10-0-0-151_03-25-2022_18.40.00.9602_1	1	Log	03/25/2022 6:40:23 PM	Not Applicable	False	Not Cataloged	1891375
ip-10-0-0-151_03-25-2022_17.55.01.0197_1	1	Log	03/25/2022 5:55:09 PM	Not Applicable	False	Not Cataloged	1789099
ip-10-0-0-151_03-25-2022_17.50.55.0853_1	1	Log	03/25/2022 5:51:12 PM	Not Applicable	False	Not Cataloged	1788879
ip-10-0-0-151_03-25-2022_17.50.55.0853_0	1	Data	03/25/2022 5:51:05 PM	Unverified	False	Not Cataloged	1788832
ip-10-0-0-151_03-25-2022_17.40.00.9758_1	1	Log	03/25/2022 5:40:08 PM	Not Applicable	False	Not Cataloged	1788110

Mount backups

Choose the host to mount the backup

ip-10-0-0-47.ec2.internal

Mount path : /var/opt/snapcenter/sco/backup_mount/ip-10-0-0-151_03-25-2022_17.50.55.0853_1/ORCL

Mount Cancel

- Highlight the snapshot copy to be cloned for the replica, and click the Clone button to start the clone procedure.

ORCL Topology

Database Settings Protect Refresh

Manage Copies

93 Backups
0 Clones
Local copies

Summary Card

93 Backups
6 Data Backups
87 Log Backups
0 Clones

Primary Backup(s)

search

Backup Name	Count	Type	End Date	Verified	Mounted	RMAN Cataloged	SCN
ip-10-0-0-151_03-25-2022_17.55.01.0197_1	1	Log	03/25/2022 5:55:09 PM	Not Applicable	False	Not Cataloged	1789099
ip-10-0-0-151_03-25-2022_17.50.55.0853_1	1	Log	03/25/2022 5:51:12 PM	Not Applicable	False	Not Cataloged	1788879
ip-10-0-0-151_03-25-2022_17.50.55.0853_0	1	Data	03/25/2022 5:51:05 PM	Unverified	False	Not Cataloged	1788832
ip-10-0-0-151_03-25-2022_17.40.00.9758_1	1	Log	03/25/2022 5:40:08 PM	Not Applicable	False	Not Cataloged	1788110
ip-10-0-0-151_03-25-2022_17.25.01.0539_1	1	Log	03/25/2022 5:25:08 PM	Not Applicable	False	Not Cataloged	1787180

5. Change the replica copy name so that it is different from the primary database name. Click Next.

The screenshot shows a wizard window titled "Clone from ORCL" with a close button (x) in the top right corner. On the left, there is a vertical sidebar with seven steps: 1 Name (highlighted in blue), 2 Locations, 3 Credentials, 4 PreOps, 5 PostOps, 6 Notification, and 7 Summary. The main area of the wizard is titled "Provide clone database SID" and contains a "Clone SID" label followed by a text input field containing the value "ORCLREAD". At the bottom right of the wizard, there are two buttons: "Previous" (disabled) and "Next" (active).

6. Change the clone host to the standby EC2 host, accept the default naming, and click Next.

Clone from ORCL
✕

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps
- 6 Notification
- 7 Summary

Select the host to create a clone

Clone host

Datafile locations ⓘ

Control files ⓘ

Redo logs ⓘ

Group	Size	Unit	Number of files
RedoGroup 1	128	MB	1
<input type="text" value="/ora_nfs_data_ORCLREAD/ORCLREAD/redolog/redo04.log"/>			
RedoGroup 2	128	MB	1

7. Change your Oracle home settings to match those configured for the target Oracle server host, and click Next.

Clone from ORCL

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Database Credentials for the clone

Credential name for sys user: None + i

Database port: 1521

Oracle Home Settings i

Oracle Home: /rdsdbbin/oracle

Oracle OS User: rdsdb

Oracle OS Group: database

Previous Next

8. Specify a recovery point using either time or the SCN and mounted archive log path.

Clone from ORCL

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps**
- 6 Notification
- 7 Summary

Recover Database

Until Cancel *i*

Date and Time *i*

Date-time format: MM/DD/YYYY hh:mm:ss

Until SCN (System Change Number) *i*

Specify external archive log locations *i*

Create new DBID *i*

Create tempfile for temporary tablespace *i*

Enter SQL queries to apply when clone is created

Enter scripts to run after clone operation *i*

Previous **Next**

9. Send the SMTP email settings if needed.

Clone from ORCL x

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps
- 6 Notification**
- 7 Summary

Provide email settings ⓘ

Email preference:

From:

To:

Subject:

Attach job report

10. Clone the job summary, and click Finish to launch the clone job.

Clone from ORCL

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification


7 Summary

Summary

Clone from backup	ip-10-0-0-151_03-25-2022_17.50.55.0853_0
Clone SID	ORCLREAD
Clone server	ip-10-0-0-47.ec2.internal
Oracle home	/rdsdbbin/oracle
Oracle OS user	rdsdb
Oracle OS group	database
Datafile mountpaths	/ora_nfs_data_ORCLREAD
Control files	/ora_nfs_data_ORCLREAD/ORCLREAD/control/control01.ctl
Redo groups	RedoGroup =1 TotalSize =128 Path =/ora_nfs_data_ORCLREAD/ORCLREAD/redolog/redo04.log RedoGroup =2 TotalSize =128 Path =/ora_nfs_data_ORCLREAD/ORCLREAD/redolog/redo03.log RedoGroup =3 TotalSize =128 Path =/ora_nfs_data_ORCLREAD/ORCLREAD/redolog/redo02.log RedoGroup =4 TotalSize =128 Path =/ora_nfs_data_ORCLREAD/ORCLREAD/redolog/redo01.log
Recovery scope	Until SCN 1788879
Prescript full path	none
Prescript arguments	
Postscript full path	none
Postscript arguments	
Send email	No

Previous Finish

11. Validate the replica clone by reviewing the clone job log.



The cloned database is registered in SnapCenter immediately.

Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Overall Status
ORCL	Single Instance	ip-10-0-0-151.ec2.internal	orcl_full_backup orcl_log_backup	Oracle full backup Oracle log backup	03/25/2022 9:10:09 PM	Backup succeeded
ORCLREAD	Single Instance	ip-10-0-0-47.ec2.internal				Not protected

12. Turn off Oracle archive log mode. Log into the EC2 instance as oracle user and execute following command:

```
sqlplus / as sysdba
```

```
shutdown immediate;
```

```
startup mount;
```

```
alter database noarchivelog;
```

```
alter database open;
```



Instead primary Oracle backup copies, a clone can also be created from replicated secondary backup copies on target FSx cluster with same procedures.

HA failover to standby and resync

The standby Oracle HA cluster provides high availability in the event of failure in the primary site, either in the compute layer or in the storage layer. One significant benefit of the solution is that a user can test and validate the infrastructure at any time or with any frequency. Failover can be user simulated or triggered by real failure. The failover processes are identical and can be automated for fast application recovery.

See the following list of failover procedures:

1. For a simulated failover, run a log snapshot backup to flush the latest transactions to the standby site, as demonstrated in the section [Taking an archive log snapshot](#). For a failover triggered by an actual failure, the last recoverable data is replicated to the standby site with the last successful scheduled log volume backup.
2. Break the SnapMirror between primary and standby FSx cluster.
3. Mount the replicated standby database volumes at the standby EC2 instance host.
4. Relink the Oracle binary if the replicated Oracle binary is used for Oracle recovery.
5. Recover the standby Oracle database to the last available archive log.
6. Open the standby Oracle database for application and user access.
7. For an actual primary site failure, the standby Oracle database now takes the role of the new primary site and database volumes can be used to rebuild the failed primary site as a new standby site with the reverse SnapMirror method.
8. For a simulated primary site failure for testing or validation, shut down the standby Oracle database after the completion of testing exercises. Then unmount the standby database volumes from the standby EC2 instance host and resync replication from the primary site to the standby site.

These procedures can be performed with the NetApp Automation Toolkit available for download at the public NetApp GitHub site.

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

Read the README instruction carefully before attempting setup and failover testing.

Database migration from on-prem to public cloud

Database migration is a challenging endeavor by any means. Migrating an Oracle database from on-premises to the cloud is no exception.

The following sections provide key factors to consider when migrating Oracle databases to the AWS public cloud with the AWS EC2 compute and FSx storage platform.

ONTAP storage is available on-premises

If the on-premises Oracle database is sitting on an ONTAP storage array, then it is easier to set up replication for database migration using the NetApp SnapMirror technology that is built into AWS FSx ONTAP storage. The migration process can be orchestrated using NetApp BlueXP console.

1. Build a target compute EC2 instance that matches the on-premises instance.
2. Provision matching, equally sized database volumes from FSx console.
3. Mount the FSx database volumes to the EC2 instance.
4. Set up SnapMirror replication between the on-premises database volumes to the target FSx database volumes. The initial sync might take some time to move the primary source data, but any following incremental updates are much quicker.
5. At the time of switchover, shut down the primary application to stop all transactions. From the Oracle sqlplus CLI interface, execute an Oracle online log switch and allow SnapMirror sync to push the last archived log to the target volume.
6. Break up the mirrored volumes, run Oracle recovery at the target, and bring up the database for service.
7. Point applications to the Oracle database in the cloud.

The following video demonstrates how to migrate an Oracle database from on-premises to AWS FSx/EC2 using the NetApp BlueXP console and SnapMirror replication.

[Migrate on-prem Oracle DB to AWS](#)

ONTAP storage is not available on premises

If the on-premises Oracle database is hosted on third-party storage other than ONTAP, database migration is based on the restore of a Oracle database backup copy. You must play the archive log to make it current before switching over.

AWS S3 can be used as a staging storage area for database move and migration. See the following high level steps for this method:

1. Provision a new, matching EC2 instance that is comparable with the on-premises instance.

2. Provision equal database volumes from FSx storage and mount the volumes to the EC2 instance.
3. Create a disk-level Oracle backup copy.
4. Move the backup copy to AWS S3 storage.
5. Recreate the Oracle control file and restore and recover the database by pulling data and the archive log from S3 storage.
6. Sync the target Oracle database with the on-premises source database.
7. At switchover, shut down the application and source Oracle database. Copy the last few archive logs and apply them to the target Oracle database to bring it up to date.
8. Start up the target database for user access.
9. Redirect application to the target database to complete the switchover.

Migrate on-premises Oracle databases to AWS FSx/EC2 using PDB relocation with maximum availability

This migration approach is best suited to Oracle databases that are already deployed in PDB/CDB multitenant model, and ONTAP storage is not available on-premises. The PDB relocation method utilizes Oracle PDB hot clone technology to move PDBs between a source CDB and a target CDB while minimizing service interruption.

First, create CDB in the AWS FSx/EC2 with sufficient storage to host PDBs to be migrated from on-premises. Multiple on-premises PDBs can be relocated one at a time.

1. If the on-premises database is deployed in a single instance rather than in the multitenant PDB/CDB model, follow the instructions in [Converting a single instance non-CDB to a PDB in a multitenant CDB](#) to convert the single instance to multitenant PDB/CDB. Then follow the next step to migrate the converted PDB to CDB in AWS FSx/EC2.
2. If the on-premises database is already deployed in the multitenant PDB/CDB model, follow the instructions in [Migrate on-premises Oracle databases to cloud with PDB relocation](#) to perform the migration.

The following video demonstrates how an Oracle database (PDB) can be migrated to FSx/EC2 using PDB relocation with maximum availability.

[Migrate on-prem Oracle PDB to AWS CDB with max availability](#)



Although the instructions in step 1 and 2 are illustrated in the context of Azure public cloud, the procedures are applicable to AWS cloud without any changes.

The NetApp Solutions Automation team provides a migration toolkit that can facilitate Oracle database migration from on-premises to the AWS cloud. Use following command to download the Oracle database migration toolkit for PDB relocation.

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

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

Oracle VLDB Incremental Merge via RMAN on ANF



NetApp

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 >](#)

4. 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](#)

<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"/>




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,noi  
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_0g2m6br1 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>

```

- 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-

```

```

SYSAUX_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-
SYSAUX_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-
SYSAUX_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

```

- 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_0q1sd7cm
destination for restore of datafile 00023: /nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-23_041sd6s5
destination for restore of datafile 00027: /nfsanf/oracopy/data_D-
NTAP1_I-2441823937_TS-SOE_FNO-27_081sd70i
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_0s2m6n11
/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=ntapl_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 similiarly 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,rsz=262144,wsz=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_1s2m9kkg
-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-
SYSAUX_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-
SYSAUX_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_1s2m9kkg
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_1s2m9kkg"
```

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.

- 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/onlinelog/redo03.log  
/u03/orareco/NTAP1/onlinelog/redo02.log  
/u03/orareco/NTAP1/onlinelog/redo01.log
```

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

```
Database altered.
```

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

```
Database altered.
```

```
SQL> alter database rename file  
'/u03/orareco/NTAP1/onlinelog/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_1s2m9kkg

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_PDB2,NTAP1_PDB3)	/u01, /u02, /u03 NFS mounts on ANF capacity pool
ora-02	NTAP2(NTAP2_PDB1,NTAP2_PDB2,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 /tmp/ and /root/ 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,rsize=65536,wsiz=65536 0 0
172.30.136.68:/ora-01-u02 /u02 nfs
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsiz=65536 0 0
172.30.136.68:/ora-01-u03 /u03 nfs
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=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.

The screenshot displays the Oracle Enterprise Manager Database Express interface. At the top, there is a login section with the following fields: Username (pre-filled with 'system'), Password (masked with dots), and Container Name. A 'Log In' button is positioned below these fields. The main dashboard area is titled 'Database Home' and shows the instance 'NTAP1 (19.18.0.0.0)'. Key metrics include:

- Status:** Up Time: 34 minutes, 43 seconds; Type: Single Instance (NTAP1); 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:** A line chart showing activity over time, with categories for CPU, User I/O, System I/O, Network, and Other.
- Resources:** Four bar charts showing:
 - Host CPU: 1% usage.
 - Active Sessions: 60 sessions.
 - Memory: Total 14 GB, with sub-categories like total_sga, total_pga, shared_pool, large_pool, buffer cache, and Shared IO P...
 - Data Storage: Total 3.7 GB, with sub-categories for NTAP1_PDB3, NTAP1_PDB2, and NTAP1_PDB1.
- SQL Monitor:** A table showing the top 20 SQL queries by last active time, with columns for Status, Duration, SQL ID, SQL Plan Hash, User Name, Parallel, Database Time, I/O Requests, and 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/ladbi/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/ladbi/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.

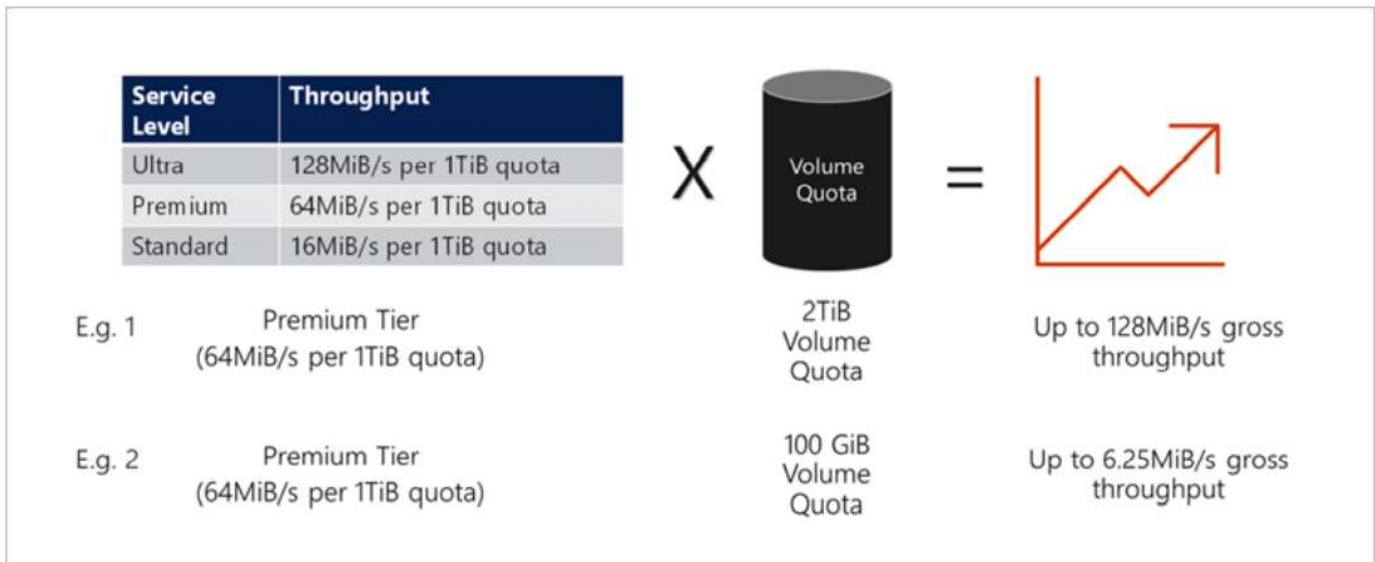
For example, a 10TiB-provisioned single-capacity pool with the premium service level provides an overall available throughput for all volumes in this capacity pool of 10x 64MBps, so 640MBps with 40,000 (16K) IOPs or 80,000 (8K) IOPs.

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,rsize=65536,wsiz=65536</code>
<ul style="list-style-type: none"> ORACLE_HOME ORACLE_BASE 	<code>rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=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**.

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
ANFAVVal2H	Virtual machine	Hybrid Cloud TME Onprem	ANFAVSVAL2	West Europe	Running	Windows	Standard_DS2_v2	20.229.80.88	1
ANFAVSfg01	Virtual machine	Hybrid Cloud TME Onprem	anfavgsg	South Central US	Stopped (deallocated)	Linux	Standard_D32ds_v4	-	1
ANFAVSfgAZ1	Virtual machine	Hybrid Cloud TME Onprem	anfavgsg	South Central US	Running	Linux	Standard_E32as_v4	40.124.74.246	1
ANFAVSfgAZ2	Virtual machine	Hybrid Cloud TME Onprem	anfavgsg	South Central US	Stopped (deallocated)	Linux	Standard_E32as_v4	40.124.178.111	1
ANFAVSfgAZ3	Virtual machine	Hybrid Cloud TME Onprem	anfavgsg	South Central US	Stopped (deallocated)	Linux	Standard_E32as_v4	40.124.194.32	1
ANFAVSvalDC	Virtual machine	Hybrid Cloud TME Onprem	anfavgsg	South Central US	Stopped (deallocated)	Windows	Standard_B4ms	-	1
ANFAVSvalH	Virtual machine	Hybrid Cloud TME Onprem	anfavgsg	South Central US	Running	Windows	Standard_B2ms	70.37.66.218	1
ANFAVSvalH2	Virtual machine	Hybrid Cloud TME Onprem	anfavgsg	South Central US	Running	Windows	Standard_B2s	20.225.210.195	1
ANFCVOCM	Virtual machine	Hybrid Cloud TME Onprem	anfavsval2	West Europe	Running	Linux	Standard_DS3_v2	-	1
ANFCVODRDC2	Virtual machine	Hybrid Cloud TME Onprem	anfavsval2	West Europe	Running	Windows	Standard_B2s	-	1
ANFCVODRDRDemo	Virtual machine	Hybrid Cloud TME Onprem	anfvodrdemo-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 * ⓘ

Resource group * ⓘ [Create new](#)

Instance details

Virtual machine name * ⓘ

Region * ⓘ

Availability options ⓘ

Security type ⓘ

Image * ⓘ [See all images](#) | [Configure VM generation](#)

Run with Azure Spot discount ⓘ

Size * ⓘ [See all sizes](#)

Administrator account

Authentication type ⓘ

SSH public key

Password

Review + create

< Previous

Next : Disks >

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

Create a virtual machine ...

Size * ⓘ See all sizes

Administrator account

Authentication type ⓘ SSH public key
 Password

Username * ⓘ ✓

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 *

⚠ 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 >](#)

3. 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.

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

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 * Delete with VM Enable encryption at host

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

Encryption type * 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
-----	------	------------	-----------	--------------	----------------

[Create and attach a new disk](#) [Attach an existing disk](#)[Advanced](#)[Review + create](#)[< Previous](#)[Next : Networking >](#)





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

Delete public IP and NIC when VM is deleted ⓘ

Enable accelerated networking ⓘ

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?

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

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

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

Create a virtual machine

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

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

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

 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

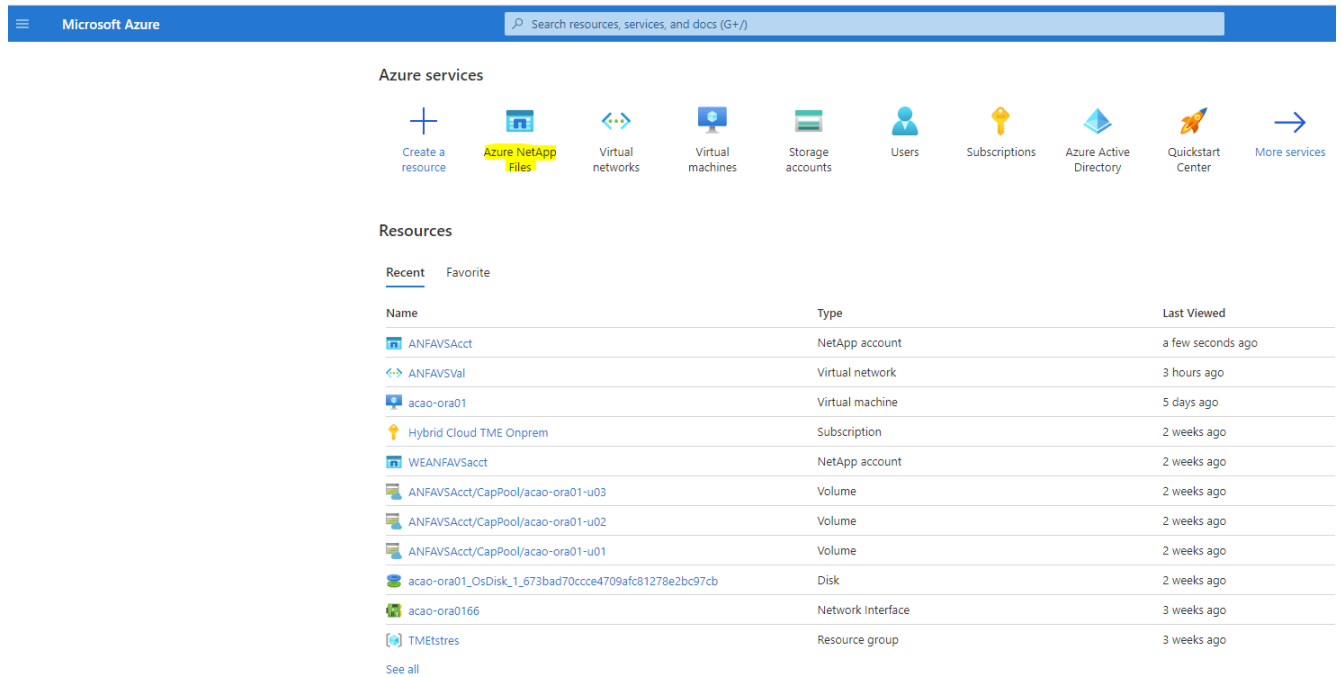
Next : Review + create >

- 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.

- 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.

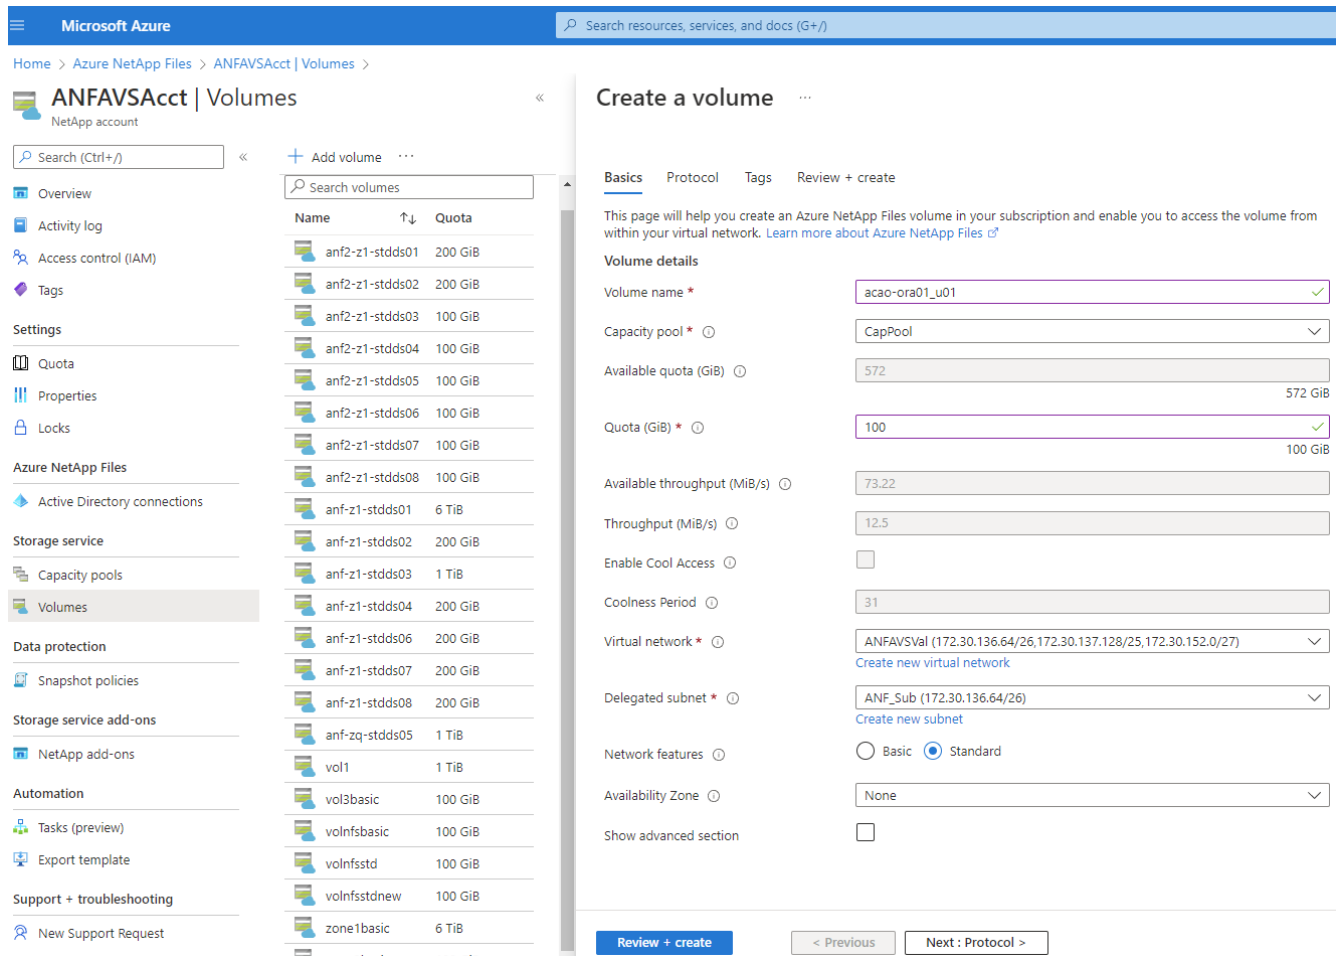


- Under your NetApp storage account, click **Volumes**, and then **Add volume** to create new Oracle volumes.





- 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**.



- 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+/) Add volume

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
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 *

Versions *

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

<input type="checkbox"/>	Index	Allowed clients	Access	Root Access	...
<input type="checkbox"/>	1	0.0.0.0	Read & Write	On	...
<input type="checkbox"/>	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

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
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+/) << + Add volume ...

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
vol3basic	100 GiB
volnfsstd	100 GiB
volnfsstdnew	100 GiB
zone1basic	6 TiB
zone2basic	100 GiB

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 acao-ora01-u01
 Capacity pool CapPool
 Service level Ultra
 Quota 100 GiB
 Encryption key source Microsoft.NetApp
 Availability Zone None

Networking

Virtual network ANFAVVal (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 acao-ora01-u01

Tags

database 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 AzAcSnap 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 `ConnectionString` 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 AzPasswdl
```

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 shows the Azure portal interface for a NetApp volume. The left sidebar shows the navigation menu with 'Volumes' selected. The main content area is titled 'acao-ora01-u01 (ANFAVSAcct/CapPool/acao-ora01-u01) | Snapshots'. It features a search bar and a table of snapshots.

Name	Location	Created
acao-ora01-data_2022-09-09T165255-8258852Z	South Central US	09/09/2022, 12:53:22 PM
acao-ora01-data_2022-09-12T160536-9809839Z	South Central US	09/12/2022, 12:05:55 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 shows the Azure NetApp Files dashboard. On the left, the 'Volumes' page is visible, listing various volumes including 'acao-ora01-u03'. The main pane shows the 'Snapshots' page for the selected volume. A table of snapshots is displayed, with the most recent snapshot, 'acao-ora01-log_2022-09-13T134502-044919Z', highlighted. A context menu is open over this snapshot, showing the 'Revert volume' option.

Name	Location	Created
acao-ora01-data_2022-09-12T160628-8754796Z	South Central US	09/12/2022, 12:06:31 PM
acao-ora01-log_2022-09-13T120122-8173645Z	South Central US	09/13/2022, 08:01:25 AM
acao-ora01-log_2022-09-13T121501-6760417Z	South Central US	09/13/2022, 08:15:04 AM
acao-ora01-log_2022-09-13T123001-8000548Z	South Central US	09/13/2022, 08:30:05 AM
acao-ora01-log_2022-09-13T124501-7472967Z	South Central US	09/13/2022, 08:45:04 AM
acao-ora01-log_2022-09-13T131501-6270448Z	South Central US	09/13/2022, 09:15:04 AM
acao-ora01-log_2022-09-13T133001-7773619Z	South Central US	09/13/2022, 09:30:04 AM
acao-ora01-log_2022-09-13T134502-044919Z	South Central US	09/13/2022, 09:45:04 AM


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.ct1
/u02/oradata/ORATST/control01.ct1.bk
[oracle@acao-ora01 ~]$ cp /u03/orareco/ORATST/control02.ct1
/u02/oradata/ORATST/control01.ct1
```

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](#).

On-Premises/Hybrid Cloud

TR-4996: Oracle SI Deployment and Protection in VCF with vVols

Allen Cao, Niyaz Mohamed, NetApp

The solution provides an overview and details for Oracle deployment and protection in VMware Cloud Foundation (VCF) with vSphere Virtual Volumes (vVols) as primary database storage and Oracle database in a single instance (SI) configuration..

Purpose

VMware vSphere Virtual Volumes (vVols) is a SAN/NAS management and integration framework that exposes virtual disks as native storage objects and enables array-based operations at the virtual disk level. In other words, vVols make SAN/NAS devices VM-aware and unlocks the ability to leverage array-based data services with a VM-centric approach at the granularity of a single virtual disk. vVols allows customers to leverage the unique capabilities of their current storage investments and transition without disruption to a simpler and more efficient operational model optimized for virtual environments that work across all storage types.

This documentation demonstrates the deployment and protection of an Oracle single instance database in a VMware Cloud Foundation environment with vVols as primary database storage in a NetApp ONTAP storage cluster. Oracle database is configured as if it is deployed in local file systems on a local storage system. This technical report focuses on steps in creating vVols in VCF for Oracle deployment. We also demonstrate how to use the NetApp SnapCenter UI tool to backup, restore, and clone an Oracle database for dev/test or other use cases for storage-efficient database operation in VCF.

This solution addresses the following use cases:

- Oracle SI database deployment in VCF with vVols datastore on NetApp ONTAP AFF as primary database storage
- Oracle database backup and restore in VCF with vVols datastore using NetApp SnapCenter UI tool
- Oracle database clone for dev/test or other use cases in VCF with vVols datastore using NetApp SnapCenter UI tool

Audience

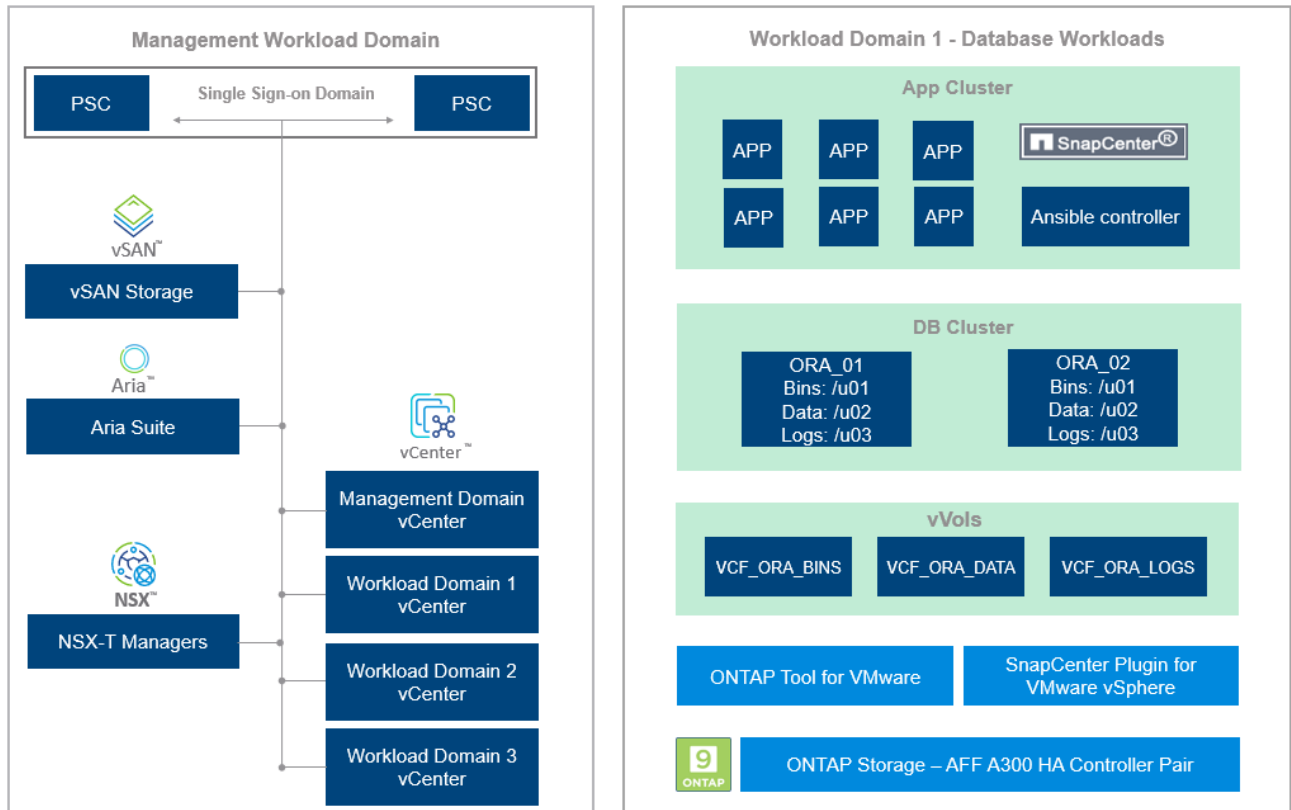
This solution is intended for the following people:

- A DBA who would like to deploy Oracle in VCF with vVols datastore on NetApp ONTAP AFF as primary database storage
- A database solution architect who would like to test Oracle workloads in VCF with vVols datastore on NetApp ONTAP AFF storage
- A storage administrator who would like to deploy and manage an Oracle database deployed to VCF with vVols datastore on NetApp ONTAP AFF storage
- An application owner who would like to stand up an Oracle database in VCF with vVol datastore

Solution test and validation environment

The testing and validation of this solution was performed in a lab environment with VCF with vVols datastore on NetApp ONTAP AFF storage that might not match the final deployment environment. For more information, see the section [Key factors for deployment consideration](#).

Oracle Single Instance Deployment and Protection in VCF with vVols



NetApp

Hardware and software components

Hardware		
NetApp ONTAP AFF A300	Version 9.14.1P4	DS224 shelf with 24 NVMe disks, total capacity 35.2 TiB
VMware VSphere cluster	Version 8.02	12 CPU(s) x Intel® Xeon® Gold 5218 CPU @ 2.30GHz, 8 nodes (4 management and 4 workload domains)
Software		
RedHat Linux	RHEL-8.6, 4.18.0-372.9.1.el8.x86_64 kernel	Hosting Oracle DB servers, deployed RedHat subscription for testing
Windows Server	2022 Standard, 10.0.20348 Build 20348	Hosting SnapCenter server
Centos Linux	CentOS Linux release 8.5.2111	Hosting Ansible controller
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 6.0	Workgroup deployment
SnapCenter Plug-in for VMware vSphere	Version 6.0	Deployed as an ova VM to vSphere cluster
ONTAP tool for VMware vSphere	Version 9.13	Deployed as an ova VM to vSphere cluster
Open JDK	Version java-11-openjdk-11.0.23.0.9-3.el8.x86_64	SnapCenter plugin requirement on DB VMs

Oracle database configuration in VCF

Server	Database	DB Storage
ora_01	NTAP1(NTAP1_pdb1,NTAP1_pdb2 ,NTAP1_pdb3)	vVols datastores on NetApp ONTAP AFF A300
ora_02	NTAP2(NTAP2_pdb1,NTAP2_pdb2 ,NTAP2_pdb3), NTAP1CLN	vVols datastores on NetApp ONTAP AFF A300

Key factors for deployment consideration

- **Protocol for vVols to ONTAP cluster connectivity.** NFS or iSCSI are good choices. The performance levels are equivalent. In this solution demonstration, we used NFS as a storage protocol for vVols connectivity to the underlined ONTAP storage cluster. If VCF infrastructure supports, FC/FCoE, NVMe/FC protocols are also supported for vVols datastores on NetApp ONTAP.
- **Oracle storage layout on vVols datastores.** In our tests and validations, we deployed three vVols datastores for Oracle binary, Oracle data, and Oracle log files. It's good practice to separate different types of Oracle files into their datastore so that database backup, recovery, or clone can be easily managed and executed. Create dedicate vVols for large databases and share vVols for smaller databases or databases with similar QoS profile.
- **Credential for ONTAP storage authentication.** Only use ONTAP cluster-level credentials for the ONTAP storage cluster authentication, including SnapCenter connectivity to the ONTAP storage cluster or ONTAP tool connectivity to the ONTAP storage cluster.
- **Provision storage from vVols datastore to database VM.** Add only one disk at a time to the database VM from the vVols datastore. Adding Multiple disks from vVols datastores at the same time is not supported at this time.
- **Database protection.** 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 Oracle 19c deployment in VCF with vVols datastores on NetApp ONTAP storage in an Oracle single instance configuration.

Prerequisites for deployment

Deployment requires the following prerequisites.

1. A VMware VCF has been setup. For information or instruction on how to create a VCF, please refer to VMware documentation [VMware Cloud Foundation Documentation](#).
2. Provision three Linux VMs, two VMs for Oracle database and one VM for Ansible controller within VCF workload domain. Provision one Windows server VM for running NetApp SnapCenter server. For information on setting up Ansible controller for automated Oracle database deployment, referring to following resources [Getting Started with NetApp solution automation](#).
3. SnapCenter plugin version 6.0 for VMware vSphere has been deployed in VCF. Refer to following resources for the plugin deployment: [SnapCenter Plug-in for VMware vSphere documentation](#).
4. ONTAP tool for VMware vSphere has been deployed in VCF. Refer to following resources for the ONTAP tool for VMware vSphere deployment: [ONTAP tools for VMware vSphere documentation](#)



Ensure that you have allocated at least 50G in Oracle VM root volume in order to have sufficient space to stage Oracle installation files.

Create storage capability profile


First, create a custom storage capability profile for the underlined ONTAP storage that is hosting the vVols datastore.

1. From vSphere client shortcuts, open NetApp ONTAP tool. Ensure that ONTAP storage cluster has been added to Storage Systems as part of ONTAP tool deployment.

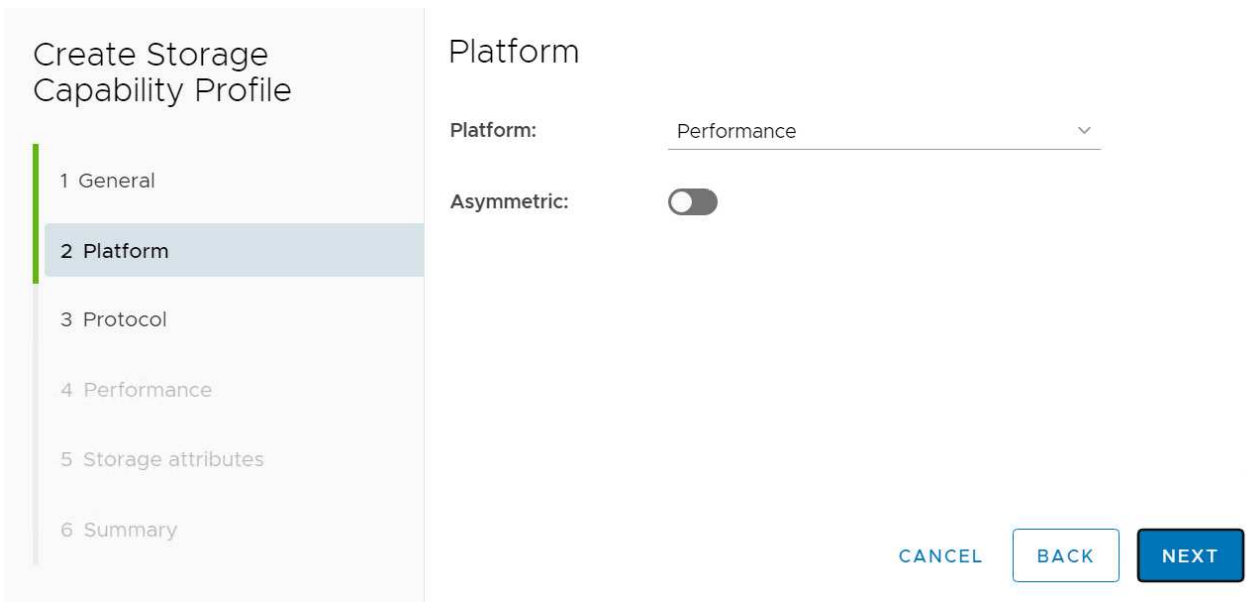
The screenshot shows the vSphere Client interface. The top navigation bar includes the vSphere Client logo and a search bar. Below the navigation bar, there are several sections: Shortcuts, Inventories, Monitoring, Plugins, and Administration. The Plugins section is highlighted, showing two plugins: SnapCenter Plug-in for VMware vSphere and NetApp ONTAP tools. The NetApp ONTAP tools plugin is highlighted in orange. Below the Plugins section, the Administration section is visible, showing the Licensing option. The bottom part of the screenshot shows the NetApp ONTAP tools interface. The top bar of this interface includes the vSphere Client logo, a search bar, and the instance ID: INSTANCE 172.21.166.149:8443. The main content area is titled "Storage Systems" and contains a table with the following data:

Name	Type	IP Address	ONTAP Release	Status	Capacity	NFS VAAI	Supported Protocols
ntaph0-4300x9u25	Cluster	172.16.9.25	9.14.1	Normal	43.76%		

2. Click on Storage capability profile to add a custom profile for Oracle. Name the profile and add a brief description.



3. Choose storage controller category: performance, capacity, or hybrid.



4. Select the protocol.

Create Storage Capability Profile

- 1 General
- 2 Platform
- 3 Protocol**
- 4 Performance
- 5 Storage attributes
- 6 Summary

Protocol

Protocol:

CANCEL

BACK

NEXT

5. Define a QoS policy if desired.

Create Storage Capability Profile

- 1 General
- 2 Platform
- 3 Protocol
- 4 Performance**
- 5 Storage attributes
- 6 Summary

Performance

- None ⓘ
- QoS policy group ⓘ

Min IOPS:

Max IOPS:

Unlimited

CANCEL

BACK

NEXT

6. Additional storage attributes for the profile. Be sure that the encryption is enabled on the NetApp controller if you want to have the encryption capability or it may cause issues when applying the profile.

Create Storage Capability Profile

- 1 General
- 2 Platform
- 3 Protocol
- 4 Performance
- 5 Storage attributes**
- 6 Summary

Storage attributes

Deduplication:	Yes	▼
Compression:	Yes	▼
Space reserve:	Thin	▼
Encryption:	Yes	▼
Tiering policy (FabricPool):	None	▼

CANCEL

BACK

NEXT

7. Review the summary and finish the storage capability profile creation.

Create Storage Capability Profile

- 1 General
- 2 Platform
- 3 Protocol
- 4 Performance
- 5 Storage attributes
- 6 Summary**

Summary

Name:	ASA_ORA
Description:	ASA for Oracle.
Platform:	Performance
Asymmetric:	No
Protocol:	Any
Performance:	None
Space reserve:	Thin
Deduplication:	Yes
Compression:	Yes
Encryption:	Yes
Tiering policy (FabricPool):	None

CANCEL

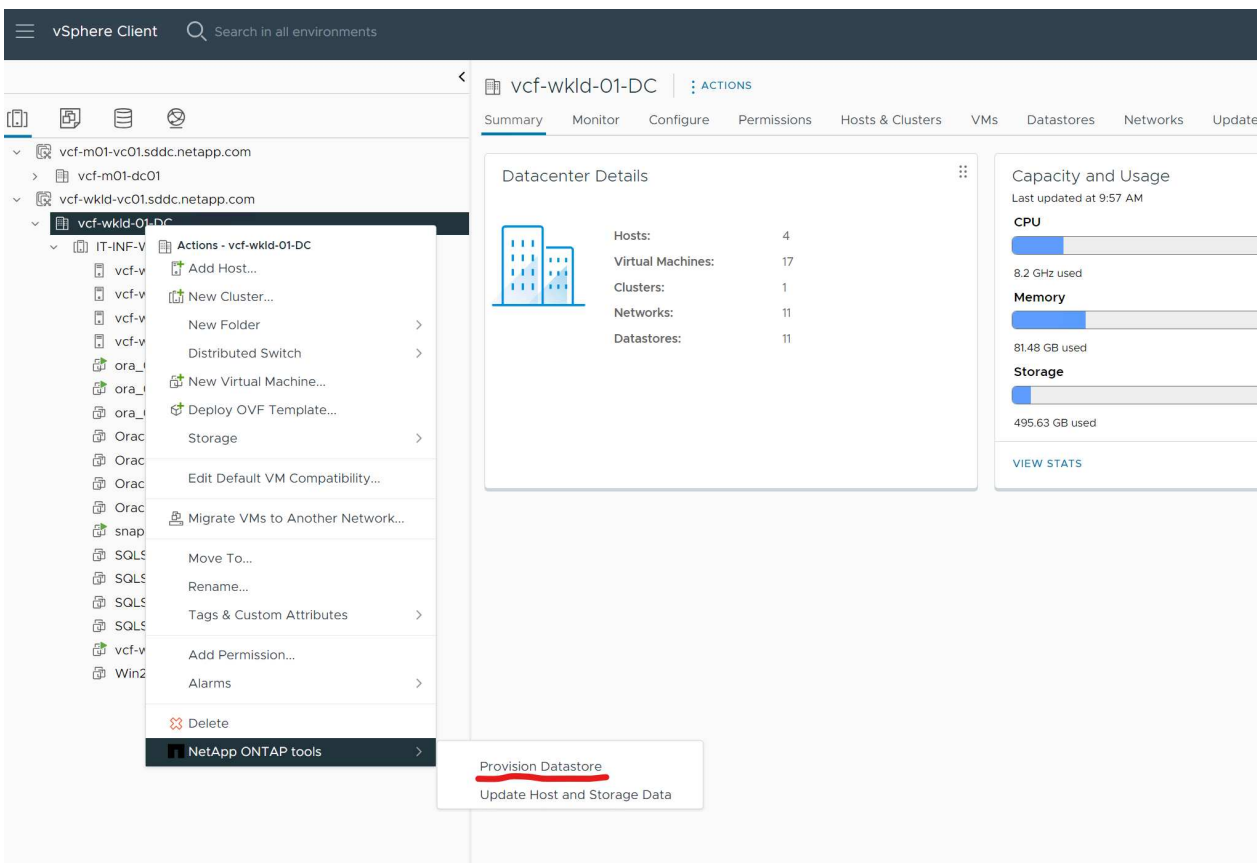
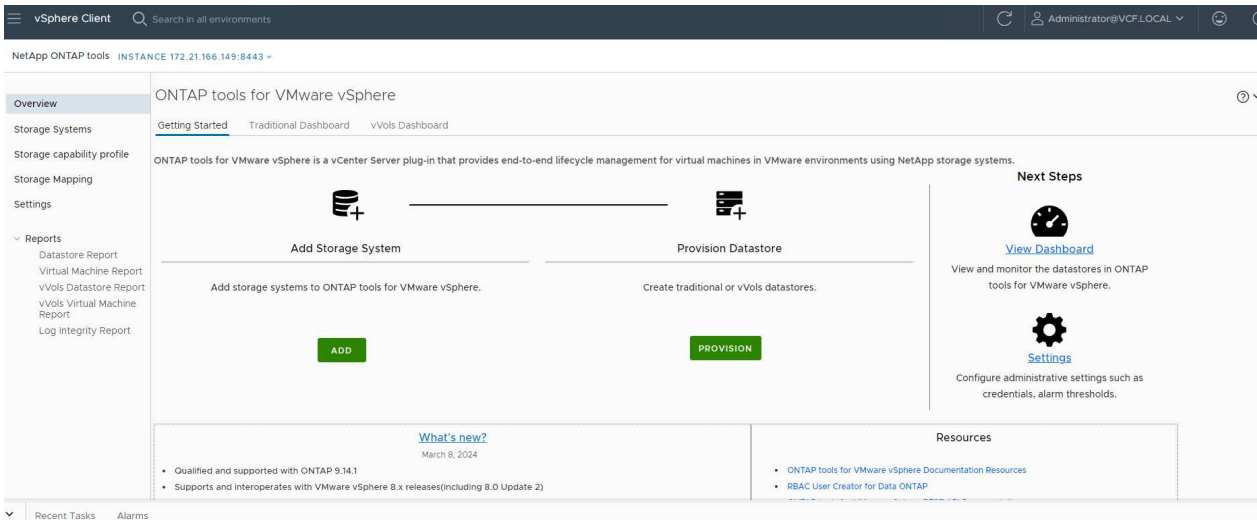
BACK

FINISH

Create and configure vVols datastore

With the prerequisites completed, login to the VCF as an admin user via vSphere client, navigating to workload domain. Do not use built-in VMware storage option to create vVols. Instead, use NetApp ONTAP tool to create vVols. Following demonstrates the procedures to create and configure vVols.

1. The vVols creation workflow can be triggered either from ONTAP tool interface or from VCF workload domain cluster.



2. Filling in general information for datastore including provisioning destination, type, name, and protocol.

New Datastore


1 General

2 Storage system

3 Storage attributes

4 Summary

General

Specify the details of the datastore to provision 

Provisioning destination: [BROWSE](#)

Type: NFS VMFS vVols

Name:

Description:

Protocol: NFS iSCSI FC / FCoE NVMe/FC

[CANCEL](#)

[NEXT](#)

3. Select the custom storage capability profile created from previous step, the Storage system, and Storage VM, where vVols are to be created.

New Datastore

1 General

2 Storage system

3 Storage attributes

4 Summary

Storage system

Specify the storage capability profiles and the storage system you want to use.

Storage capability profiles:

Custom profiles

Storage system:

Storage VM:

[CANCEL](#)

[BACK](#)

[NEXT](#)

4. Choose **Create new volumes**, fill in the volume name and size and click on **ADD** then **NEXT** to move to the summary page.

New Datastore

- 1 General
- 2 Storage system
- 3 Storage attributes
- 4 Summary

Storage attributes

Specify the storage details for provisioning the datastore.

Volumes: Create new volumes Select volumes

Create new volumes

Name	Size	Storage Capability Profile	Aggregate
 FlexVol volumes are not added.			

Name	Size(GB)	Storage capability profile	Aggregates	Space reserve
vcf_ora_bins	150	ASA_ORA	EHCAGgr02 - (17899.73 G)	Thin

ADD

CANCEL

BACK

NEXT



You could add more than one volume to a vVols datastore or span a vVols datastore volumes across ONTAP controller nodes for performance.

5. Click **Finish** to create vVols datastore for Oracle binary.

New Datastore

- 1 General
- 2 Storage system
- 3 Storage attributes
- 4 Summary

Summary

General

vCenter server:	vcf-wkld-vc01.sddc.netapp.com
Provisioning destination:	vcf-wkld-01-DC
Datastore name:	VCF_ORA_BINS
Datastore type:	vVols
Protocol:	NFS
Storage capability profile:	ASA_ORA

Storage system details

Storage system:	ntaphci-a300e9u25
SVM:	VCF_NFS

Storage attributes

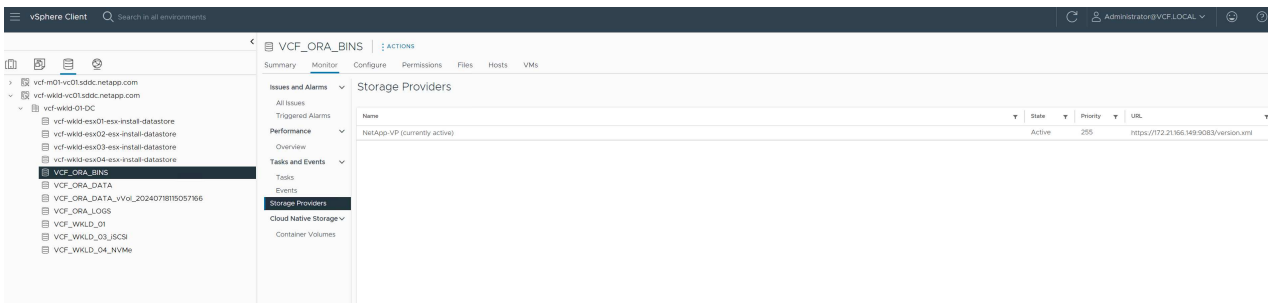
New FlexVol Name	New FlexVol Size	Aggregate	Storage Capability Profile
------------------	------------------	-----------	----------------------------

CANCEL

BACK

FINISH

6. Repeat the same procedures to create vVols datastore for Oracle data and log.



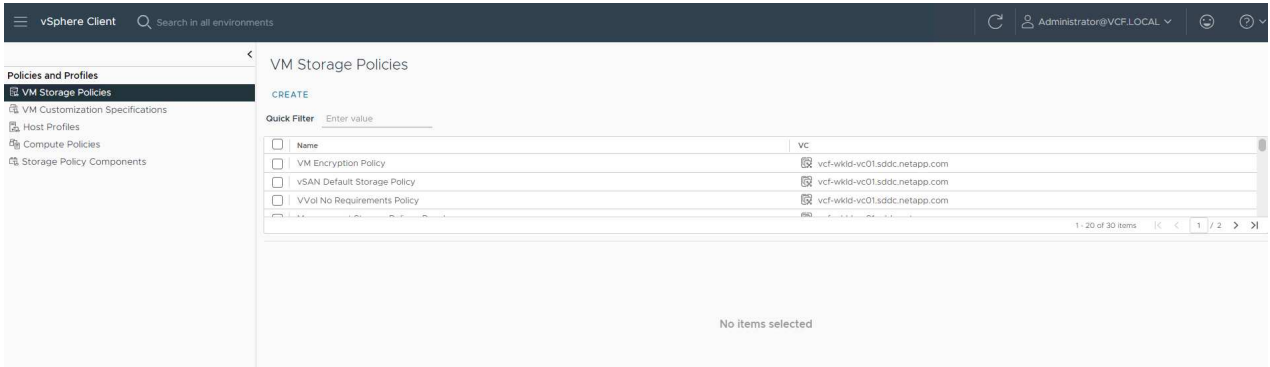


Notice that when an Oracle database is cloned, an additional vVols for Data is added to the vVols list.


Create VM storage policy based on storage capability profile

Before provisioning storage from vVols datastore to database VM, add a VM storage policy based on storage capability profile created from previous step. Following are the procedures.

1. From vSphere client menus, open Policies and Profiles and highlight VM Storage Policies. Click Create to open VM Storage Policies workflow.



2. Name the VM storage policy.



3. In Datastore specific rules, check Enable rules for "NetAPP.clustered.Data.ONTAP.VP.vvol" storage

Edit VM Storage Policy

1 Name and description

2 Policy structure

3 NetApp.clustered.Data.ONTAP.VP.vvol rules

4 Storage compatibility

5 Review and finish

Policy structure



Host based services

Create rules for data services provided by hosts. Available data services could include encryption, I/O control, caching, etc. Host based services will be applied in addition to any datastore specific rules.

Enable host based rules

Datastore specific rules

Create rules for a specific storage type to configure data services provided by the datastores. The rules will be applied when VMs are placed on the specific storage type.

Enable rules for "vSAN" storage

Enable rules for "vSANDirect" storage

Enable rules for "VMFS" storage

Enable rules for "NetApp.clustered.Data.ONTAP.VP.vvol" storage

Enable tag based placement rules

Storage topology

Create rules for storage consumption domain topology. The storage topology will be applied to all datastore specific rules.

Enable consumption domain

CANCEL

BACK

NEXT

4. For NetApp.clustered.Data.ONTAP.VP.vvol rules Placement, select the custom storage capacity profile created from the previous step.

Create VM Storage Policy

1 Name and description

2 Policy structure

3 **NetApp.clustered.Data.ONTAP.VP.vvol rules**

4 Storage compatibility

5 Review and finish

NetApp.clustered.Data.ONTAP.VP.vvol rules

×

Placement Replication Tags

ProfileName ⓘ

ASA_ORA

CANCEL

BACK

NEXT

5. For NetApp.clustered.Data.ONTAP.VP.vvol rules Replication, choose Disabled if vVols are not replicated.

Create VM Storage Policy

1 Name and description

2 Policy structure

3 **NetApp.clustered.Data.ONTAP.VP.vvol rules**

4 Storage compatibility

5 Review and finish

NetApp.clustered.Data.ONTAP.VP.vvol rules

×

Placement Replication Tags

Disabled

Custom

CANCEL

BACK

NEXT

6. Storage compatibility page displays the compatible vVols datastores in VCF environment.

Create VM Storage Policy

- 1 Name and description
- 2 Policy structure
- 3 NetApp.clustered.Data.ONTAP.VP.vvol rules
- 4 Storage compatibility**
- 5 Review and finish

Storage compatibility



COMPATIBLE INCOMPATIBLE

Expand datastore clusters

Compatible storage 650 GB (650 GB free)

Quick Filter

Name	Datacenter	Type	Free Space	Capacity	Warnings
VCF_ORA_BINS	vcf-wkld-01-DC	vVol	150.00 GB	150.00 GB	
VCF_ORA_DATA	vcf-wkld-01-DC	vVol	250.00 GB	250.00 GB	
VCF_ORA_LOGS	vcf-wkld-01-DC	vVol	250.00 GB	250.00 GB	

Manage Columns

3 Items

CANCEL

BACK

NEXT

7. Review and finish to create the VM Storage Policy.

Create VM Storage Policy

- 1 Name and description
- 2 Policy structure
- 3 NetApp.clustered.Data.ONTAP.VP.vvol rules
- 4 Storage compatibility
- 5 Review and finish

Review and finish

General

Name: vVol_database

Description: vCenter Server: vcf-wkld-vc01.sddc.netapp.com

NetApp.clustered.Data.ONTAP.VP.vvol rules

Placement: ProfileName: ASA_ORA

×

CANCEL
BACK
FINISH

8. Validate the VM Storage Policy just created.

The screenshot shows the vSphere Client interface. On the left, the 'Policies and Profiles' sidebar is visible, with 'VM Storage Policies' selected. The main area displays a list of storage policies. The 'vVol_database' policy is highlighted. Below the list, the details for the selected policy are shown:

Rules	VM Compliance	VM Template	Storage Compatibility
General			
Name	vVol_database		
Description	vCenter Server: vcf-wkld-vc01.sddc.netapp.com		
Rule-set 1: NetApp.clustered.Data.ONTAP.VP.vvol			
Placement	ProfileName: ASA_ORA		
Storage Type	NetApp.clustered.Data.ONTAP.VP.vvol		
ProfileName	ASA_ORA		

Allocate disks to DB VM from vVols datastores and configure DB storage

From vSphere client, add three disks from the vVols datastores to database VM by editing VM settings. Then, login to VM to format and mount the disks to mount points /u01, /u02, and /u03. The following demonstrates the exact steps and tasks.

1. Add a disk to VM for Oracle binary storage.

Edit Settings | ora_01

Virtual Hardware VM Options Advanced Parameters

ADD NEW DEVICE ▾

> CPU	4 ▾ ⓘ	
> Memory	16	GB ▾
> Hard disk 1	50	GB ▾
▾ New Hard disk *	50	GB ▾
Maximum Size	142.5 GB	
VM storage policy	vVol_database ▾	
Location	VCF_ORA_BINS ▾	
Disk Provisioning	Thin Provision ▾	
Sharing	No sharing ▾	
Disk Mode	Dependent ▾	

CANCEL

OK

2. Add a disk to VM for Oracle data storage.

Edit Settings | ora_01



Virtual Hardware | VM Options | Advanced Parameters

ADD NEW DEVICE ▾

> CPU	4 ▾	
> Memory	16	GB ▾
> Hard disk 1	50	GB ▾
> New Hard disk *	50	GB ▾
▼ New Hard disk 2 *	100	GB ▾
Maximum Size	475 GB	
VM storage policy	vVol_database ▾	
Location	VCF_ORA_DATA ▾	
Disk Provisioning	Thin Provision ▾	
Sharing	No sharing ▾	

CANCEL

OK

3. Add a disk to VM for Oracle log storage.

ADD NEW DEVICE ▾

> CPU	4 ▾ ⓘ		
> Memory	16	▾ GB ▾	
> Hard disk 1	50	GB ▾	⋮
> New Hard disk *	50	GB ▾	⋮
> New Hard disk 2 *	100	GB ▾	⋮
▾ New Hard disk 3 *	100	GB ▾	⋮
Maximum Size	285 GB		
VM storage policy	vVol_database ▾		
Location	VCF_ORA_LOGS ▾		
Disk Provisioning	Thin Provision ▾		
Sharing	No sharing ▾		

CANCEL

OK

4. From VM Edit Settings, Advanced Parameters, add Attribute `disk.enableuuid` with Value `TRUE`. The VM needs to be down to add the advanced parameter. Setting this option enables SnapCenter to precisely identify the vVol in your environment.

Advanced Configuration Parameters

Modify or add configuration parameters as needed for experimental features or as instructed by technical support. Empty values will be removed (supported on ESXi 6.0 and later).

Attribute Value

Attribute	Value
sched.cpu.latencySensitivity	normal
tools.guest.desktop.autoLock	TRUE
svga.present	TRUE
pciBridge0.present	TRUE
pciBridge4.present	TRUE
pciBridge4.virtualDev	pcieRootPort
pciBridge4.functions	8
pciBridge5.present	TRUE
pciBridge5.virtualDev	pcieRootPort
pciBridge5.functions	8
pciBridge6.present	TRUE

5. Now, restart the VM. Login to VM as an admin user via ssh to review the newly added disk drives.


```
[admin@ora_01 ~]$ sudo fdisk -l
```

```
Disk /dev/sdb: 50 GiB, 53687091200 bytes, 104857600 sectors  
Units: sectors of 1 * 512 = 512 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

```
Disk /dev/sdc: 100 GiB, 107374182400 bytes, 209715200 sectors  
Units: sectors of 1 * 512 = 512 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

```
Disk /dev/sdd: 100 GiB, 107374182400 bytes, 209715200 sectors  
Units: sectors of 1 * 512 = 512 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

```
.  
. .  
.
```

6. Partition the drives as a primary and single partition by simply accepting the default choices.

```
sudo fdisk /dev/sdb
```

```
sudo fdisk /dev/sdc
```

```
sudo fdisk /dev/sdd
```

7. Format the partitioned disks as xfs file systems.

```
sudo mkfs.xfs /dev/sdb1
```

```
sudo mkfs.xfs /dev/sdc1
```

```
sudo mkfs.xfs /dev/sdd1
```

8. Mount the drives to mount point /u01, /u02, and /u03.

```
sudo mount -t xfs /dev/sdb1 /u01
```

```
sudo mount -t xfs /dev/sdc1 /u02
```

```
sudo mount -t xfs /dev/sdd1 /u03
```

```
[admin@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	782M	7.0G	10%	/run
tmpfs	7.8G	0	7.8G	0%	/sys/fs/cgroup
/dev/mapper/rhel-root	44G	19G	26G	43%	/
/dev/sda1	1014M	258M	757M	26%	/boot
tmpfs	1.6G	12K	1.6G	1%	/run/user/42
tmpfs	1.6G	4.0K	1.6G	1%	/run/user/1000
/dev/sdb1	50G	390M	50G	1%	/u01
/dev/sdc1	100G	746M	100G	1%	/u02
/dev/sdd1	100G	746M	100G	1%	/u03

9. Add mount points to `/etc/fstab` so that disk drives will be mounted when VM reboots.

```
sudo vi /etc/fstab
```

```

[oracle@ora_01 ~]$ cat /etc/fstab

#
# /etc/fstab
# Created by anaconda on Wed Oct 18 19:43:31 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/rhel-root / xfs defaults
0 0
UUID=aff942c4-b224-4b62-807d-6a5c22f7b623 /boot
xfs defaults 0 0
/dev/mapper/rhel-swap none swap defaults
0 0
/root/swapfile swap swap defaults 0 0
/dev/sdb1 /u01 xfs defaults
0 0
/dev/sdc1 /u02 xfs defaults
0 0
/dev/sdd1 /u03 xfs defaults
0 0

```

Oracle database deployment in VCF

It's recommended to leverage NetApp automation toolkit to deploy Oracle in VCF with vVols. For detailed reference on executing Oracle automated deployment on xfs file systems, refer to TR-4992: [Simplified, Automated Oracle Deployment on NetApp C-Series with NFS](#). Although the TR-4992 covers automated Oracle deployment on the NetApp C-Series with NFS, it's identical to Oracle deployment in VCF with vVols if bypassing NFS file systems mounting to database VM. We would simply skip that with specific tags. Following are step by step procedures.

1. Login to Ansible controller VM as admin user via ssh and clone a copy of automation toolkit for Oracle on NFS.

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

2. Stage the following Oracle installation files in /tmp/archive folder on database VM. The folder should allow all user access with 777 permission.

```
LINUX.X64_193000_db_home.zip  
p34765931_190000_Linux-x86-64.zip  
p6880880_190000_Linux-x86-64.zip
```

3. Configure deployment target file - hosts, global variables file - vars/vars.yml, and local DB VM variables file - host_vars/host_name.yml according to instructions in this section of TR-4992: [Parameter files configuration](#). Comment out nfs_lif variable from local DB VM variable file.
4. Set up ssh keyless authentication between Ansible controller and database VMs, which requires to generate a ssh key pair and copy the public key to database VMs admin user root directory .ssh folder authorized_keys file.

```
ssh-keygen
```

5. From Ansible controller, cloned automation toolkit home directory /home/admin/na_oracle_deploy_nf, execute prerequisites playbook.

```
ansible-playbook -i hosts 1-ansible_requirements.yml
```

6. Execute Linux configuration playbook.

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

7. Execute Oracle deployment playbook.

```
ansible-playbook -i hosts 4-oracle_config.yml -u admin -e
@vars/vars.yml --skip-tags "ora_mount_points,enable_dnfs_client"
```

8. Optionally, all above playbooks can be executed from a single playbook run as well.

```
ansible-playbook -i hosts 0-all_playbook.yml -u admin -e
@vars/vars.yml --skip-tags "ora_mount_points,enable_dnfs_client"
```

9. Login to EM express to to validate Oracle after successful playbook execution.

The screenshot shows the Oracle Enterprise Manager Database Express interface. The top section is a blue header with the Oracle logo and the text "ORACLE ENTERPRISE MANAGER DATABASE EXPRESS". Below this is a login form with fields for Username (system), Password (masked), and Container Name, and a "Log In" button. The bottom section shows the dashboard for a database instance (NTAP1). The dashboard includes a "Status" panel with details like Up Time (6 days 3 hours 17 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), and Incident(s) (0). The "Performance" panel shows a line graph for CPU usage over time. The "Resources" panel contains four bar charts: Host CPU (2%), Active Sessions (0.21), Memory (14 GB), and Data Storage (3.7 GB). The "SQL Monitor - Last Hour (20 max)" panel shows a table with columns for Status, Duration, SQL ID, SQL Plan Hash, User Name, Parallel, Database Time, I/O Requests, and SQL Text.

10. Optionally, execute destroy playbook to remove database from DB VM.

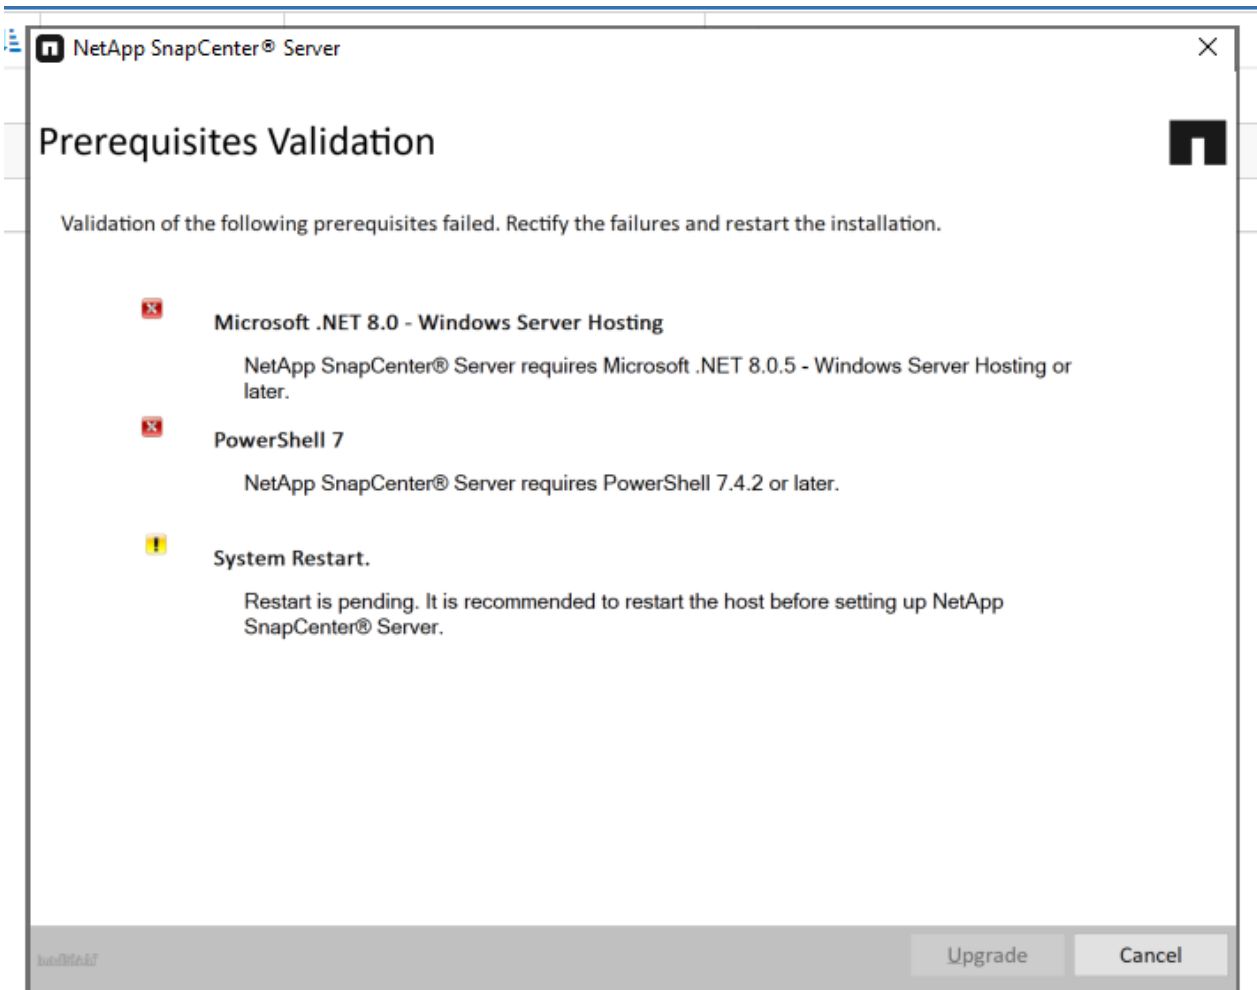
```
ansible-playbook -i hosts 5-destroy.yml -u admin -e @vars/vars.yml
```

Oracle backup, restore, and clone in VCF with SnapCenter

SnapCenter Setup

SnapCenter version 6 has many feature enhancements over version 5, including support for VMware vVols datastore. SnapCenter relies on a host-side plug-in on a database VM to perform application-aware data protection management activities. For detailed information on NetApp SnapCenter plug-in for Oracle, refer to this documentation [What can you do with the Plug-in for Oracle Database](#). The following provides high-level steps to set up SnapCenter version 6 for Oracle database backup, recovery, and clone in VCF.

1. Download the version 6 of SnapCenter software from NetApp support site: [NetApp Support Downloads](#).
2. Login to the SnapCenter hosting Windows VM as administrator. Install prerequisites for SnapCenter 6.0.




3. As administrator, install latest java JDK from [Get Java for desktop applications](#).




If Windows server is deployed in a domain environment, add a domain user to SnapCenter server local administrators group and run SnapCenter installation with the domain user.

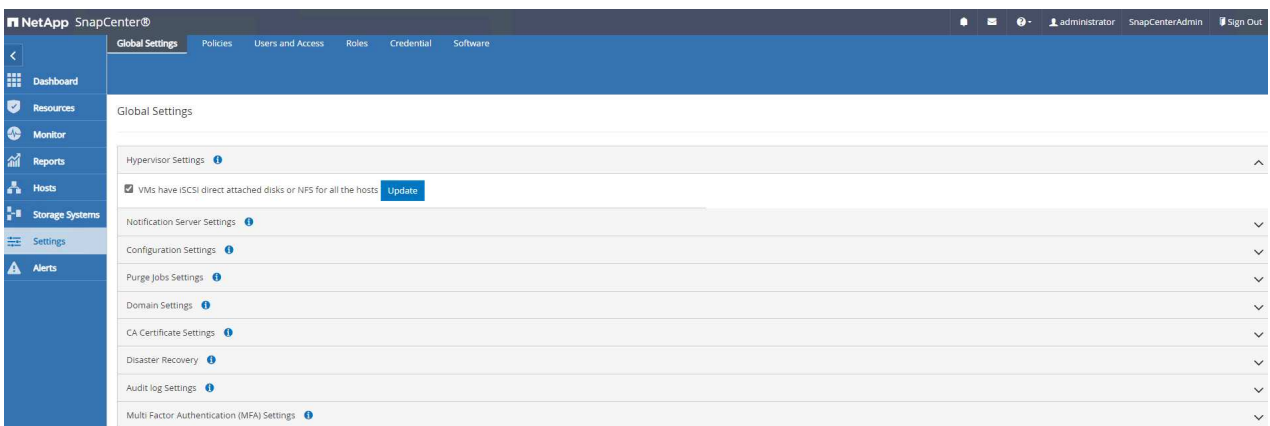
4. Login to SnapCenter UI via HTTPS port 8846 as installation user to configure SnapCenter for Oracle.





5. Review Get Started menu to get up to speed on SnapCenter if you are a new user.




6. Update Hypervisor Settings in global settings.




7. Add ONTAP storage cluster to Storage Systems with cluster management IP and authenticated via cluster admin user ID.



8. Add database VM and vSphere plugin VM `Credential` for SnapCenter access to DB VM and vSphere plugin VM. The credential should have sudo privilege on the Linux VMs. You may create different credentials for different management user IDs for the VMs.



9. Add Oracle database VM in VCF to `Hosts` with DB VM credential created in previous step.



Confirm Fingerprint ✕

Authenticity of the host cannot be determined i

Host name	Fingerprint	Valid
ora_01.sddc.netapp.com	ssh-ed25519 256 36:60:E6:D0:2B:F2:F8:90:97:A1:D3:83:40:7E:E3:7A	

Confirm and Submit
Close

NetApp SnapCenter®

Managed Hosts | Disks | Shares | Initiator Groups | iSCSI Session

Search by Name

<input type="checkbox"/>	Name	Type	System	Plug-in	Version	Overall Status
<input type="checkbox"/>	ora_01.sddc.netapp.com	Linux	Stand-alone	Oracle Database, UNIX	6.0	Running
<input type="checkbox"/>	ora_02.sddc.netapp.com	Linux	Stand-alone	Oracle Database, UNIX	6.0	Running

View all the alerts and its details.

- Similarly, add NetApp VMware plugin VM to `Hosts` with vSphere plugin VM credential created in previous step.

NetApp SnapCenter®

Managed Hosts

Search by Name

Add Host

Host Type: vSphere

Host Name: 172.21.166.14

Credentials: admin

Submit Cancel

NetApp SnapCenter®

Managed Hosts | Disks | Shares | Initiator Groups | iSCSI Session

Search by Name

<input type="checkbox"/>	Name	Type	System	Plug-in	Version	Overall Status
<input type="checkbox"/>	ora_01.sddc.netapp.com	Linux	Stand-alone	Oracle Database, UNIX	6.0	Running
<input type="checkbox"/>	ora_02.sddc.netapp.com	Linux	Stand-alone	Oracle Database, UNIX	6.0	Running
<input type="checkbox"/>	vef001.sv0.sddc.netapp.com	vSphere	Stand-alone	VMware vSphere	6.0	Running

- Finally, after Oracle database is discovered on DB VM, back to `Settings-Policies` to create Oracle database backup policies. Ideally, create a separate archive log backup policy to allow more frequent backup interval to minimize data loss in the event of a failure.

The screenshot shows the NetApp SnapCenter interface. The top navigation bar includes 'Global Settings', 'Policies', 'Users and Access', 'Roles', 'Credential', and 'Software'. The left sidebar contains 'Dashboard', 'Resources', 'Monitor', 'Reports', 'Hosts', 'Storage Systems', 'Settings', and 'Alerts'. The main content area displays a table of backup policies for 'Oracle Database'.

Name	Backup Type	Schedule Type	Replication	Verification
Oracle Archive Logs Backup	LOG, ONLINE	Hourly		
Oracle Online Full Backup	FULL, ONLINE	Hourly		




Ensure that the SnapCenter server name can be resolved to the IP address from the DB VM and vSphere plugin VM. Equally, the DB VM name and vSphere plugin VM name can be resolved to the IP address from the SnapCenter server.

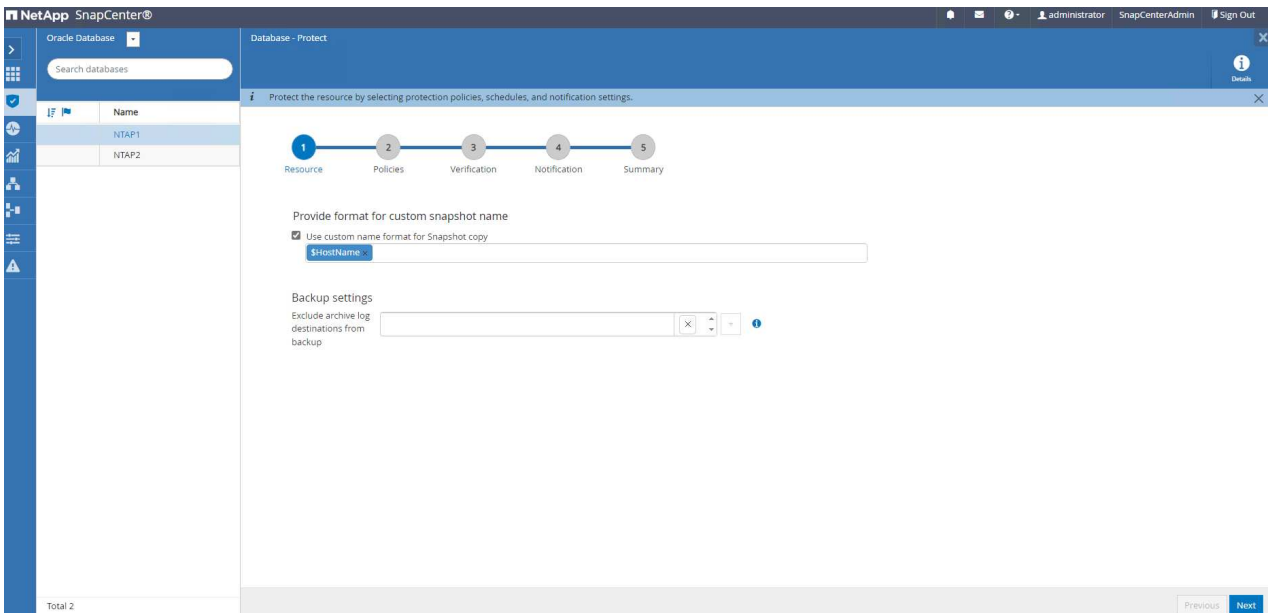
Database backup

SnapCenter leverages ONTAP volume snapshot for much quicker database backup, restore, or clone compared with traditional RMAN based methodology. The snapshots are application-consistent as the database is put in Oracle backup mode before a snapshot.


1. From the Resources tab, any databases on the VM are auto-discovered after the VM is added to SnapCenter. Initially, the database status shows as Not protected.



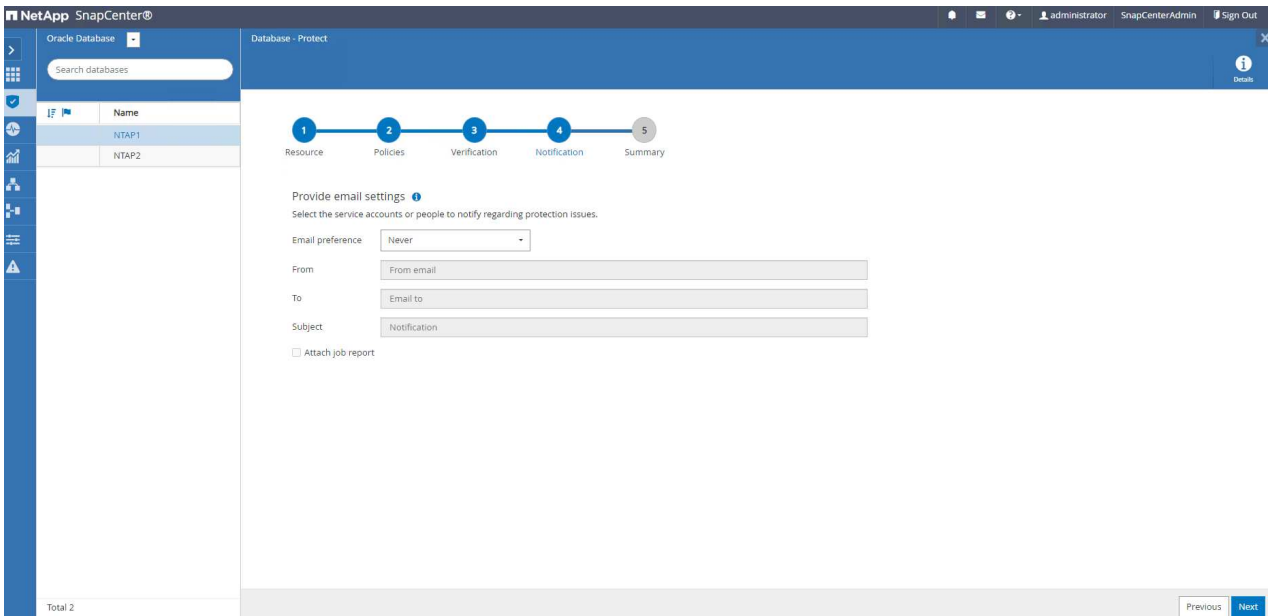
2. Click on database to start a workflow to enable protection for the database.




3. Apply backup policy, setup scheduling if needed.




4. Setup backup job notification if required.



5. Review the summary and finish to enable database protection.



6. On-demand backup job can be triggered by click on **Back up Now**.



Backup



Create a backup for the selected resource

Resource Name

NTAP1

Policy

Oracle Online Backup



Verify after backup

Cancel

Backup

7. The backup job can be monitored at the `Monitor` tab by clicking on the running job.

Job Details



Backup of Resource Group 'ora_01_sddc_netapp_com_NTAP1' with policy 'Oracle Online Backup'

✓ ▾ Backup of Resource Group 'ora_01_sddc_netapp_com_NTAP1' with policy 'Oracle Online Backup'

✓ ▾ ora_01.sddc.netapp.com

- ✓ ▶ Prescripts
- ✓ ▶ Preparing for Oracle Database Backup
- ✓ ▶ Preparing for File-System Backup
- ✓ ▶ Backup datafiles and control files
- ✓ ▶ Backup archive logs
- ✓ ▶ Finalizing Oracle Database Backup
- ✓ ▶ Finalizing File-System Backup
- ✓ ▶ Postscripts
- ✓ ▶ Data Collection
- ✓ ▶ Send EMS Messages

📘 Task Name: ora_01.sddc.netapp.com Start Time: 07/16/2024 5:39:33 PM End Time: 07/16/2024 5:40:23 PM

View Logs

Cancel Job

Close

8. Click on database to review the backup sets completed for each database.

Manage Copies

4 Backups
1 Clone
Local copies

Summary Card

4 Backups
2 Data Backups
2 Log Backups
1 Clone
0 Snapshots Locked

Primary Backup(s)

search

Backup Name	Snapshot Lock Expiration	Count	Type	End Date	Verified	Mounted	RMAN Cataloged	SCN
ora_01_07-18-2024_11.17.20.8165_1		1	Log	07/18/2024 11:17:55 AM	Not Applicable	False	Not Cataloged	2874360
ora_01_07-18-2024_11.17.20.8165_0		1	Data	07/18/2024 11:17:41 AM	Unverified	False	Not Cataloged	2874313
ora_01_07-18-2024_11.09.08.6002_1		1	Log	07/18/2024 11:09:44 AM	Not Applicable	False	Not Cataloged	2873909
ora_01_07-18-2024_11.09.08.6002_0		1	Data	07/18/2024 11:09:30 AM	Unverified	False	Not Cataloged	2873861

Database restore/recovery

SnapCenter provides a number of restore and recovery options for Oracle databases from snapshot backup. In this example, we demonstrate to restore from a older snapshot backup, then roll forward the database to the last available log.

1. First, run a snapshot backup. Then, create a test table and insert a row into table to validate that recovered database from snapshot image before test table creation regains the test table.

```
[oracle@ora_01 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Wed Jul 17 10:20:10
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> sho 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;

SQL> select * from test;

no rows selected

SQL> insert into test values (1, sysdate, 'test oracle
backup/restore/clone on VMware Cloud Foundation vVols');

1 row created.

SQL> commit;

Commit complete.

SQL> select * from test;
```

```

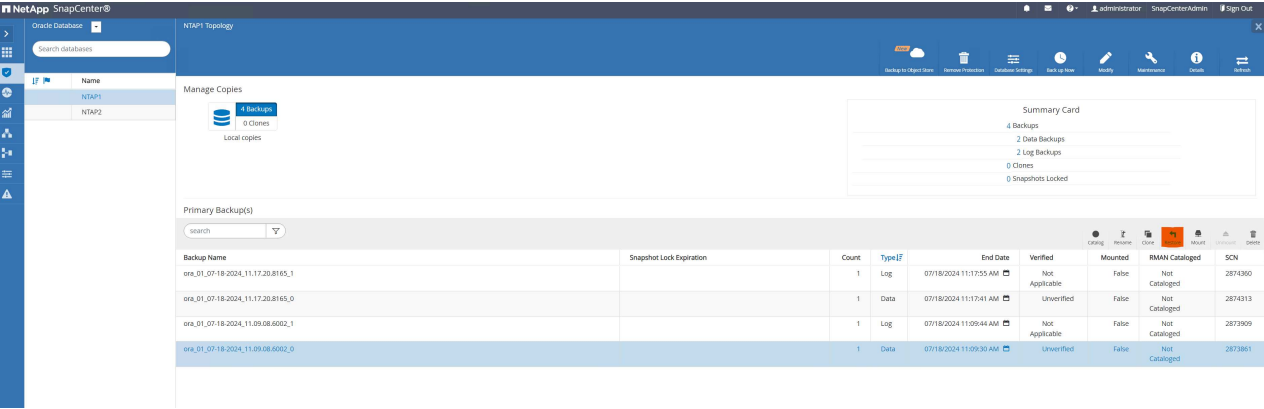
ID
-----
DT
-----
EVENT
-----

1
18-JUL-24 11.15.03.000000 AM
test oracle backup/restore/clone on VMware Cloud Foundation vVols

SQL>

```

2. From SnapCenter Resources tab, open the database NTAP1 backup topology page. Highlight the snapshot data backup set before the creation of test table. Click on Restore to launch restore-recover workflow.



3. Choose restore scope.

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Restore Scope ⓘ

- All Datafiles
- Pluggable databases (PDBs)
- Pluggable database (PDB) tablespaces

Control files

Database State

Change database state if needed for restore and recovery

Restore Mode ⓘ

Force in place restore

In place restore will skip the foreign files(files which are not part of the database) validation check. The Oracle database and the ASM disk group will be restored to the point when the backup was created.

Previous

Next

4. Choose recovery scope to All Logs.

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Choose Recovery Scope

- All Logs ?
- Until SCN (System Change Number)
- Date and Time
- No recovery

Specify external archive log files locations + - ?

Previous

Next

5. Specify any optional pre-scripts to run.

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Specify optional scripts to run before performing a restore job ⓘ

Prescript full path

Arguments

Script timeout

Previous

Next

6. Specify any optional after-script to run.

Restore NTAP1



1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Specify optional scripts to run after performing a restore job ⓘ

Postscript full path

Arguments

Open the database or container database in READ-WRITE mode after recovery

Previous

Next

7. Send a job report if desired.

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Provide email settings ⓘ

Email preference

From

To

Subject

Attach job report

Previous

Next

8. Review the summary and click on `Finish` to launch the restoration and recovery.

Restore NTAP1



1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Summary

Backup name	ora_01_07-16-2024_17.39.32.7534_0
Backup date	07/16/2024 5:40:02 PM
Restore scope	All DataFiles
Recovery scope	All Logs
Options	Change database state if necessary , Open the database or container database in READ-WRITE mode after recovery
Prescript full path	None
Prescript arguments	
Postscript full path	None
Postscript arguments	
Send email	No

Previous

Finish

9. From Monitor tab, open the job to review the details.

Job Details



Restore 'ora_01.sddc.netapp.com\NTAP1'

- ✓ ▾ Restore 'ora_01.sddc.netapp.com\NTAP1'
- ✓ ▾ ora_01.sddc.netapp.com
 - ✓ ▶ Prescripts
 - ✓ ▶ Mount log backups
 - ✓ ▶ Pre Restore
 - ✓ ▶ Restore
 - ✓ ▶ Post Restore
 - ✓ ▶ Unmount log backups
 - ✓ ▶ Postscripts
 - ✓ ▶ Post Restore Cleanup
 - ✓ ▶ Data Collection
 - ✓ ▶ Send EMS Messages

i Task Name: ora_01.sddc.netapp.com Start Time: 07/18/2024 11:26:50 AM End Time: 07/18/2024 11:40:25 AM

View Logs

Cancel Job

Close

10. From DB VM ora_01, validate that a successful restore/recovery of database rolled forward to its most recent state and recovered the test table.

```
[oracle@ora_01 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Thu Jul 18 11:42:58
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 from v$database;

NAME          OPEN_MODE
-----
NTAP1         READ WRITE

SQL> alter session set container=ntap1_pdb1;

Session altered.

SQL> select * from test;


          ID
-----
DT
-----
EVENT
-----
          1
18-JUL-24 11.15.03.000000 AM
test oracle backup/restore/clone on VMware Cloud Foundation vVols

SQL>
```

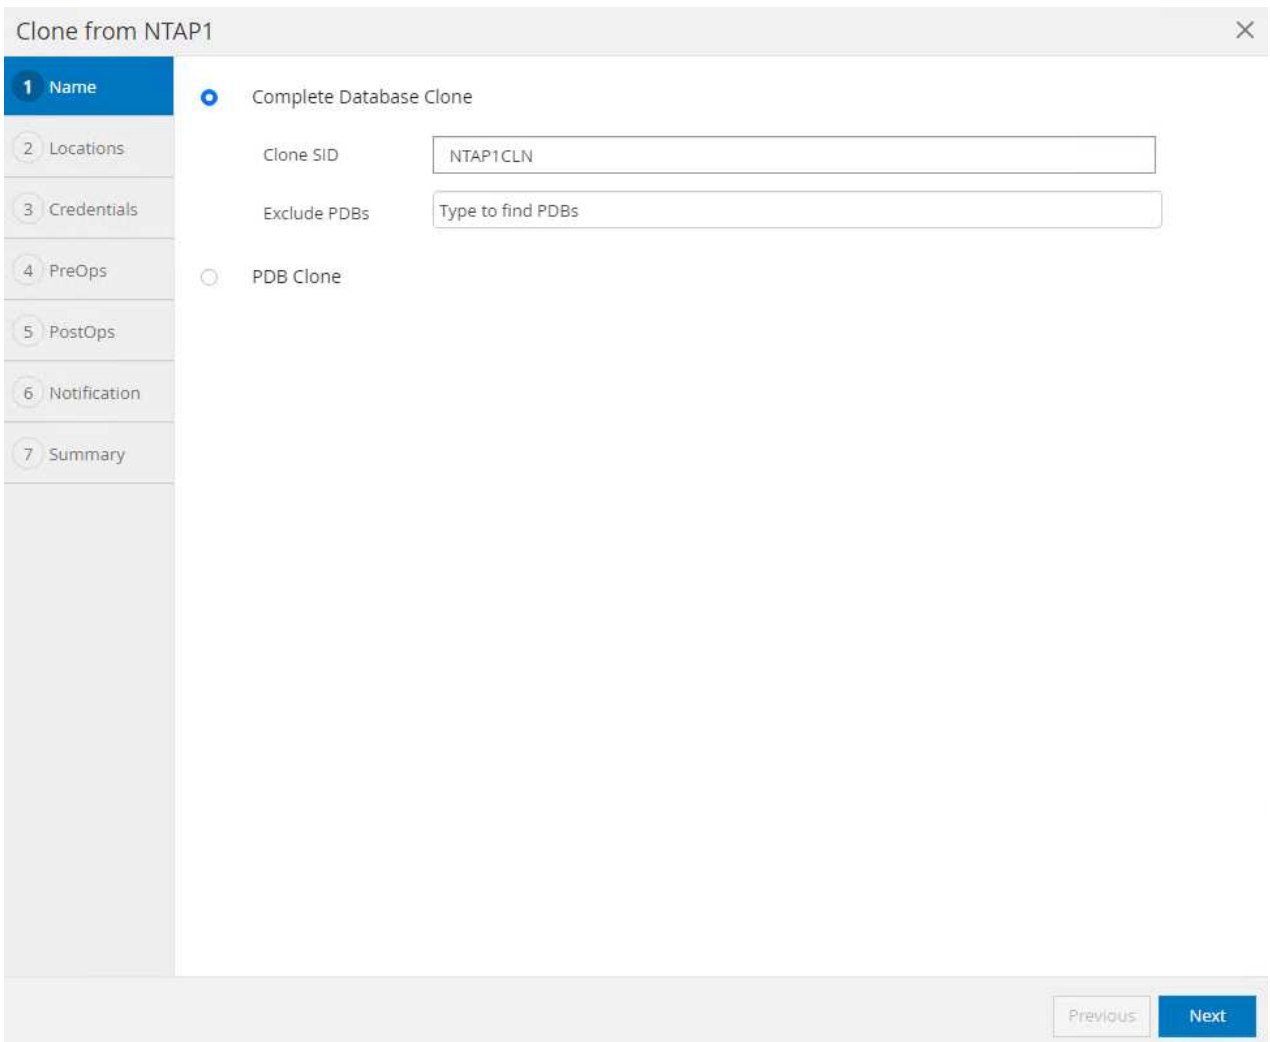
Database clone

In this example, the most recent backup sets is used to clone a database on the DB VM ora_02 in a different software installation and ORACLE_HOME in VCF.

1. Again, open the database NTAP1 backup list. Select the most recent data backup set, click on Clone button to launch database clone workflow.



2. Name the clone database SID.



3. Select ora_02 in VCF as the target database clone host. Identical Oracle database software should have been installed and configured on the host.

✕
Clone from NTAP1

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Select the host to create a clone

Clone host

☏ Datafile locations ℹ

 Reset

☏ Control files ℹ

/u02_NTAP1CLN/NTAP1CLN/control/control01.ctl	✕	↑	+
/u02_NTAP1CLN/NTAP1CLN/control/control02.ctl	✕	↓	Reset

☏ Redo logs ℹ

Group	Size	Unit	Number of files			
▶ RedoGroup 1	200	MB	1	✕	↑	+
▶ RedoGroup 2	200	MB	1	✕	↓	+
▶ RedoGroup 3	200	MB	1	✕	↓	+

Previous Next

4. Select the proper ORACLE_HOME, user and group on the target host. Keep credential at default.

1 Name

Database Credentials for the clone

2 Locations

Credential name for sys user

None



3 Credentials

Database port

1521

4 PreOps

5 PostOps

Oracle Home Settings i

6 Notification

Oracle Home

/u01/app/oracle/product/19.0.0/NTAP2

7 Summary

Oracle OS User

oracle

Oracle OS Group

oinstall

Previous

Next

5. You may change clone database parameters to meet configuration or resources requirements for the clone database.

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Specify scripts to run before clone operation ⓘ

Prescript full path Arguments Script timeout

Database Parameter settings

processes	320	✕	▲
remote_login_passwordfile	EXCLUSIVE	✕	+
sga_target	4G	✕	
undo_tablespace	UNDOTBS1	✕	▼

Previous

Next

- Choose recovery scope. `Until Cancel` recovers the clone up to last available log file in the backup set.

Clone from NTAP1



1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Recover Database

Until Cancel




Date and Time



Date-time format: MM/DD/YYYY hh:mm:ss

Until SCN (System Change Number)



Specify external archive log locations   

Create new DBID 

Create tempfile for temporary tablespace 

 Enter SQL queries to apply when clone is created

 Enter scripts to run after clone operation 

Previous

Next

7. Review the summary and launch the clone job.

Clone from NTAP1



1 Name	Summary	
2 Locations	Clone from backup	ora_01_07-18-2024_11.17.20.8165_0
3 Credentials	Clone SID	NTAP1CLN
4 PreOps	Clone server	ora_02.sddc.netapp.com
5 PostOps	Exclude PDBs	none
6 Notification	Oracle home	/u01/app/oracle/product/19.0.0/NTAP2
7 Summary	Oracle OS user	oracle
	Oracle OS group	oinstall
	Datafile mountpaths	/u02_NTAP1CLN
	Control files	/u02_NTAP1CLN/NTAP1CLN/control/control01.ctl /u02_NTAP1CLN/NTAP1CLN/control/control02.ctl
	Redo groups	RedoGroup =1 TotalSize =200 Path =/u02_NTAP1CLN/NTAP1CLN/redolog/redo01_01.log RedoGroup =2 TotalSize =200 Path =/u02_NTAP1CLN/NTAP1CLN/redolog/redo02_01.log RedoGroup =3 TotalSize =200 Path =/u02_NTAP1CLN/NTAP1CLN/redolog/redo03_01.log
	Recovery scope	Until Cancel
	Prescript full path	none
	Prescript arguments	
	Postscript full path	none
	Postscript arguments	
	Send email	No

Previous

Finish

8. Monitor the clone job execution from **Monitor** tab.

Job Details



Clone from backup 'ora_01_07-18-2024_11.17.20.8165_0'

✔ ▼ Clone from backup 'ora_01_07-18-2024_11.17.20.8165_0'

✔ ▼ ora_02.sddc.netapp.com

- ✔ ▶ Prescripts
- ✔ ▶ Query Host Information
- ✔ ▶ Prepare for Cloning
- ✔ ▶ Cloning Resources
- ✔ ▶ FileSystem Clone
- ✔ ▶ Application Clone
- ✔ ▶ Postscripts
- ✔ ▶ Register Clone
- ✔ ▶ Unmount Clone
- ✔ ▶ Data Collection

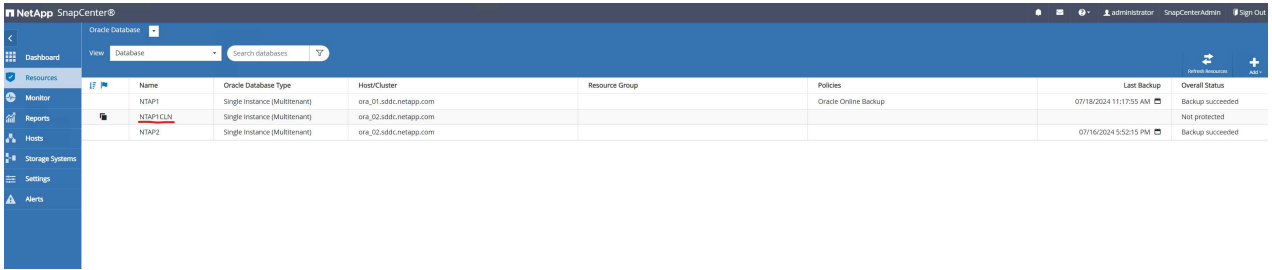
📘 Task Name: ora_02.sddc.netapp.com Start Time: 07/18/2024 11:50:41 AM End Time: 07/18/2024 12:02:34 PM

View Logs

Cancel Job

Close

9. Cloned database is immediately registered in SnapCenter.



The screenshot shows the NetApp SnapCenter interface. The main content area displays a table of Oracle Database resources. The table has columns for Name, Oracle Database Type, Host/Cluster, Resource Group, Policies, Last Backup, and Overall Status. The resources listed are NTAP1, NTAP1CLN, and NTAP2.

Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Overall Status
NTAP1	Single Instance (Multitenant)	ora_01.adfnc.netapp.com		Oracle Online Backup	07/18/2024 11:17:55 AM	Backup succeeded
NTAP1CLN	Single Instance (Multitenant)	ora_02.adfnc.netapp.com				Not protected
NTAP2	Single Instance (Multitenant)	ora_02.adfnc.netapp.com			07/16/2024 5:52:15 PM	Backup succeeded

10. From DB VM ora_02, validate the cloned database and query the test table.

```
[oracle@ora_02 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Thu Jul 18 12:06:48
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
-----
NTAP1CLN     READ WRITE        ARCHIVELOG

SQL> select instance_name, host_name from v$instance;

INSTANCE_NAME
-----
HOST_NAME
-----
NTAP1CLN
ora_02

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
      2 ;

Session altered.

SQL> select * from test;

          ID
-----
DT
-----
EVENT
-----
          1
18-JUL-24 11.15.03.000000 AM
test oracle backup/restore/clone on VMware Cloud Foundation vVols

SQL>

```

This completes the demonstration of SnapCenter backup, restore, and clone of Oracle database in VCF.

Where to find additional information

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

- VMware Cloud Foundation

<https://www.vmware.com/products/cloud-infrastructure/vmware-cloud-foundation>

- NetApp Enterprise Database Solutions

<https://docs.netapp.com/us-en/netapp-solutions/databases/index.html>

- SnapCenter Software 6.0

https://docs.netapp.com/us-en/snapcenter/concept/concept_snapcenter_overview.html

TR-4992: Simplified, Automated Oracle Deployment on NetApp C-Series with NFS

Allen Cao, Niyaz Mohamed, NetApp

This solution provides an overview and details for automated Oracle deployment in NetApp AFF C-Series as primary database storage with NFS protocol. The Oracle database deploys as a container database with dNFS enabled.

Purpose

NetApp AFF C-Series is a capacity flash storage that makes all-flash more accessible and affordable for unified storage. It is sufficient performance-wise for many tier 1 or tier 2 Oracle database workloads. Powered by NetApp ONTAP® data management software, AFF C-Series systems deliver industry-leading efficiency, superior flexibility, best-in-class data services, and cloud integration to help you scale your IT infrastructure, simplify your data management, and reduce storage cost and power consumption.

This documentation demonstrates the simplified deployment of Oracle databases in NetApp C-Series 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 solution provides the best practices in setting up storage networking and storage virtual machine (SVM) with NFS protocol on C-Series storage controllers. The solution also includes information on fast Oracle database backup, restore, and clone with the NetApp SnapCenter UI tool.

This solution addresses the following use cases:

- Automated Oracle container database deployment on NetApp C-Series storage controllers.
- Oracle database protection and clone on C-Series with SnapCenter UI tool.

Audience

This solution is intended for the following people:


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

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

Simplified, Automated Oracle Database Deployment on NetApp C-Series with NFS



NetApp

Hardware and software components

Hardware		
NetApp C-Series C400	ONTAP Version 9.13.1P3	Two disk shelves / 24 disks with 278 TiB capacity
VM for DB server	4 vCPUs, 16GiB RAM	Two Linux VM instances for concurrent deployment
VM for SnapCenter	4 vCPUs, 16GiB RAM	One Windows VM instance
Software		
RedHat Linux	RHEL Linux 8.6 (LVM) - x64 Gen2	Deployed RedHat subscription for testing
Windows Server	2022 DataCenter 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 C400 volumes
ora_02	NTAP2(NTAP2_PDB1,NTAP2_PD B2,NTAP2_PDB3)	/u01, /u02, /u03 NFS mounts on C400 volumes

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 a DB VM can drive significantly more I/O than the native NFS client. Automated Oracle deployment configures dNFS on NFSv3 by default.
- **Load balancing on C400 controller pair.** Place Oracle database volumes on C400 controller nodes evenly to balance the workload. DB1 on controller 1, DB2 on controller 2, and so on. Mount the DB volumes to its local lif address.
- **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 information for Oracle database protection and clone after deployment.

Prerequisites for deployment

Deployment requires the following prerequisites.

1. A NetApp C-Series storage controller pair is racked, stacked, and latest version of ONTAP operating system is installed and configured. Refer to this setup guide as necessary: [Detailed guide - AFF C400](#)
2. Provision two Linux VMs as Oracle DB servers. See the architecture diagram in the previous section for details about the environment setup.
3. 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.

Enable ssh public/private key authentication between Ansible controller and database VMs.

5. From Ansible controller admin user home directory, 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
```

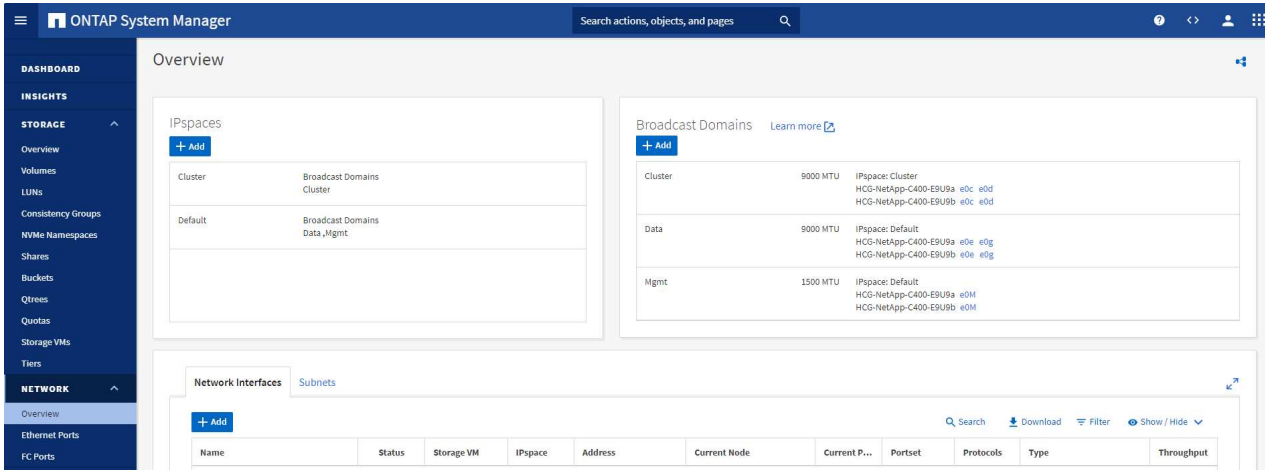
6. Stage following Oracle 19c installation files on DB VM /tmp/archive directory with 777 permission.

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

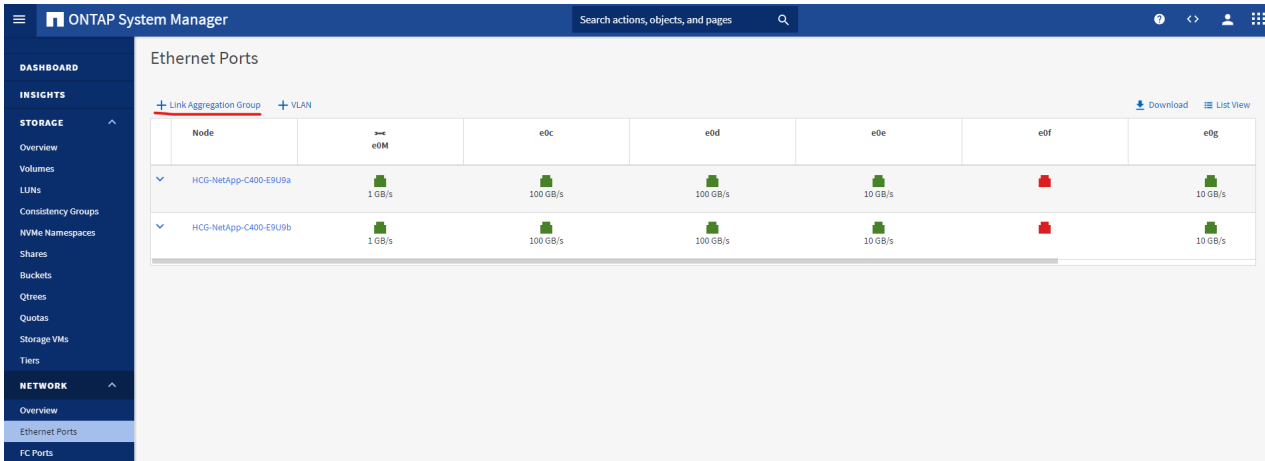
Configure Networking and SVM on C-Series for Oracle

This section of deployment guide demonstrates best practices to set up networking and storage virtual machine (SVM) on C-Series controller for Oracle workload with NFS protocol using ONTAP System Manager UI.

1. Login to ONTAP System Manager to review that after initial ONTAP cluster installation, broadcast domains have been configured with ethernet ports properly assigned to each domain. Generally, there should be a broadcast domain for cluster, a broadcast domain for management, and a broadcast domain for workload such as data.



2. From NETWORK - Ethernet Ports, click Link Aggregate Group to create a LACP link aggregate group port a0a, which provides load balance and failover among the member ports in the aggregate group port. There are 4 data ports - e0e, e0f, e0g, e0h available on C400 controllers.



3. Select the ethernet ports in the group, LACP for mode, and Port for load distribution.

Add Link Aggregation Group



NODE

HCG-NetApp-C400-E9U9a

BROADCAST DOMAIN

Automatically select a broadcast domain (recommended)

PORTS TO INCLUDE

The following ports are down: e0f, e0h.

e0e e0f e0g e0h

MODE

- Single
Only one port is used at a time.
- Multiple
All ports can be used simultaneously.
- LACP
The LACP protocol determines the ports that can be used.

LOAD DISTRIBUTION

- IP based
Network traffic is distributed based on the destination IP address.
- MAC based
Network traffic is distributed based on the next-hop MAC addresses.
- Sequential
Network traffic is distributed by round-robin over the outbound links.
- Port
Network traffic is distributed based on the transport layer (TCP/UDP) ports.

Save

Cancel

4. Validate LACP port a0a created and broadcast domain Data is now operating on LACP port.

The screenshot shows the ONTAP System Manager interface. The left sidebar contains navigation menus for Dashboard, Insights, Storage, Network, FC Ports, Events & Jobs, Protection, Hosts, and Cluster. The main content area is titled "Ethernet Ports" and displays a table of port configurations for two nodes: HCG-NetApp-C400-E9U9b and HCG-NetApp-C400-E9U9a. The table has columns for Node, a0a, e0M, e0c, and e0d. Each port is shown with a green status icon and a speed of 100 GB/s. The table also includes a "Link Aggregation Group" column and a "VLAN" column, both currently empty.

Node	a0a	e0M	e0c	e0d
HCG-NetApp-C400-E9U9b		1 GB/s	100 GB/s	100 GB/s
HCG-NetApp-C400-E9U9a		1 GB/s	100 GB/s	100 GB/s

ONTAP System Manager Search actions, objects, and pages

DASHBOARD

INSIGHTS

STORAGE ▾

NETWORK ▾

Overview

Ethernet Ports

FC Ports

EVENTS & JOBS ▾

PROTECTION ▾

HOSTS ▾

CLUSTER ▾

Overview

IPspaces + Add

Cluster	Broadcast Domains Cluster
Default	Broadcast Domains Data ,Mgmt

Broadcast Domains [Learn more](#) + Add

Cluster	9000 MTU	IPspace: Cluster HCG-NetApp-C400-E9U9a e0c e0d HCG-NetApp-C400-E9U9b e0c e0d
Data	9000 MTU	IPspace: Default HCG-NetApp-C400-E9U9a a0a HCG-NetApp-C400-E9U9b a0a
Mgmt	1500 MTU	IPspace: Default HCG-NetApp-C400-E9U9a e0M

- From Ethernet Ports, click VLAN to add a VLAN on each controller node for Oracle workload on NFS protocol.

Add VLAN



NODE

HCG-NetApp-C400-E9U9a



BROADCAST DOMAIN

Automatically select a broadcast domain (recommended)



PORT

a0a



VLAN ID

3277

Cancel

Save

ONTAP System Manager Search actions, objects, and pages

Ethernet Ports

+ Link Aggregation Group + VLAN Download List View

Node	a0a	a0a-3277	e0M	e0c
HCG-NetApp-C400-E9U9b			1 GB/s	100 GB/s
HCG-NetApp-C400-E9U9a			1 GB/s	100 GB/s

ONTAP System Manager Search actions, objects, and pages

Overview

IPspaces

+ Add

Cluster	Broadcast Domains
Cluster	Cluster
Default	Data, Mgmt

Broadcast Domains

Learn more

+ Add

Cluster	9000 MTU	IPspace: Cluster
HCG-NetApp-C400-E9U9a	e0c e0d	HCG-NetApp-C400-E9U9b e0c e0d
Data	9000 MTU	IPspace: Default
HCG-NetApp-C400-E9U9a	a0a a0a-3277	HCG-NetApp-C400-E9U9b a0a a0a-3277
Mgmt	1500 MTU	IPspace: Default
HCG-NetApp-C400-E9U9a	e0M	

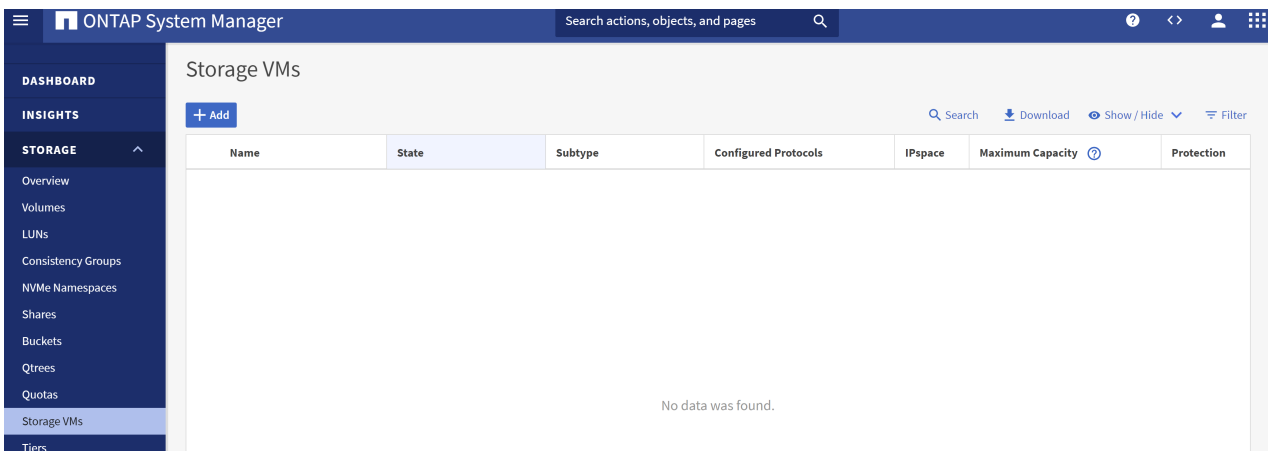
6. Login to C-Series controllers from cluster management IP via ssh to validate that network failover groups are configured correctly. ONTAP create and manage failover groups automatically.

```

HCG-NetApp-C400-E9U9::> net int failover-groups show
(network interface failover-groups show)
                                Failover
Vserver          Group          Targets
-----
Cluster
                Cluster
                                HCG-NetApp-C400-E9U9a:e0c,
                                HCG-NetApp-C400-E9U9a:e0d,
                                HCG-NetApp-C400-E9U9b:e0c,
                                HCG-NetApp-C400-E9U9b:e0d
HCG-NetApp-C400-E9U9
                Data
                                HCG-NetApp-C400-E9U9a:a0a,
                                HCG-NetApp-C400-E9U9a:a0a-3277,
                                HCG-NetApp-C400-E9U9b:a0a,
                                HCG-NetApp-C400-E9U9b:a0a-3277
                Mgmt
                                HCG-NetApp-C400-E9U9a:e0M,
                                HCG-NetApp-C400-E9U9b:e0M
3 entries were displayed.

```

7. From STORAGE - Storage VMs, click +Add to create a SVM for Oracle.



8. Name your Oracle SVM, check Enable NFS and Allow NFS client access.

Add Storage VM



STORAGE VM NAME

oracle

Access Protocol

SMB/CIFS, NFS, S3 [iSCSI](#) [FC](#) [NVMe](#)

Enable SMB/CIFS

Enable NFS

Allow NFS client access

Add at least one rule to allow NFS clients to access volumes in this storage VM. [?](#)

EXPORT POLICY

Default

RULES

No data

[+](#) Add

Enable S3

DEFAULT LANGUAGE [?](#)

c.utf_8

9. Add NFS export policy Default rules.

New Rule



CLIENT SPECIFICATION

172.21.21.0/255.255.255.0


ACCESS PROTOCOLS

SMB/CIFS

FlexCache

NFS NFSv3 NFSv4

ACCESS DETAILS

Type	Read-only Access	Read/Write Access	Superuser Access
All	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
All (As anonymous user) 	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
UNIX	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Kerberos 5	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Kerberos 5i	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Kerberos 5p	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
NTLM	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Cancel

Save

10. In NETWORK INTERFACE, fill in IP address on each node for NFS lif addresses.

NETWORK INTERFACE

Use multiple network interfaces when client traffic is high.

HCG-NetApp-C400-E9U9a

IP ADDRESS

172.21.21.100

SUBNET MASK

255.255.255.0

GATEWAY

[Add optional gateway](#)

BROADCAST DOMAIN AND PORT

Data

Use the same subnet mask and gateway for all of the following interfaces

HCG-NetApp-C400-E9U9b

IP ADDRESS

172.21.21.101

SUBNET MASK

255.255.255.0

GATEWAY

[Add optional gateway](#)

BROADCAST DOMAIN AND PORT

Data

Storage VM Administration

- Enable maximum capacity limit
The maximum capacity that all volumes in this storage VM can allocate. [Learn More](#)
- Manage administrator account

Save

Cancel

11. Validate SVM for Oracle is up/running and NFS lifs status is active.

The screenshot shows the ONTAP System Manager interface. The left sidebar is expanded to 'STORAGE' > 'Storage VMs'. The main content area displays a table of Storage VMs. One VM named 'oracle' is listed with the following details:

Name	State	Subtype	Configured Protocols	IPspace	Maximum Capacity	Protection
oracle	running	default	NFS	Default	The maximum capacity is disabled	

Network Interfaces [Subnets](#)

[+ Add](#) [Search](#) [Download](#) [Filter](#) [Show / Hide](#)

Name	Status	Storage VM	IPspace	Address	Current Node	Current P...	Portset	Protocols	T
HCG-NetApp-C400-E9U9a_clu s1	✔		Cluster	169.254.47.43	HCG-NetApp-C400-E9U9a	e0c			C
HCG-NetApp-C400-E9U9b_clu s1	✔		Cluster	169.254.152.124	HCG-NetApp-C400-E9U9b	e0c			C
HCG-NetApp-C400-E9U9b_clu s2	✔		Cluster	169.254.107.230	HCG-NetApp-C400-E9U9b	e0d			C
HCG-NetApp-C400-E9U9b_mg mt1	✔		Default	10.61.180.109	HCG-NetApp-C400-E9U9b	e0M			C
lif_oracle_145	✔	oracle	Default	172.21.21.100	HCG-NetApp-C400-E9U9a	a0a-3277		NFS	D
lif_oracle_37	✔	oracle	Default	172.21.21.101	HCG-NetApp-C400-E9U9b	a0a-3277		NFS	D

Showing 1 - 9 of 9 Network Interfaces [←](#) **1** [→](#)

12. From STORAGE-Volumes tab to add NFS volumes for Oracle database.

ONTAP System Manager [Search actions, objects, and pages](#)

Volumes [+ Add](#) [More](#) [Search](#) [Download](#) [Show / Hide](#) [Filter](#)

Name	Storage VM	Status	Capacity	IOPS	Latency (ms)	Throughput (MB/s)	Protection
oracle_root	oracle	✔ Online	<div style="display: flex; align-items: center;"> <div style="width: 292px; height: 10px; background-color: #ccc; margin-right: 5px;"></div> <div style="width: 973px; height: 10px; background-color: #ccc; margin-right: 5px;"></div> <div style="width: 100px; height: 10px; background-color: #0070c0; margin-right: 5px;"></div> 1 GiB </div> <div style="display: flex; align-items: center; margin-top: 2px;"> 292 KiB used 973 MiB available </div>	0	0	0	✔ ☁

13. Name your volume, assign capacity, and performance level.

Add Volume



NAME



Add as a cache for a remote volume (FlexCache)

Simplifies file distribution, reduces WAN latency, and lowers WAN bandwidth costs.

Storage and Optimization

CAPACITY

GiB



PERFORMANCE SERVICE LEVEL



Not sure? [Get help selecting type](#)

OPTIMIZATION OPTIONS



Distribute volume data across the cluster (FlexGroup)



14. In `Access Permission`, choose the default policy created from previous step. Uncheck `Enable Snapshot Copies` as we prefer to use SnapCenter to create application consistent snapshots.

Access Permissions

Export via NFS

GRANT ACCESS TO HOST

default

Create a new export policy, or select an existing export policy.

Rule Index	Clients	Access Protocols	Read-Only Rule	Read/W
1	172.21.21.0/255.255.255.0	NFSv3, NFSv4, NFS	Sys	Sys

SnapLock


[SnapLock Considerations](#)

Enable SnapLock

With SnapLock, files can be stored and committed to a non-erasable, non-rewritable state either forever or for a designated retention period.

Protection

Enable Snapshot Copies (Local)

Enable Snapshot locking 

Enables the ability to lock Snapshot copies that were created either manually or by Snapshot policies. The Snapshot copies are locked only when a retention period is specified.

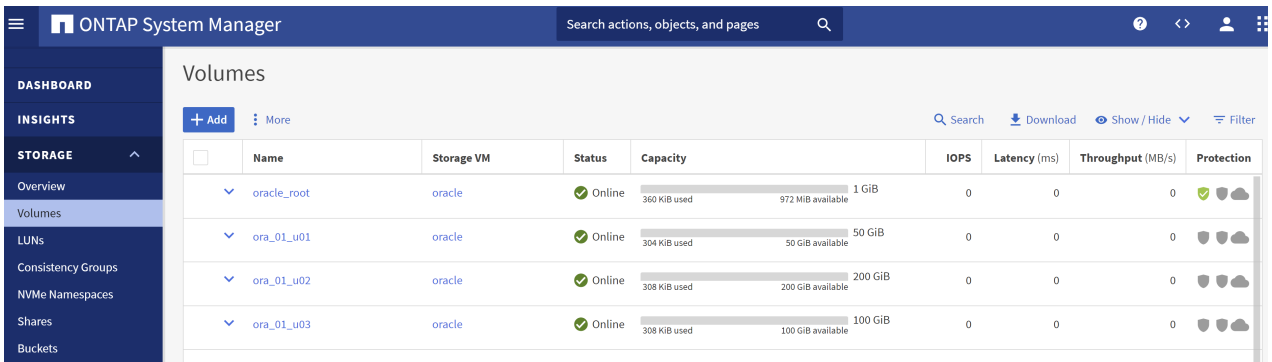
Enable SnapMirror (Local or Remote)

Save

Cancel

 Save to Ansible Playbook

15. Create three DB volumes for each DB server: `server_name_u01` - binary, `server_name_u02` - data, `server_name_u03` - logs.



The screenshot shows the ONTAP System Manager interface. The left sidebar contains navigation options: DASHBOARD, INSIGHTS, STORAGE (selected), Overview, Volumes, LUNs, Consistency Groups, NVMe Namespaces, Shares, and Buckets. The main content area is titled 'Volumes' and displays a table with columns: Name, Storage VM, Status, Capacity, IOPS, Latency (ms), Throughput (MB/s), and Protection. There are four rows of volumes listed:

Name	Storage VM	Status	Capacity	IOPS	Latency (ms)	Throughput (MB/s)	Protection
oracle_root	oracle	Online	360 KiB used / 972 MiB available / 1 GiB	0	0	0	Protected
ora_01_u01	oracle	Online	304 KiB used / 50 GiB available / 50 GiB	0	0	0	Protected
ora_01_u02	oracle	Online	308 KiB used / 200 GiB available / 200 GiB	0	0	0	Protected
ora_01_u03	oracle	Online	308 KiB used / 100 GiB available / 100 GiB	0	0	0	Protected



The DB volume naming convention should strictly follow format as stated above to ensure automation to work correctly.

This completes the C-series controller configuration for Oracle.

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 admin -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 admin -e
@vars/vars.yml
```

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

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

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

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

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

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

4. Undo the environment

```
ansible-playbook -i hosts 5-destroy.yml -u admin -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 DB VM ora_01 or ora_02.

1. Validate NFS mounts

```
[admin@ora_01 ~]$ cat /etc/fstab

#
# /etc/fstab
# Created by anaconda on Wed Oct 18 19:43:31 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/rhel-root / xfs defaults
0 0
UUID=aff942c4-b224-4b62-807d-6a5c22f7b623 /boot
xfs defaults 0 0
/dev/mapper/rhel-swap none swap defaults
0 0
/root/swapfile swap swap defaults 0 0
172.21.21.100:/ora_01_u01 /u01 nfs
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsiz=65536 0 0
172.21.21.100:/ora_01_u02 /u02 nfs
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsiz=65536 0 0
172.21.21.100:/ora_01_u03 /u03 nfs
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsiz=65536 0 0

[admin@ora_01 tmp]$ 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      18M  7.8G  1% /run
tmpfs                     7.8G         0  7.8G  0% /sys/fs/cgroup
/dev/mapper/rhel-root     44G        28G  17G  62% /
/dev/sda1                 1014M     258M  757M  26% /boot
tmpfs                     1.6G        12K  1.6G  1% /run/user/42
tmpfs                     1.6G         4.0K  1.6G  1% /run/user/1000
```

```

172.21.21.100:/ora_01_u01 50G 8.7G 42G 18% /u01
172.21.21.100:/ora_01_u02 200G 384K 200G 1% /u02
172.21.21.100:/ora_01_u03 100G 320K 100G 1% /u03

[admin@ora_02 ~]$ 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      18M  7.8G  1% /run
tmpfs                     7.8G         0  7.8G  0% /sys/fs/cgroup
/dev/mapper/rhel-root     44G        28G  17G  63% /
/dev/sda1                 1014M    258M  757M  26% /boot
tmpfs                     1.6G       12K  1.6G  1% /run/user/42
tmpfs                     1.6G       4.0K  1.6G  1% /run/user/1000
172.21.21.101:/ora_02_u01 50G 7.8G 43G 16% /u01
172.21.21.101:/ora_02_u02 200G 320K 200G 1% /u02
172.21.21.101:/ora_02_u03 100G 320K 100G 1% /u03

```

2. Validate Oracle listener

```

[admin@ora_02 ~]$ sudo su
[root@ora_02 admin]# su - oracle
[oracle@ora_02 ~]$ lsnrctl status listener.ntap2

LSNRCTL for Linux: Version 19.0.0.0.0 - Production on 29-MAY-2024
12:13:30

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

Connecting to
 (DESCRIPTION=(ADDRESS=(PROTOCOL=TCP) (HOST=ora_02.cie.netapp.com) (PORT=1521)))
STATUS of the LISTENER
-----
Alias                     LISTENER.NTAP2
Version                   TNSLSNR for Linux: Version 19.0.0.0.0 -
Production
Start Date                23-MAY-2024 16:13:03
Uptime                    5 days 20 hr. 0 min. 26 sec
Trace Level               off
Security                  ON: Local OS Authentication
SNMP                      OFF
Listener Parameter File
/u01/app/oracle/product/19.0.0/NTAP2/network/admin/listener.ora
Listener Log File

```

```

/u01/app/oracle/diag/tnslsnr/ora_02/listener.ntap2/alert/log.xml
Listening Endpoints Summary...

(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=ora_02.cie.netapp.com)(PORT=1521)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=EXTPROC1521)))

(DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)(HOST=ora_02.cie.netapp.com)(PORT=5500))(Security=(my_wallet_directory=/u01/app/oracle/product/19.0.0/NTAP2/admin/NTAP2/xdb_wallet))(Presentation=HTTP)(Session=RAW))
Services Summary...
Service "192551f1d7e65fc3e06308b43d0a63ae.solutions.netapp.com" has
1 instance(s).
  Instance "NTAP2", status READY, has 1 handler(s) for this
service...
Service "1925529a43396002e06308b43d0a2d5a.solutions.netapp.com" has
1 instance(s).
  Instance "NTAP2", status READY, has 1 handler(s) for this
service...
Service "1925530776b76049e06308b43d0a49c3.solutions.netapp.com" has
1 instance(s).
  Instance "NTAP2", status READY, has 1 handler(s) for this
service...
Service "NTAP2.solutions.netapp.com" has 1 instance(s).
  Instance "NTAP2", status READY, has 1 handler(s) for this
service...
Service "NTAP2XDB.solutions.netapp.com" has 1 instance(s).
  Instance "NTAP2", status READY, has 1 handler(s) for this
service...
Service "ntap2_pdb1.solutions.netapp.com" has 1 instance(s).
  Instance "NTAP2", status READY, has 1 handler(s) for this
service...
Service "ntap2_pdb2.solutions.netapp.com" has 1 instance(s).
  Instance "NTAP2", status READY, has 1 handler(s) for this
service...
Service "ntap2_pdb3.solutions.netapp.com" has 1 instance(s).
  Instance "NTAP2", status READY, has 1 handler(s) for this
service...
The command completed successfully
[oracle@ora_02 ~]$

```

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
      5 NTAP1_PDB3                    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
```

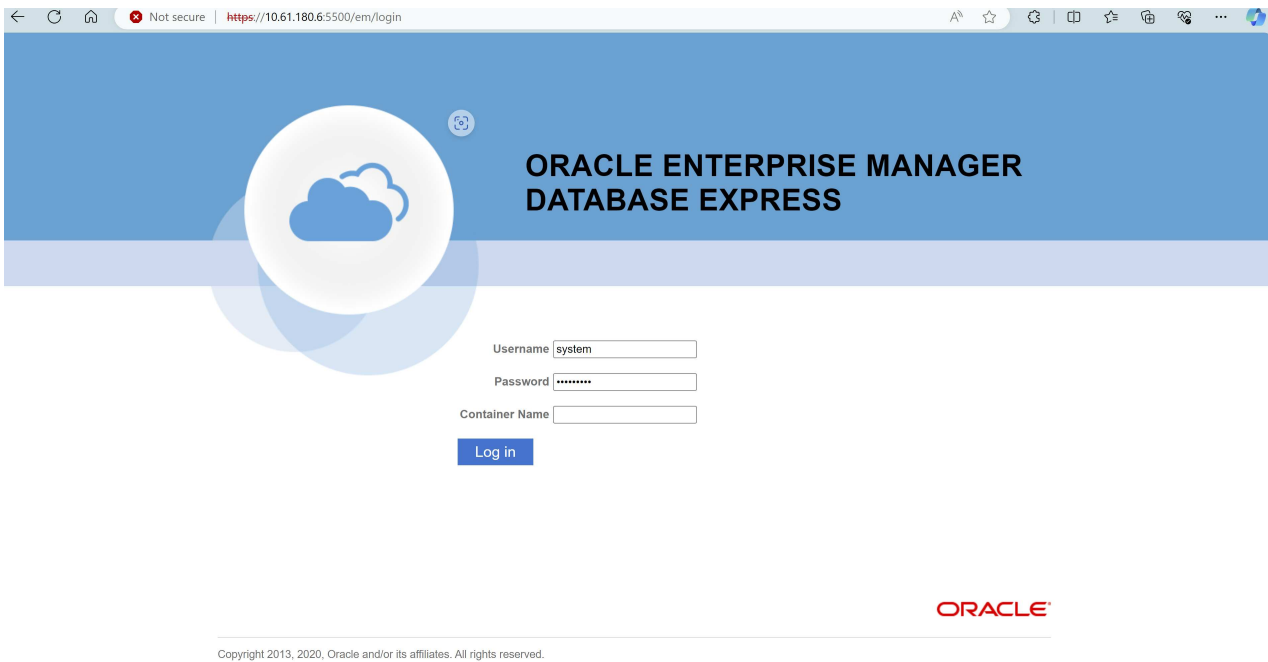
```
SQL> select svrname, dirname from v$dnfs_servers;
```

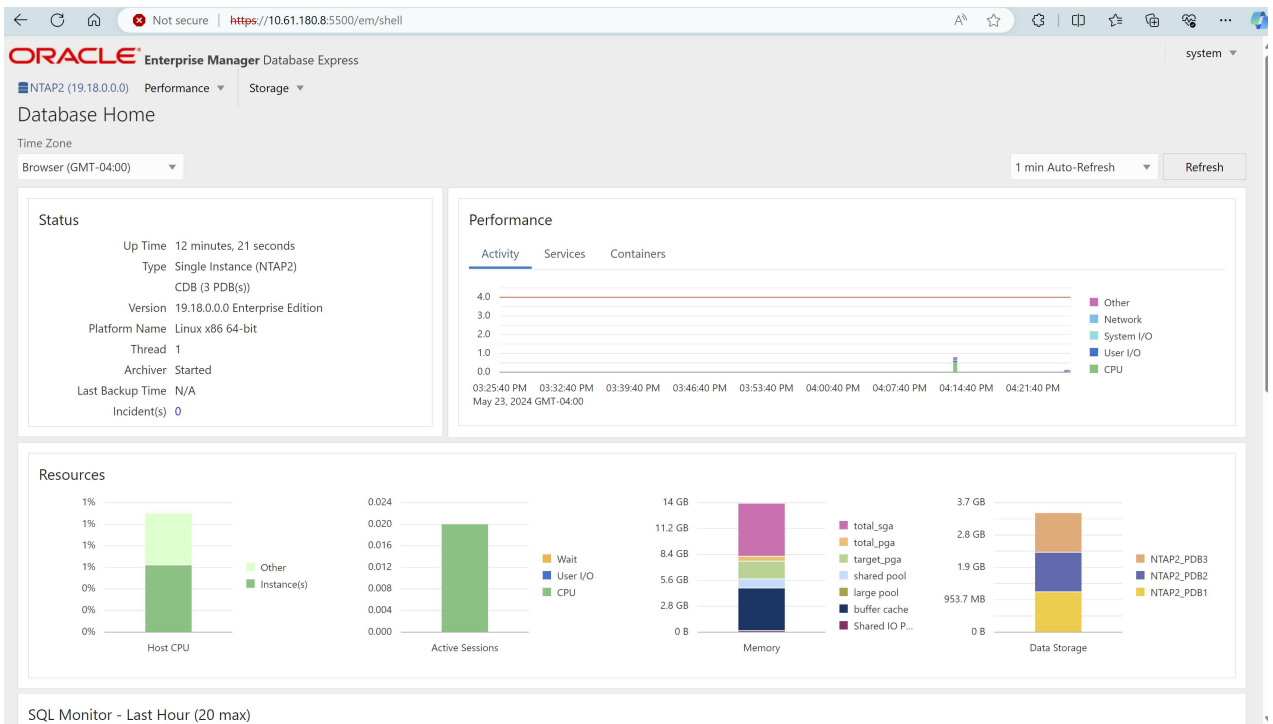
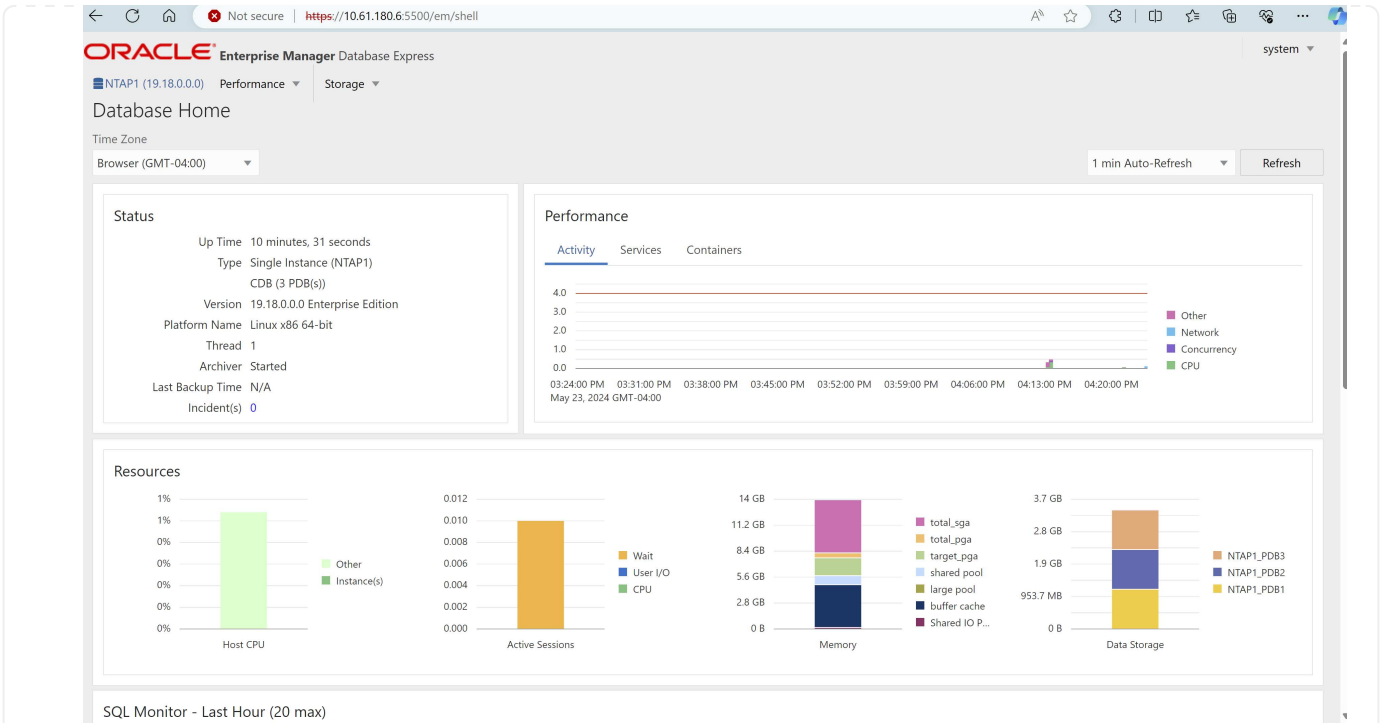
SVRNAME

DIRNAME

```
-----  
-----  
172.21.21.100  
/ora_01_u02  
  
172.21.21.100  
/ora_01_u03  
  
172.21.21.100  
/ora_01_u01
```

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





Oracle backup, restore, and clone with SnapCenter

NetApp recommends SnapCenter UI tool to manage Oracle database deployed in C-Series. Refer to TR-4979 [Simplified, Self-managed Oracle in VMware Cloud on AWS with guest-mounted FSx ONTAP](#) section Oracle backup, restore, and clone with SnapCenter for details on setting up SnapCenter and executing the database backup, restore, and clone workflows.

Where to find additional information

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

- NetApp AFF C-Series

<https://www.netapp.com/pdf.html?item=/media/81583-da-4240-aff-c-series.pdf>

- NetApp Enterprise Database Solutions

<https://docs.netapp.com/us-en/netapp-solutions/databases/index.html>

- Deploying Oracle Direct NFS

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

TR-4983: Simplified, Automated Oracle Deployment on NetApp ASA with iSCSI

Allen Cao, Niyaz Mohamed, NetApp

This solution provides overview and details for automated Oracle deployment and protection in NetApp ASA array as primary database storage with iSCSI protocol and Oracle database configured in standalone ReStart using asm as volume manager.

Purpose

NetApp ASA systems deliver modern solutions to your SAN infrastructure. They simplify at scale and enable you to accelerate your business-critical applications such as databases, make sure that your data is always available (99.9999% uptime), and reduce TCO and carbon footprint. The NetApp ASA systems include A-Series models designed for the most performance-demanding applications and C-Series models optimized for cost-effective, large-capacity deployments. Together, the ASA A-Series and C-Series systems deliver exceptional performance to improve customer experience and reduce time to results, keep business-critical data available, protected, and secure, and provide more effective capacity for any workload, backed by the industry's most effective guarantee.

This documentation demonstrates the simplified deployment of Oracle databases in a SAN environment built with ASA systems using Ansible automation. The Oracle database is deployed in a standalone ReStart configuration with iSCSI protocol for data access and Oracle ASM for database disks management on the ASA storage array. It also provides information on Oracle database backup, restore, and clone using the NetApp SnapCenter UI tool for storage-efficient database operation in NetApp ASA systems.

This solution addresses the following use cases:

- Automated Oracle database deployment in NetApp ASA systems as primary database storage
- Oracle database backup and restore in NetApp ASA systems using NetApp SnapCenter tool
- Oracle database clone for dev/test or other use cases in NetApp ASA systems using NetApp SnapCenter tool

Audience

This solution is intended for the following people:

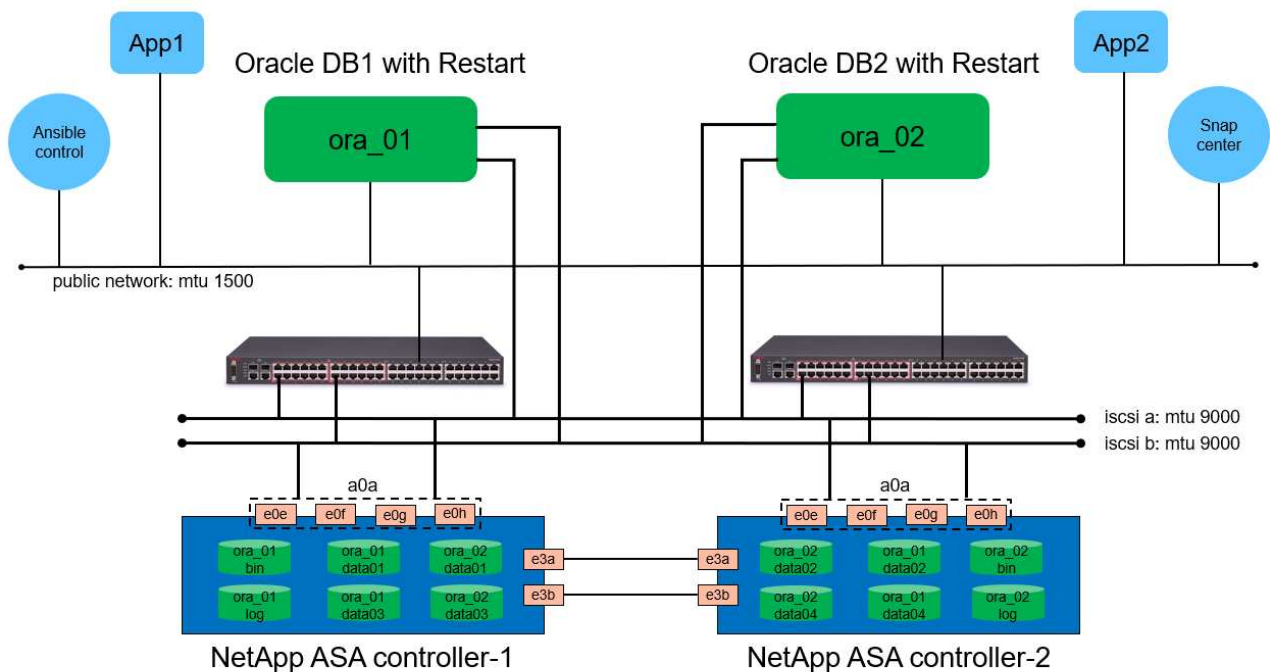
- A DBA who would like to deploy Oracle in NetApp ASA systems.
- A database solution architect who would like to test Oracle workloads in NetApp ASA systems.
- A storage administrator who would like to deploy and manage an Oracle database on NetApp ASA systems.
- An application owner who would like to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

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

Simplified, Automated Oracle Database Deployment on NetApp ASA with iSCSI



NetApp

Hardware and software components

Hardware		
NetApp ASA A400	Version 9.13.1P1	2 NS224 shelves, 48 NVMe AFF drives with total 69.3 TiB capacity
UCSB-B200-M4	Intel® Xeon® CPU E5-2690 v4 @ 2.60GHz	4-node VMware ESXi cluster
Software		
RedHat Linux	RHEL-8.6, 4.18.0-372.9.1.el8.x86_64 kernel	Deployed RedHat subscription for testing
Windows Server	2022 Standard, 10.0.20348 Build 20348	Hosting SnapCenter server

Oracle Grid Infrastructure	Version 19.18	Applied RU patch p34762026_190000_Linux-x86-64.zip
Oracle Database	Version 19.18	Applied RU patch p34765931_190000_Linux-x86-64.zip
Oracle OPatch	Version 12.2.0.1.36	Latest patch p6880880_190000_Linux-x86-64.zip
SnapCenter Server	Version 4.9P1	Workgroup deployment
VMware vSphere Hypervisor	version 6.5.0.20000	VMware Tools, Version: 11365 - Linux, 12352 - Windows
Open JDK	Version java-1.8.0-openjdk.x86_64	SnapCenter plugin requirement on DB VMs

Oracle database configuration in the lab environment

Server	Database	DB Storage
ora_01	NTAP1(NTAP1_PDB1,NTAP1_PDB2,NTAP1_PDB3)	iSCSI luns on ASA A400
ora_02	NTAP2(NTAP2_PDB1,NTAP2_PDB2,NTAP2_PDB3)	iSCSI luns on ASA A400

Key factors for deployment consideration

- **Oracle database storage layout.** In this automated Oracle deployment, we provision four database volumes to host Oracle binary, data, and logs by default. We then create two ASM disk groups from data and logs luns. Within the +DATA asm disk group, we provision two data luns in a volume on each ASA A400 cluster node. Within the +LOGS asm disk group, we create two luns in a log volume on a single ASA A400 node. Multiple luns laid out within an ONTAP volume provides better performance in general.
- **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. In the event of multi-DB server deployments, the playbook builds with an algorithm to place database luns on dual controllers of ASA A400 optimally. The binary and logs luns of odd number DB server in server hosts index place on controller 1. The binary and logs luns of even number DB server in the server hosts index place on controller 2. The DB data luns evenly distributed to two controllers. Oracle ASM combines the data luns on two controllers into a single ASM disk group to fully utilize the processing power of both controllers.
- **iSCSI configuration.** The database VMs connect to ASA storage with the iSCSI protocol for storage access. You should configure dual paths on each controller node for redundancy and set up iSCSI multi-path on the DB server for multi-path storage access. Enable jumbo frame on storage network to maximize performance and throughput.
- **Oracle ASM redundancy level to use for each Oracle ASM disk group that you create.** Because the ASA A400 configures storage in RAID DP for data protection at the cluster disk level, you should use `External Redundancy`, which means that the option does not allow Oracle ASM to mirror the contents of the disk group.

- **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 protection in NetApp ASA A400 with directly mounted database luns via iSCSI to DB VM in a single node Restart configuration with Oracle ASM as database volume manager.

Prerequisites for deployment

Deployment requires the following prerequisites.

1. It is assumed that the NetApp ASA storage array has been installed and configured. This includes iSCSI broadcast domain, LACP interface groups a0a on both controller nodes, iSCSI VLAN ports (a0a-<iscsi-a-vlan-id>, a0a-<iscsi-b-vlan-id>) on both controller nodes. The following link provides detailed step-by-step instructions if help is needed. [Detailed guide - ASA A400](#)
2. Provision a Linux VM as an 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.
3. Clone a copy of the NetApp Oracle deployment automation toolkit for iSCSI.

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

4. 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](#)
5. Build two RHEL Oracle DB servers either bare metal or virtualized VM. Create an admin user on DB servers with sudo without password privilege and enable SSH private/public key authentication between Ansible host and Oracle DB server hosts. Stage following Oracle 19c installation files on DB servers /tmp/archive directory.

```
installer_archives:  
- "LINUX.X64_193000_grid_home.zip"  
- "p34762026_190000_Linux-x86-64.zip"  
- "LINUX.X64_193000_db_home.zip"  
- "p34765931_190000_Linux-x86-64.zip"  
- "p6880880_190000_Linux-x86-64.zip"
```



Ensure that you have allocated at least 50G in Oracle VM root volume to have sufficient space to stage Oracle installation files.

6. Watch the following video:

[Simplified and automated Oracle deployment on NetApp ASA with iSCSI](#)

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 local 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 the user-defined variable files are configured.

Parameter files configuration

1. Ansible target hosts file configuration:

```
# Enter NetApp ASA controller management IP address
[ontap]
172.16.9.32

# 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 global user
configurable variables          #####
#####          Consolidate all variables from ONTAP, linux
and oracle          #####
#####
#####
#####          ONTAP env specific config variables
#####
#####
#####
#####
#####
#####          ONTAP platform: on-prem, aws-fsx.
ontap_platform: on-prem

# Enter ONTAP cluster management user credentials
username: "xxxxxxxx"
password: "xxxxxxxx"

##### on-prem platform specific user defined variables #####

# Enter Oracle SVM iSCSI lif addresses. Each controller configures
```



```

with dual paths iscsi_a, iscsi_b for redundancy
ora_iscsi_lif_mgmt:
  - {name: '{{ svm_name }}_mgmt', address: 172.21.253.220, netmask:
255.255.255.0, vlan_name: ora_mgmt, vlan_id: 3509}

ora_iscsi_lifs_node1:
  - {name: '{{ svm_name }}_lif_1a', address: 172.21.234.221,
netmask: 255.255.255.0, vlan_name: ora_iscsi_a, vlan_id: 3490}
  - {name: '{{ svm_name }}_lif_1b', address: 172.21.235.221,
netmask: 255.255.255.0, vlan_name: ora_iscsi_b, vlan_id: 3491}
ora_iscsi_lifs_node2:
  - {name: '{{ svm_name }}_lif_2a', address: 172.21.234.223,
netmask: 255.255.255.0, vlan_name: ora_iscsi_a, vlan_id: 3490}
  - {name: '{{ svm_name }}_lif_2b', address: 172.21.235.223,
netmask: 255.255.255.0, vlan_name: ora_iscsi_b, vlan_id: 3491}

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

# Enter RHEL subscription to enable repo
redhat_sub_username: xxxxxxxx
redhat_sub_password: "xxxxxxx"

#####
#####
###
Oracle DB env specific config variables
###
#####
#####

# Enter Database domain name
db_domain: solutions.netapp.com

# Enter initial password for all required Oracle passwords. Change
them after installation.
initial_pwd_all: xxxxxxxx

```

3. Local DB server host_vars/host_name.yml configuration

```
# 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
```

Playbook execution

There are a total of six 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.
3-ontap_config.yml - configure ONTAP svm/volumes/luns for Oracle
database and grant DB server access to luns.
4-oracle_config.yml - install and configure Oracle on DB servers for
grid infrastructure 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 admin -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 admin -e
@vars/vars.yml
```

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

```
ansible-playbook -i hosts 3-ontap_config.yml -u admin -e
@vars/vars.yml
```

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

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

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

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

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

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

4. Undo the environment

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

Post execution validation

After the playbook run, login to the Oracle DB server as oracle user to validate that Oracle grid infrastructure and database are created successfully. Following is an example of Oracle database validation on host ora_01.

1. Validate the grid infrastructure and resources created.

```
[oracle@ora_01 ~]$ df -h
Filesystem                Size      Used Avail Use% Mounted on
devtmpfs                  7.7G      40K    7.7G   1% /dev
tmpfs                     7.8G     1.1G    6.7G  15% /dev/shm
tmpfs                     7.8G     312M    7.5G   4% /run
tmpfs                     7.8G        0    7.8G   0% /sys/fs/cgroup
/dev/mapper/rhel-root      44G      38G    6.8G  85% /
/dev/sda1                 1014M    258M    757M  26% /boot
tmpfs                     1.6G      12K    1.6G   1% /run/user/42
tmpfs                     1.6G     4.0K    1.6G   1% /run/user/1000
/dev/mapper/ora_01_biny_01p1 40G     21G    20G  52% /u01
[oracle@ora_01 ~]$ asm
[oracle@ora_01 ~]$ crsctl stat res -t
-----
-----
Name                Target  State          Server          State
details
-----
-----
Local Resources
-----
-----
ora.DATA.dg
                ONLINE  ONLINE        ora_01          STABLE
ora.LISTENER.lsnr
                ONLINE  INTERMEDIATE  ora_01          Not All
Endpoints Re
gistered, STABLE
ora.LOGS.dg
                ONLINE  ONLINE        ora_01          STABLE
ora.asm
                ONLINE  ONLINE        ora_01
Started, STABLE
ora.ons
                OFFLINE OFFLINE        ora_01          STABLE
-----
-----
Cluster Resources
-----
```

```

-----
ora.cssd
  1          ONLINE  ONLINE      ora_01      STABLE
ora.diskmon
  1          OFFLINE OFFLINE
ora.driver.afd
  1          ONLINE  ONLINE      ora_01      STABLE
ora.evmd
  1          ONLINE  ONLINE      ora_01      STABLE
ora.ntap1.db
  1          ONLINE  ONLINE      ora_01
Open,HOME=/u01/app/o

racle/product/19.0.0

/NTAP1, STABLE
-----
-----
[oracle@ora_01 ~]$

```



Ignore the Not All Endpoints Registered in State details. This results from a conflict of manual and dynamic database registration with the listener and can be safely ignored.

2. Validate ASM filter driver is working as expected.


```

[oracle@ora_01 ~]$ asmcmd
ASMCMD> lsdg
State      Type      Rebal  Sector  Logical_Sector  Block      AU
Total_MB  Free_MB  Req_mir_free_MB  Usable_file_MB  Offline_disks
Voting_files  Name
MOUNTED  EXTERN  N      512     512    4096  4194304
327680   318644          0      318644          0
N  DATA/
MOUNTED  EXTERN  N      512     512    4096  4194304
81920   78880          0      78880          0
N  LOGS/
ASMCMD> lsdisk
Path
AFD:ORA_01_DAT1_01
AFD:ORA_01_DAT1_03
AFD:ORA_01_DAT1_05
AFD:ORA_01_DAT1_07
AFD:ORA_01_DAT2_02
AFD:ORA_01_DAT2_04
AFD:ORA_01_DAT2_06
AFD:ORA_01_DAT2_08
AFD:ORA_01_LOGS_01
AFD:ORA_01_LOGS_02
ASMCMD> afd_state
ASMCMD-9526: The AFD state is 'LOADED' and filtering is 'ENABLED' on
host 'ora_01'
ASMCMD>

```

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

← → ↻ Not secure | https://10.61.180.21:5500/em/login



ORACLE ENTERPRISE MANAGER DATABASE EXPRESS

Username

Password

Container Name

Log in

ORACLE

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

← → ↻ Not secure | https://10.61.180.21:5500/em/shell

ORACLE Enterprise Manager Database Express system ▾

NTAP1 (19.18.0.0.0) Performance ▾ Storage ▾

Database Home

Time Zone: 1 min Auto-Refresh ▾ Refresh

Status

Up Time 1 hours, 7 minutes, 23 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

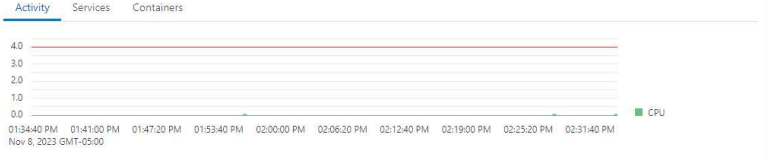
Archiver Stopped

Last Backup Time N/A

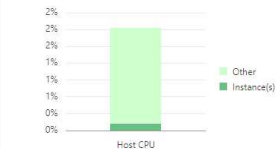
Incident(s) ✖ 4

Performance


Activity Services Containers



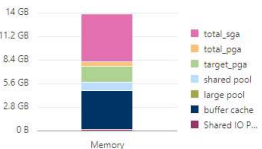
Resources



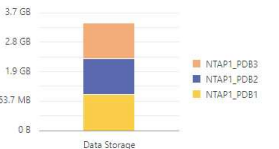
Host CPU: 2%



Active Sessions: 0



Memory: 14 GB



Data Storage: 3.7 GB

SQL Monitor - Last Hour (20 max)

Top 20 by Filter by Status, SQL ID or User Name

Enable additional port from sqlplus for login to individual container database or PDBs.

```
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> select dbms_xdb_config.gethttpsport() from dual;
```

```
DBMS_XDB_CONFIG.GETHTTPS()
-----
                                0
```


```
SQL> exec DBMS_XDB_CONFIG.SETHTTPS(5501);
```

PL/SQL procedure successfully completed.

```
SQL> select dbms_xdb_config.gethttpsport() from dual;
```

```
DBMS_XDB_CONFIG.GETHTTPS()
-----
                                5501
```

login to NTAP1_PDB1 from port 5501.



Oracle backup, restore, and clone with SnapCenter

Refer to TR-4979 [Simplified, self-managed Oracle in VMware Cloud on AWS with guest-mounted FSx ONTAP](#) section Oracle backup, restore, and clone with SnapCenter for details on setting up SnapCenter and executing the database backup, restore, and clone workflows.

Where to find additional information

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

- NETAPP ASA: ALL-FLASH SAN ARRAY

<https://www.netapp.com/data-storage/all-flash-san-storage-array/>

- Installing Oracle Grid Infrastructure for a Standalone Server with a New Database Installation

<https://docs.oracle.com/en/database/oracle/oracle-database/19/ladbi/installing-oracle-grid-infrastructure-for-a-standalone-server-with-a-new-database-installation.html#GUID-0B1CEE8C-C893-46AA-8A6A-7B5FAAEC72B3>

- Installing and Configuring Oracle Database Using Response Files

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

- Use Red Hat Enterprise Linux 8.2 with ONTAP

https://docs.netapp.com/us-en/ontap-sanhost/hu_rhel_82.html#all-san-array-configurations

NVA-1155: Oracle 19c RAC databases on FlexPod Datacenter with Cisco UCS and NetApp AFF A800 over FC - Design and deployment guide

Allen Cao, NetApp

This design and deployment guide for Oracle 19c RAC databases on FlexPod Datacenter with Cisco UCS and NetApp AFF A800 over FC provides details of the solution design as well as step-by-step deployment processes for hosting Oracle RAC databases on most recent FlexPod Datacenter infrastructure with the Oracle Linux 8.2 operating system and a Red Hat compatible kernel.

[NVA-1155: Oracle 19c RAC databases on FlexPod Datacenter with Cisco UCS and NetApp AFF A800 over FC](#)

TR-4250: SAP with Oracle on UNIX and NFS with NetApp Clustered Data ONTAP and SnapManager for SAP 3.4

Nils Bauer, NetApp

TR-4250 addresses the challenges of designing storage solutions to support SAP business suite products using an Oracle database. The primary focus of this document is the common storage infrastructure design, deployment, operation, and management challenges faced by business and IT leaders who use the latest generation of SAP solutions. The recommendations in this document are generic; they are not specific to an SAP application or to the size and scope of the SAP implementation. TR-4250 assumes that the reader has a basic understanding of the technology and operation of NetApp and SAP products. TR-4250 was developed based on the interaction of technical staff from NetApp, SAP, Oracle, and our customers.

[TR-4250: SAP with Oracle on UNIX and NFS with NetApp Clustered Data ONTAP and SnapManager for SAP 3.4](#)

Deploying Oracle Database

Solution Overview

This page describes the Automated method for deploying Oracle19c on NetApp ONTAP storage.

Automated Deployment of Oracle19c for ONTAP on NFS

Organizations are automating their environments to gain efficiencies, accelerate deployments, and reduce manual effort. Configuration management tools like Ansible are being used to streamline enterprise database operations. In this solution, we demonstrate how you can use Ansible to automate the provisioning and configuration of Oracle 19c with NetApp ONTAP. By enabling storage administrators, systems administrators, and DBAs to consistently and rapidly deploy new storage, configure database servers, and install Oracle 19c software, you achieve the following benefits:

- Eliminate design complexities and human errors, and implement a repeatable consistent deployment and best practices
- Decrease time for provisioning of storage, configuration of DB hosts, and Oracle installation

- Increase database administrators, systems and storage administrators productivity
- Enable scaling of storage and databases with ease

NetApp provides customers with validated Ansible modules and roles to accelerate deployment, configuration, and lifecycle management of your Oracle database environment. This solution provides instruction and Ansible playbook code, to help you:

- Create and configure ONTAP NFS storage for Oracle Database
- Install Oracle 19c on RedHat Enterprise Linux 7/8 or Oracle Linux 7/8
- Configure Oracle 19c on ONTAP NFS storage

For more details or to begin, please see the overview videos below.

AWX/Tower Deployments

Part 1: Getting Started, Requirements, Automation Details and Initial AWX/Tower Configuration

[AWX Deployment](#)

Part 2: Variables and Running the Playbook

[AWX Playbook Run](#)

CLI Deployment

Part 1: Getting Started, Requirements, Automation Details and Ansible Control Host Setup

[CLI Deployment](#)

Part 2: Variables and Running the Playbook

[CLI Playbook Run](#)

Getting started

This solution has been designed to be run in an AWX/Tower environment or by CLI on an Ansible control host.

AWX/Tower

For AWX/Tower environments, you are guided through creating an inventory of your ONTAP cluster management and Oracle server (IPs and hostnames), creating credentials, configuring a project that pulls the Ansible code from NetApp Automation Github, and the Job Template that launches the automation.

1. Fill out the variables specific to your environment, and copy and paste them into the Extra Vars fields in your job template.
2. After the extra vars have been added to your job template, you can launch the automation.
3. The job template is run in three phases by specifying tags for `ontap_config`, `linux_config`, and `oracle_config`.

CLI via the Ansible control host

1. To configure the Linux host so that it can be used as an Ansible control host [click here for detailed instructions](#)
2. After the Ansible control host is configured, you can git clone the Ansible Automation repository.
3. Edit the hosts file with the IPs and/or hostnames of your ONTAP cluster management and Oracle server's management IPs.
4. Fill out the variables specific to your environment, and copy and paste them into the `vars.yml` file.
5. Each Oracle host has a variable file identified by its hostname that contains host-specific variables.
6. After all variable files have been completed, you can run the playbook in three phases by specifying tags for `ontap_config`, `linux_config`, and `oracle_config`.

Requirements

Environment	Requirements
Ansible environment	AWX/Tower or Linux host to be the Ansible control host
	Ansible v.2.10 and higher
	Python 3
	Python libraries - netapp-lib - xmltodict - jmespath
ONTAP	ONTAP version 9.3 - 9.7
	Two data aggregates
	NFS vlan and ifgrp created
Oracle server(s)	RHEL 7/8
	Oracle Linux 7/8
	Network interfaces for NFS, public, and optional mgmt
	Oracle installation files on Oracle servers

Automation Details

This automated deployment is designed with a single Ansible playbook that consists of three separate roles. The roles are for ONTAP, Linux, and Oracle configurations. The following table describes which tasks are being automated.

Role	Tasks
ontap_config	Pre-check of the ONTAP environment
	Creation of NFS based SVM for Oracle
	Creation of export policy
	Creation of volumes for Oracle
	Creation of NFS LIFs
linux_config	Create mount points and mount NFS volumes
	Verify NFS mounts
	OS specific configuration
	Create Oracle directories
	Configure hugepages
	Disable SELinux and firewall daemon
	Enable and start chronyd service
	increase file descriptor hard limit
	Create pam.d session file
oracle_config	Oracle software installation
	Create Oracle listener
	Create Oracle databases
	Oracle environment configuration
	Save PDB state
	Enable instance archive mode
	Enable DNFS client
	Enable database auto startup and shutdown between OS reboots

Default parameters

To simplify automation, we have preset many required Oracle deployment parameters with default values. It is generally not necessary to change the default parameters for most deployments. A more advanced user can make changes to the default parameters with caution. The default parameters are located in each role folder under defaults directory.

Deployment instructions

Before starting, download the following Oracle installation and patch files and place them in the `/tmp/archive` directory with read, write, and execute access for all users on each DB server to be deployed. The automation tasks look for the named installation files in that particular directory for Oracle installation and configuration.

```
LINUX.X64_193000_db_home.zip -- 19.3 base installer
p31281355_190000_Linux-x86-64.zip -- 19.8 RU patch
p6880880_190000_Linux-x86-64.zip -- opatch version 12.2.0.1.23
```

License

You should read license information as stated in the Github repository. By accessing, downloading, installing, or using the content in this repository, you agree the terms of the license laid out [here](#).

Note that there are certain restrictions around producing and/or sharing any derivative works with the content in this repository. Please make sure you read the terms of the [License](#) before using the content. If you do not agree to all of the terms, do not access, download, or use the content in this repository.

After you are ready, click [here for detailed AWX/Tower deployment procedures](#) or [here for CLI deployment](#).

Step-by-step deployment procedure

This page describes the Automated method for deploying Oracle19c on NetApp ONTAP storage.

AWX/Tower deployment Oracle 19c Database

1. Create the inventory, group, hosts, and credentials for your environment

This section describes the setup of inventory, groups, hosts, and access credentials in AWX/Ansible Tower that prepare the environment for consuming NetApp automated solutions.

1. Configure the inventory.
 - a. Navigate to Resources → Inventories → Add, and click Add Inventory.
 - b. Provide the name and organization details, and click Save.
 - c. On the Inventories page, click the inventory created.
 - d. If there are any inventory variables, paste them in the variables field.
 - e. Navigate to the Groups sub-menu and click Add.
 - f. Provide the name of the group for ONTAP, paste the group variables (if any) and click Save.
 - g. Repeat the process for another group for Oracle.
 - h. Select the ONTAP group created, go to the Hosts sub-menu and click Add New Host.
 - i. Provide the IP address of the ONTAP cluster management IP, paste the host variables (if any), and click Save.
 - j. This process must be repeated for the Oracle group and Oracle host(s) management IP/hostname.
2. Create credential types. For solutions involving ONTAP, you must configure the credential type to match username and password entries.
 - a. Navigate to Administration → Credential Types, and click Add.
 - b. Provide the name and description.
 - c. Paste the following content in Input Configuration:

```
fields:
  - id: username
    type: string
    label: Username
  - id: password
    type: string
    label: Password
    secret: true
  - id: vsadmin_password
    type: string
    label: vsadmin_password
    secret: true
```

a. Paste the following content into Injector Configuration:

```
extra_vars:
  password: '{{ password }}'
  username: '{{ username }}'
  vsadmin_password: '{{ vsadmin_password }}'
```

1. Configure the credentials.

- a. Navigate to Resources → Credentials, and click Add.
- b. Enter the name and organization details for ONTAP.
- c. Select the custom Credential Type you created for ONTAP.
- d. Under Type Details, enter the username, password, and vsadmin_password.
- e. Click Back to Credential and click Add.
- f. Enter the name and organization details for Oracle.
- g. Select the Machine credential type.
- h. Under Type Details, enter the Username and Password for the Oracle hosts.
- i. Select the correct Privilege Escalation Method, and enter the username and password.

2. Create a project

1. Go to Resources → Projects, and click Add.
 - a. Enter the name and organization details.
 - b. Select Git in the Source Control Credential Type field.
 - c. enter https://github.com/NetApp-Automation/na_oracle19c_deploy.git as the source control URL.
 - d. Click Save.
 - e. The project might need to sync occasionally when the source code changes.

3. Configure Oracle host_vars

The variables defined in this section are applied to each individual Oracle server and database.

1. Input your environment-specific parameters in the following embedded Oracle hosts variables or host_vars form.



The items in blue must be changed to match your environment.

Host VARS Config

```
#####
##### Host Variables Configuration #####
#####

# Add your Oracle Host
ansible_host: "10.61.180.15"

# Oracle db log archive mode: true - ARCHIVELOG or false - NOARCHIVELOG
log_archive_mode: "true"

# Number of pluggable databases per container instance identified by sid.
Pdb_name specifies the prefix for container database naming in this case
cdb2_pdb1, cdb2_pdb2, cdb2_pdb3
oracle_sid: "cdb2"
pdb_num: "3"
pdb_name: "{{ oracle_sid }}_pdb"

# CDB listener port, use different listener port for additional CDB on
same host
listener_port: "1523"

# CDB is created with SGA at 75% of memory_limit, MB. Consider how many
databases to be hosted on the node and how much ram to be allocated to
each DB. The grand total SGA should not exceed 75% available RAM on node.
memory_limit: "5464"

# Set "em_configuration: DBEXPRESS" to install enterprise manager express
and choose a unique port from 5500 to 5599 for each sid on the host.
# Leave them black if em express is not installed.
em_configuration: "DBEXPRESS"
em_express_port: "5501"

# {{groups.oracle[0]}} represents first Oracle DB server as defined in
Oracle hosts group [oracle]. For concurrent multiple Oracle DB servers
deployment, [0] will be incremented for each additional DB server. For
example, {{groups.oracle[1]}}" represents DB server 2,
```

"{{groups.oracle[2]}}" represents DB server 3 ... As a good practice and the default, minimum three volumes is allocated to a DB server with corresponding /u01, /u02, /u03 mount points, which store oracle binary, oracle data, and oracle recovery files respectively. Additional volumes can be added by click on "More NFS volumes" but the number of volumes allocated to a DB server must match with what is defined in global vars file by volumes_nfs parameter, which dictates how many volumes are to be created for each DB server.

```
host_datastores_nfs:
  - {vol_name: "{{groups.oracle[0]}}_u01", aggr_name: "aggr01_node01",
    lif: "172.21.94.200", size: "25"}
  - {vol_name: "{{groups.oracle[0]}}_u02", aggr_name: "aggr01_node01",
    lif: "172.21.94.200", size: "25"}
  - {vol_name: "{{groups.oracle[0]}}_u03", aggr_name: "aggr01_node01",
    lif: "172.21.94.200", size: "25"}
```

- a. Fill in all variables in the blue fields.
- b. After completing variables input, click the Copy button on the form to copy all variables to be transferred to AWX or Tower.
- c. Navigate back to AWX or Tower and go to Resources → Hosts, and select and open the Oracle server configuration page.
- d. Under the Details tab, click edit and paste the copied variables from step 1 to the Variables field under the YAML tab.
- e. Click Save.
- f. Repeat this process for any additional Oracle servers in the system.

4. Configure global variables

Variables defined in this section apply to all Oracle hosts, databases, and the ONTAP cluster.

1. Input your environment-specific parameters in following embedded global variables or vars form.



The items in blue must be changed to match your environment.

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

#####
### Ontap env specific config variables ###
#####

#Inventory group name
#Default inventory group name - 'ontap'
```

```

#Change only if you are changing the group name either in inventory/hosts
file or in inventory groups in case of AWX/Tower
hosts_group: "ontap"

#CA_signed_certificates (ONLY CHANGE to 'true' IF YOU ARE USING CA SIGNED
CERTIFICATES)
ca_signed_certs: "false"

#Names of the Nodes in the ONTAP Cluster
nodes:
  - "AFF-01"
  - "AFF-02"

#Storage VLANs
#Add additional rows for vlans as necessary
storage_vlans:
  - {vlan_id: "203", name: "infra_NFS", protocol: "NFS"}
More Storage VLANsEnter Storage VLANs details

#Details of the Data Aggregates that need to be created
#If Aggregate creation takes longer, subsequent tasks of creating volumes
may fail.
#There should be enough disks already zeroed in the cluster, otherwise
aggregate create will zero the disks and will take long time
data_aggregates:
  - {aggr_name: "aggr01_node01"}
  - {aggr_name: "aggr01_node02"}

#SVM name
svm_name: "ora_svm"

# SVM Management LIF Details
svm_mgmt_details:
  - {address: "172.21.91.100", netmask: "255.255.255.0", home_port: "e0M"}

# NFS storage parameters when data_protocol set to NFS. Volume named after
Oracle hosts name identified by mount point as follow for oracle DB server
1. Each mount point dedicates to a particular Oracle files: u01 - Oracle
binary, u02 - Oracle data, u03 - Oracle redo. Add additional volumes by
click on "More NFS volumes" and also add the volumes list to corresponding
host_vars as host_datastores_nfs variable. For multiple DB server
deployment, additional volumes sets needs to be added for additional DB
server. Input variable "{{groups.oracle[1]}}_u01",
 "{{groups.oracle[1]}}_u02", and "{{groups.oracle[1]}}_u03" as vol_name for
second DB server. Place volumes for multiple DB servers alternatingly
between controllers for balanced IO performance, e.g. DB server 1 on

```

controller node1, DB server 2 on controller node2 etc. Make sure match lif address with controller node.

volumes_nfs:

```
- {vol_name: "{{groups.oracle[0]}}_u01", aggr_name: "aggr01_node01",  
lif: "172.21.94.200", size: "25"}  
- {vol_name: "{{groups.oracle[0]}}_u02", aggr_name: "aggr01_node01",  
lif: "172.21.94.200", size: "25"}  
- {vol_name: "{{groups.oracle[0]}}_u03", aggr_name: "aggr01_node01",  
lif: "172.21.94.200", size: "25"}
```

#NFS LIFs IP address and netmask

nfs_lifs_details:

```
- address: "172.21.94.200" #for node-1  
  netmask: "255.255.255.0"  
- address: "172.21.94.201" #for node-2  
  netmask: "255.255.255.0"
```

#NFS client match

client_match: "172.21.94.0/24"

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

#NFS Mount points for Oracle DB volumes

mount_points:

```
- "/u01"  
- "/u02"  
- "/u03"
```

Up to 75% of node memory size divided by 2mb. Consider how many databases to be hosted on the node and how much ram to be allocated to each DB.

Leave it blank if hugepage is not configured on the host.

hugepages_nr: "1234"

RedHat subscription username and password

```
redhat_sub_username: "xxx"  
redhat_sub_password: "xxx"
```

```
#####
```

```

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

db_domain: "your.domain.com"

# Set initial password for all required Oracle passwords. Change them
after installation.

initial_pwd_all: "netapp123"

```

1. Fill in all variables in blue fields.
2. After completing variables input, click the Copy button on the form to copy all variables to be transferred to AWX or Tower into the following job template.

5. Configure and launch the job template.

1. Create the job template.
 - a. Navigate to Resources → Templates → Add and click Add Job Template.
 - b. Enter the name and description
 - c. Select the Job type; Run configures the system based on a playbook, and Check performs a dry run of a playbook without actually configuring the system.
 - d. Select the corresponding inventory, project, playbook, and credentials for the playbook.
 - e. Select the all_playbook.yml as the default playbook to be executed.
 - f. Paste global variables copied from step 4 into the Template Variables field under the YAML tab.
 - g. Check the box Prompt on Launch in the Job Tags field.
 - h. Click Save.
2. Launch the job template.
 - a. Navigate to Resources → Templates.
 - b. Click the desired template and then click Launch.
 - c. When prompted on launch for Job Tags, type in requirements_config. You might need to click the Create Job Tag line below requirements_config to enter the job tag.



requirements_config ensures that you have the correct libraries to run the other roles.

- a. Click Next and then Launch to start the job.
- b. Click View → Jobs to monitor the job output and progress.
- c. When prompted on launch for Job Tags, type in ontap_config. You might need to click the Create "Job Tag" line right below ontap_config to enter the job tag.
- d. Click Next and then Launch to start the job.
- e. Click View → Jobs to monitor the job output and progress
- f. After the ontap_config role has completed, run the process again for linux_config.
- g. Navigate to Resources → Templates.

- h. Select the desired template and then click Launch.
- i. When prompted on launch for the Job Tags type in `linux_config`, you might need to select the Create "job tag" line right below `linux_config` to enter the job tag.
- j. Click Next and then Launch to start the job.
- k. Select View → Jobs to monitor the job output and progress.
- l. After the `linux_config` role has completed, run the process again for `oracle_config`.
- m. Go to Resources → Templates.
- n. Select the desired template and then click Launch.
- o. When prompted on launch for Job Tags, type `oracle_config`. You might need to select the Create "Job Tag" line right below `oracle_config` to enter the job tag.
- p. Click Next and then Launch to start the job.
- q. Select View → Jobs to monitor the job output and progress.

6. Deploy additional database on same Oracle host

The Oracle portion of the playbook creates a single Oracle container database on an Oracle server per execution. To create additional container databases on the same server, complete the following steps.

1. Revise `host_vars` variables.
 - a. Go back to step 2 - Configure Oracle `host_vars`.
 - b. Change the Oracle SID to a different naming string.
 - c. Change the listener port to different number.
 - d. Change the EM Express port to a different number if you are installing EM Express.
 - e. Copy and paste the revised host variables to the Oracle Host Variables field in the Host Configuration Detail tab.
2. Launch the deployment job template with only the `oracle_config` tag.
3. Log in to Oracle server as oracle user and execute the following commands:

```
ps -ef | grep ora
```



This will list oracle processes if installation completed as expected and oracle DB started

4. Log in to the database to check the db configuration settings and the PDBs created with the following command sets.

```

[oracle@localhost ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Thu May 6 12:52:51 2021
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>

SQL> select name, log_mode from v$database;
NAME          LOG_MODE
-----
CDB2          ARCHIVELOG

SQL> show pdbs

          CON_ID CON_NAME                                OPEN MODE  RESTRICTED
-----
          2 PDB$SEED                                READ ONLY  NO
          3 CDB2_PDB1                            READ WRITE NO
          4 CDB2_PDB2                            READ WRITE NO
          5 CDB2_PDB3                            READ WRITE NO

col svrname form a30
col dirname form a30
select svrname, dirname, nfsversion from v$dnfs_servers;

SQL> col svrname form a30
SQL> col dirname form a30
SQL> select svrname, dirname, nfsversion from v$dnfs_servers;

SVRNAME          DIRNAME                                NFSVERSION
-----
172.21.126.200   /rhelora03_u02                        NFSv3.0
172.21.126.200   /rhelora03_u03                        NFSv3.0
172.21.126.200   /rhelora03_u01                        NFSv3.0

```

This confirms that dNFS is working properly.

5. Connect to database via listener to check the Oracle listener configuration with the following command. Change to the appropriate listener port and database service name.

```
[oracle@localhost ~]$ sqlplus
system@//localhost:1523/cdb2_pdb1.cie.netapp.com

SQL*Plus: Release 19.0.0.0.0 - Production on Thu May 6 13:19:57 2021
Version 19.8.0.0.0

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

Enter password:
Last Successful login time: Wed May 05 2021 17:11:11 -04:00

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

SQL> show user
USER is "SYSTEM"
SQL> show con_name
CON_NAME
CDB2_PDB1
```

This confirms that Oracle listener is working properly.

Where to go for help?

If you need help with the toolkit, please join the [NetApp Solution Automation community support slack channel](#) and look for the solution-automation channel to post your questions or inquires.

Step-by-step deployment procedure

This document details the deployment of Oracle 19c using the automation command line interface (cli).

CLI deployment Oracle 19c Database

This section covers the steps required to prepare and deploy Oracle19c Database with the CLI. Make sure that you have reviewed the [Getting Started and Requirements section](#) and prepared your environment accordingly.

Download Oracle19c repo

1. From your ansible controller, run the following command:

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

2. After downloading the repository, change directories to na_oracle19c_deploy <cd na_oracle19c_deploy>.

Edit the hosts file

Complete the following before deployment:

1. Edit your hosts file `na_oracle19c_deploy` directory.
2. Under `[ontap]`, change the IP address to your cluster management IP.
3. Under the `[oracle]` group, add the oracle hosts names. The host name must be resolved to its IP address either through DNS or the hosts file, or it must be specified in the host.
4. After you have completed these steps, save any changes.

The following example depicts a host file:

```
#ONTAP Host

[ontap]

"10.61.184.183"

#Oracle hosts

[oracle]

"rtpora01"

"rtpora02"
```

This example executes the playbook and deploys oracle 19c on two oracle DB servers concurrently. You can also test with just one DB server. In that case, you only need to configure one host variable file.



The playbook executes the same way regardless of how many Oracle hosts and databases you deploy.

Edit the `host_name.yml` file under `host_vars`

Each Oracle host has its host variable file identified by its host name that contains host-specific variables. You can specify any name for your host. Edit and copy the `host_vars` from the Host VARS Config section and paste it into your desired `host_name.yml` file.



The items in blue must be changed to match your environment.

Host VARS Config

```
#####
##### Host Variables Configuration #####
#####

# Add your Oracle Host
```

```

ansible_host: "10.61.180.15"

# Oracle db log archive mode: true - ARCHIVELOG or false - NOARCHIVELOG
log_archive_mode: "true"

# Number of pluggable databases per container instance identified by sid.
Pdb_name specifies the prefix for container database naming in this case
cdb2_pdb1, cdb2_pdb2, cdb2_pdb3
oracle_sid: "cdb2"
pdb_num: "3"
pdb_name: "{{ oracle_sid }}_pdb"

# CDB listener port, use different listener port for additional CDB on
same host
listener_port: "1523"

# CDB is created with SGA at 75% of memory_limit, MB. Consider how many
databases to be hosted on the node and how much ram to be allocated to
each DB. The grand total SGA should not exceed 75% available RAM on node.
memory_limit: "5464"

# Set "em_configuration: DBEXPRESS" to install enterprise manager express
and choose a unique port from 5500 to 5599 for each sid on the host.
# Leave them blank if em express is not installed.
em_configuration: "DBEXPRESS"
em_express_port: "5501"

# {{groups.oracle[0]}} represents first Oracle DB server as defined in
Oracle hosts group [oracle]. For concurrent multiple Oracle DB servers
deployment, [0] will be incremented for each additional DB server. For
example, {{groups.oracle[1]}}" represents DB server 2,
"{{groups.oracle[2]}}" represents DB server 3 ... As a good practice and
the default, minimum three volumes is allocated to a DB server with
corresponding /u01, /u02, /u03 mount points, which store oracle binary,
oracle data, and oracle recovery files respectively. Additional volumes
can be added by click on "More NFS volumes" but the number of volumes
allocated to a DB server must match with what is defined in global vars
file by volumes_nfs parameter, which dictates how many volumes are to be
created for each DB server.
host_datastores_nfs:
  - {vol_name: "{{groups.oracle[0]}}_u01", aggr_name: "aggr01_node01",
lif: "172.21.94.200", size: "25"}
  - {vol_name: "{{groups.oracle[0]}}_u02", aggr_name: "aggr01_node01",
lif: "172.21.94.200", size: "25"}
  - {vol_name: "{{groups.oracle[0]}}_u03", aggr_name: "aggr01_node01",
lif: "172.21.94.200", size: "25"}

```

Edit the vars.yml file

The `vars.yml` file consolidates all environment-specific variables (ONTAP, Linux, or Oracle) for Oracle deployment.

1. Edit and copy the variables from the VARS section and paste these variables into your `vars.yml` file.

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

#####
### Ontap env specific config variables ###
#####

#Inventory group name
#Default inventory group name - 'ontap'
#Change only if you are changing the group name either in inventory/hosts
file or in inventory groups in case of AWX/Tower
hosts_group: "ontap"

#CA_signed_certificates (ONLY CHANGE to 'true' IF YOU ARE USING CA SIGNED
CERTIFICATES)
ca_signed_certs: "false"

#Names of the Nodes in the ONTAP Cluster
nodes:
  - "AFF-01"
  - "AFF-02"

#Storage VLANs
#Add additional rows for vlans as necessary
storage_vlans:
  - {vlan_id: "203", name: "infra_NFS", protocol: "NFS"}
More Storage VLANsEnter Storage VLANs details

#Details of the Data Aggregates that need to be created
#If Aggregate creation takes longer, subsequent tasks of creating volumes
may fail.
#There should be enough disks already zeroed in the cluster, otherwise
aggregate create will zero the disks and will take long time
data_aggregates:
  - {aggr_name: "aggr01_node01"}
  - {aggr_name: "aggr01_node02"}

#SVM name
```

```

svm_name: "ora_svm"

# SVM Management LIF Details
svm_mgmt_details:
  - {address: "172.21.91.100", netmask: "255.255.255.0", home_port: "e0M"}

# NFS storage parameters when data_protocol set to NFS. Volume named after
Oracle hosts name identified by mount point as follow for oracle DB server
1. Each mount point dedicates to a particular Oracle files: u01 - Oracle
binary, u02 - Oracle data, u03 - Oracle redo. Add additional volumes by
click on "More NFS volumes" and also add the volumes list to corresponding
host_vars as host_datastores_nfs variable. For multiple DB server
deployment, additional volumes sets needs to be added for additional DB
server. Input variable "{{groups.oracle[1]}}_u01",
 "{{groups.oracle[1]}}_u02", and "{{groups.oracle[1]}}_u03" as vol_name for
second DB server. Place volumes for multiple DB servers alternately
between controllers for balanced IO performance, e.g. DB server 1 on
controller node1, DB server 2 on controller node2 etc. Make sure match lif
address with controller node.

volumes_nfs:
  - {vol_name: "{{groups.oracle[0]}}_u01", aggr_name: "aggr01_node01",
lif: "172.21.94.200", size: "25"}
  - {vol_name: "{{groups.oracle[0]}}_u02", aggr_name: "aggr01_node01",
lif: "172.21.94.200", size: "25"}
  - {vol_name: "{{groups.oracle[0]}}_u03", aggr_name: "aggr01_node01",
lif: "172.21.94.200", size: "25"}

#NFS LIFs IP address and netmask

nfs_lifs_details:
  - address: "172.21.94.200" #for node-1
    netmask: "255.255.255.0"
  - address: "172.21.94.201" #for node-2
    netmask: "255.255.255.0"

#NFS client match

client_match: "172.21.94.0/24"

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

#NFS Mount points for Oracle DB volumes

mount_points:

```

```

- "/u01"
- "/u02"
- "/u03"

# Up to 75% of node memory size divided by 2mb. Consider how many
databases to be hosted on the node and how much ram to be allocated to
each DB.
# Leave it blank if hugepage is not configured on the host.

hugepages_nr: "1234"

# RedHat subscription username and password

redhat_sub_username: "xxx"
redhat_sub_password: "xxx"

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

db_domain: "your.domain.com"

# Set initial password for all required Oracle passwords. Change them
after installation.

initial_pwd_all: "netappl23"

```

Run the playbook

After completing the required environment prerequisites and copying the variables into `vars.yml` and `your_host.yml`, you are now ready to deploy the playbooks.



<username> must be changed to match your environment.

1. Run the ONTAP playbook by passing the correct tags and ONTAP cluster username. Fill the password for ONTAP cluster, and vsadmin when prompted.

```

ansible-playbook -i hosts all_playbook.yml -u username -k -K -t
ontap_config -e @vars/vars.yml

```

2. Run the Linux playbook to execute Linux portion of deployment. Input for admin ssh password as well as sudo password.

```

ansible-playbook -i hosts all_playbook.yml -u username -k -K -t
linux_config -e @vars/vars.yml

```

3. Run the Oracle playbook to execute Oracle portion of deployment. Input for admin ssh password as well as sudo password.

```
ansible-playbook -i hosts all_playbook.yml -u username -k -K -t
oracle_config -e @vars/vars.yml
```

Deploy Additional Database on Same Oracle Host

The Oracle portion of the playbook creates a single Oracle container database on an Oracle server per execution. To create additional container database on the same server, complete the following steps:

1. Revise the `host_vars` variables.
 - a. Go back to step 3 - Edit the `host_name.yml` file under `host_vars`.
 - b. Change the Oracle SID to a different naming string.
 - c. Change the listener port to different number.
 - d. Change the EM Express port to a different number if you have installed EM Express.
 - e. Copy and paste the revised host variables to the Oracle host variable file under `host_vars`.
2. Execute the playbook with the `oracle_config` tag as shown above in [Run the playbook](#).

Validate Oracle installation

1. Log in to Oracle server as oracle user and execute the following commands:

```
ps -ef | grep ora
```



This will list oracle processes if installation completed as expected and oracle DB started

2. Log in to the database to check the db configuration settings and the PDBs created with the following command sets.

```

[oracle@localhost ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Thu May 6 12:52:51 2021
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>

SQL> select name, log_mode from v$database;
NAME          LOG_MODE
-----
CDB2          ARCHIVELOG

SQL> show pdbs

          CON_ID CON_NAME                                OPEN MODE  RESTRICTED
-----
          2 PDB$SEED                                READ ONLY  NO
          3 CDB2_PDB1                            READ WRITE NO
          4 CDB2_PDB2                            READ WRITE NO
          5 CDB2_PDB3                            READ WRITE NO

col svrname form a30
col dirname form a30
select svrname, dirname, nfsversion from v$dnfs_servers;

SQL> col svrname form a30
SQL> col dirname form a30
SQL> select svrname, dirname, nfsversion from v$dnfs_servers;

SVRNAME                                DIRNAME                                NFSVERSION
-----
172.21.126.200                        /rhelora03_u02                        NFSv3.0
172.21.126.200                        /rhelora03_u03                        NFSv3.0
172.21.126.200                        /rhelora03_u01                        NFSv3.0

```

This confirms that dNFS is working properly.

3. Connect to database via listener to check the Oracle listener configuration with the following command. Change to the appropriate listener port and database service name.

```
[oracle@localhost ~]$ sqlplus
system@//localhost:1523/cdb2_pdb1.cie.netapp.com

SQL*Plus: Release 19.0.0.0.0 - Production on Thu May 6 13:19:57 2021
Version 19.8.0.0.0

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

Enter password:
Last Successful login time: Wed May 05 2021 17:11:11 -04:00

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

SQL> show user
USER is "SYSTEM"
SQL> show con_name
CON_NAME
CDB2_PDB1
```

This confirms that Oracle listener is working properly.

Where to go for help?

If you need help with the toolkit, please join the [NetApp Solution Automation community support slack channel](#) and look for the solution-automation channel to post your questions or inquires.

Solution Overview

This page describes the Automated method for deploying Oracle19c on NetApp ONTAP storage.

Automated Data Protection for Oracle Databases

Organizations are automating their environments to gain efficiencies, accelerate deployments, and reduce manual effort. Configuration management tools like Ansible are being used to streamline enterprise database operations. In this solution, we demonstrate how you can use Ansible to automate the data protection of Oracle with NetApp ONTAP. By enabling storage administrators, systems administrators, and DBAs to consistently and rapidly setup data replication to an offsite data center or to public cloud, you achieve the following benefits:

- Eliminate design complexities and human errors, and implement a repeatable consistent deployment and best practices
- Decrease time for configuration of Intercluster replication, CVO instantiation, and recovery of Oracle databases
- Increase database administrators, systems and storage administrators productivity
- Provides database recovery workflow for ease of testing a DR scenario.

NetApp provides customers with validated Ansible modules and roles to accelerate deployment, configuration, and lifecycle management of your Oracle database environment. This solution provides instruction and Ansible playbook code, to help you:

On Prem to on prem replication

- Create intercluster lifs on source and destination
- Establish cluster and vserver peering
- Create and initialize SnapMirror of Oracle volumes
- Create a replication schedule through AWX/Tower for Oracle binaries, databases, and logs
- Restore Oracle DB on the destination, and bring database online

On Prem to CVO in AWS

- Create AWS connector
- Create CVO instance in AWS
- Add On-Prem cluster to Cloud Manager
- Create intercluster lifs on source
- Establish cluster and vserver peering
- Create and initialize SnapMirror of Oracle volumes
- Create a replication schedule through AWX/Tower for Oracle binaries, databases, and logs
- Restore Oracle DB on the destination, and bring database online

After you are ready, click [here for getting started with the solution](#).

Getting started

This solution has been designed to be run in an AWX/Tower environment.

AWX/Tower

For AWX/Tower environments, you are guided through creating an inventory of your ONTAP cluster management and Oracle server (IPs and hostnames), creating credentials, configuring a project that pulls the Ansible code from NetApp Automation Github, and the Job Template that launches the automation.

1. The solution has been designed to run in a private cloud scenario (on-premise to on-premise), and hybrid cloud (on-premise to public cloud Cloud Volumes ONTAP [CVO])
2. Fill out the variables specific to your environment, and copy and paste them into the Extra Vars fields in your job template.
3. After the extra vars have been added to your job template, you can launch the automation.
4. The automation is set to be ran three phases (Setup, Replication Schedule for Oracle Binaries, Database, Logs, and Replication Schedule just for Logs), and a forth phase to recovering the database at a DR site.
5. For detailed instructions for obtaining the keys and tokens necessary for the CVO Data Protection visit [Gather Pre-requisites For CVO and Connector Deployments](#)

Requirements

On-Prem

Environment	Requirements
Ansible environment	AWX/Tower
	Ansible v.2.10 and higher
	Python 3
	Python libraries - netapp-lib - xmldict - jmespath
ONTAP	ONTAP version 9.8 +
	Two data aggregates
	NFS vlan and ifgrp created
Oracle server(s)	RHEL 7/8
	Oracle Linux 7/8
	Network interfaces for NFS, public, and optional mgmt
	Existing Oracle environment on source, and the equivalent Linux operating system at the destination (DR Site or Public Cloud)

CVO

Environment	Requirements
Ansible environment	AWX/Tower
	Ansible v.2.10 and higher
	Python 3
	Python libraries - netapp-lib - xmldict - jmespath
ONTAP	ONTAP version 9.8 +
	Two data aggregates
	NFS vlan and ifgrp created
Oracle server(s)	RHEL 7/8
	Oracle Linux 7/8
	Network interfaces for NFS, public, and optional mgmt
	Existing Oracle environment on source, and the equivalent Linux operating system at the destination (DR Site or Public Cloud)
	Set appropriate swap space on the Oracle EC2 instance, by default some EC2 instances are deployed with 0 swap

Environment	Requirements
Cloud Manager/AWS	AWS Access/Secret Key
	NetApp Cloud Manager Account
	NetApp Cloud Manager Refresh Token

Automation Details

On-Prem |

This automated deployment is designed with a single Ansible playbook that consists of three separate roles. The roles are for ONTAP, Linux, and Oracle configurations. The following table describes which tasks are being automated.

Playbook	Tasks
ontap_setup	Pre-check of the ONTAP environment
	Creation of Intercluster LIFs on source cluster (OPTIONAL)
	Creation of Intercluster LIFs on destination cluster (OPTIONAL)
	Creation of Cluster and SVM Peering
	Creation of destination SnapMirror and Initialization of designated Oracle volumes
ora_replication_cg	Enable backup mode for each database in /etc/oratab
	Snapshot taken of Oracle Binary and Database volumes
	Snapmirror Updated
	Turn off backup mode for each database in /etc/oratab
ora_replication_log	Switch current log for each database in /etc/oratab
	Snapshot taken of Oracle Log volume
	Snapmirror Updated
ora_recovery	Break SnapMirror
	Enable NFS and create junction path for Oracle volumes on the destination
	Configure DR Oracle Host
	Mount and verify Oracle volumes
	Recover and start Oracle database

CVO

This automated deployment is designed with a single Ansible playbook that consists of three separate roles. The roles are for ONTAP, Linux, and Oracle configurations. The following table describes which tasks are being automated.

Playbook	Tasks
cvo_setup	Pre-check of the environment
	AWS Configure/AWS Access Key ID/Secret Key/Default Region
	Creation of AWS Role
	Creation of NetApp Cloud Manager Connector instance in AWS
	Creation of Cloud Volumes ONTAP (CVO) instance in AWS
	Add On-Prem Source ONTAP Cluster to NetApp Cloud Manager
	Creation of destination SnapMirror and Initialization of designated Oracle volumes
ora_replication_cg	Enable backup mode for each database in /etc/oratab
	Snapshot taken of Oracle Binary and Database volumes
	Snapmirror Updated
	Turn off backup mode for each database in /etc/oratab
ora_replication_log	Switch current log for each database in /etc/oratab
	Snapshot taken of Oracle Log volume
	Snapmirror Updated
ora_recovery	Break SnapMirror
	Enable NFS and create junction path for Oracle volumes on the destination CVO
	Configure DR Oracle Host
	Mount and verify Oracle volumes
	Recover and start Oracle database

Default parameters

To simplify automation, we have preset many required Oracle parameters with default values. It is generally not necessary to change the default parameters for most deployments. A more advanced user can make changes to the default parameters with caution. The default parameters are located in each role folder under defaults directory.

License

You should read license information as stated in the Github repository. By accessing, downloading, installing, or using the content in this repository, you agree the terms of the license laid out [here](#).

Note that there are certain restrictions around producing and/or sharing any derivative works with the content in this repository. Please make sure you read the terms of the [License](#) before using the content. If you do not agree to all of the terms, do not access, download, or use the content in this repository.

After you are ready, click [here for detailed AWX/Tower procedures](#).

Step-by-step deployment procedure

This page describes the Automated Data Protection of Oracle19c on NetApp ONTAP storage.

AWX/Tower Oracle Data Protection

Create the inventory, group, hosts, and credentials for your environment

This section describes the setup of inventory, groups, hosts, and access credentials in AWX/Ansible Tower that prepare the environment for consuming NetApp automated solutions.

1. Configure the inventory.
 - a. Navigate to Resources → Inventories → Add, and click Add Inventory.
 - b. Provide the name and organization details, and click Save.
 - c. On the Inventories page, click the inventory created.
 - d. Navigate to the Groups sub-menu and click Add.
 - e. Provide the name oracle for your first group and click Save.
 - f. Repeat the process for a second group called dr_oracle.
 - g. Select the oracle group created, go to the Hosts sub-menu and click Add New Host.
 - h. Provide the IP address of the Source Oracle host's management IP, and click Save.
 - i. This process must be repeated for the dr_oracle group and add the the DR/Destination Oracle host's management IP/hostname.



Below are instructions for creating the credential types and credentials for either On-Prem with ONTAP, or CVO on AWS.

On-Prem

1. Configure the credentials.
2. Create Credential Types. For solutions involving ONTAP, you must configure the credential type to match username and password entries.
 - a. Navigate to Administration → Credential Types, and click Add.
 - b. Provide the name and description.
 - c. Paste the following content in Input Configuration:

```
fields:  
  - id: dst_cluster_username  
    type: string  
    label: Destination Cluster Username  
  - id: dst_cluster_password  
    type: string  
    label: Destination Cluster Password  
    secret: true  
  - id: src_cluster_username  
    type: string  
    label: Source Cluster Username  
  - id: src_cluster_password  
    type: string  
    label: Source Cluster Password  
    secret: true
```

- d. Paste the following content into Injector Configuration and then click Save:

```
extra_vars:  
  dst_cluster_username: '{{ dst_cluster_username }}'  
  dst_cluster_password: '{{ dst_cluster_password }}'  
  src_cluster_username: '{{ src_cluster_username }}'  
  src_cluster_password: '{{ src_cluster_password }}'
```

3. Create Credential for ONTAP
 - a. Navigate to Resources → Credentials, and click Add.
 - b. Enter the name and organization details for the ONTAP Credentials
 - c. Select the credential type that was created in the previous step.
 - d. Under Type Details, enter the Username and Password for your Source and Destination Clusters.
 - e. Click Save
4. Create Credential for Oracle
 - a. Navigate to Resources → Credentials, and click Add.
 - b. Enter the name and organization details for Oracle

- c. Select the Machine credential type.
- d. Under Type Details, enter the Username and Password for the Oracle hosts.
- e. Select the correct Privilege Escalation Method, and enter the username and password.
- f. Click Save
- g. Repeat process if needed for a different credential for the dr_oracle host.

CVO

1. Configure the credentials.
2. Create credential types. For solutions involving ONTAP, you must configure the credential type to match username and password entries, we will also add entries for Cloud Central and AWS.
 - a. Navigate to Administration → Credential Types, and click Add.
 - b. Provide the name and description.
 - c. Paste the following content in Input Configuration:

```
fields:
  - id: dst_cluster_username
    type: string
    label: CVO Username
  - id: dst_cluster_password
    type: string
    label: CVO Password
    secret: true
  - id: cvo_svm_password
    type: string
    label: CVO SVM Password
    secret: true
  - id: src_cluster_username
    type: string
    label: Source Cluster Username
  - id: src_cluster_password
    type: string
    label: Source Cluster Password
    secret: true
  - id: regular_id
    type: string
    label: Cloud Central ID
    secret: true
  - id: email_id
    type: string
    label: Cloud Manager Email
    secret: true
  - id: cm_password
    type: string
    label: Cloud Manager Password
    secret: true
  - id: access_key
    type: string
    label: AWS Access Key
    secret: true
  - id: secret_key
    type: string
    label: AWS Secret Key
    secret: true
  - id: token
    type: string
    label: Cloud Central Refresh Token
    secret: true
```

d. Paste the following content into Injector Configuration and click Save:

```

extra_vars:
  dst_cluster_username: '{{ dst_cluster_username }}'
  dst_cluster_password: '{{ dst_cluster_password }}'
  cvo_svm_password: '{{ cvo_svm_password }}'
  src_cluster_username: '{{ src_cluster_username }}'
  src_cluster_password: '{{ src_cluster_password }}'
  regular_id: '{{ regular_id }}'
  email_id: '{{ email_id }}'
  cm_password: '{{ cm_password }}'
  access_key: '{{ access_key }}'
  secret_key: '{{ secret_key }}'
  token: '{{ token }}'

```

3. Create Credential for ONTAP/CVO/AWS

- a. Navigate to Resources → Credentials, and click Add.
- b. Enter the name and organization details for the ONTAP Credentials
- c. Select the credential type that was created in the previous step.
- d. Under Type Details, enter the Username and Password for your Source and CVO Clusters, Cloud Central/Manager, AWS Access/Secret Key and Cloud Central Refresh Token.
- e. Click Save

4. Create Credential for Oracle (Source)

- a. Navigate to Resources → Credentials, and click Add.
- b. Enter the name and organization details for Oracle host
- c. Select the Machine credential type.
- d. Under Type Details, enter the Username and Password for the Oracle hosts.
- e. Select the correct Privilege Escalation Method, and enter the username and password.
- f. Click Save

5. Create Credential for Oracle Destination

- a. Navigate to Resources → Credentials, and click Add.
- b. Enter the name and organization details for the DR Oracle host
- c. Select the Machine credential type.
- d. Under Type Details, enter the Username (ec2-user or if you have changed it from default enter that), and the SSH Private Key
- e. Select the correct Privilege Escalation Method (sudo), and enter the username and password if needed.
- f. Click Save

Create a project

1. Go to Resources → Projects, and click Add.

- a. Enter the name and organization details.
- b. Select Git in the Source Control Credential Type field.
- c. enter https://github.com/NetApp-Automation/na_oracle19c_data_protection.git as the source control URL.
- d. Click Save.
- e. The project might need to sync occasionally when the source code changes.

Configure global variables

Variables defined in this section apply to all Oracle hosts, databases, and the ONTAP cluster.

1. Input your environment-specific parameters in following embedded global variables or vars form.



The items in blue must be changed to match your environment.

On-Prem

```
# Oracle Data Protection global user configuration variables
# Ontap env specific config variables
hosts_group: "ontap"
ca_signed_certs: "false"

# Inter-cluster LIF details
src_nodes:
  - "AFF-01"
  - "AFF-02"

dst_nodes:
  - "DR-AFF-01"
  - "DR-AFF-02"

create_source_intercluster_lifs: "yes"

source_intercluster_network_port_details:
  using_dedicated_ports: "yes"
  using_ifgrp: "yes"
  using_vlans: "yes"
  failover_for_shared_individual_ports: "yes"
  ifgrp_name: "a0a"
  vlan_id: "10"
  ports:
    - "e0b"
    - "e0g"
  broadcast_domain: "NFS"
  ipspace: "Default"
  failover_group_name: "iclifs"

source_intercluster_lif_details:
  - name: "icl_1"
    address: "10.0.0.1"
    netmask: "255.255.255.0"
    home_port: "a0a-10"
    node: "AFF-01"
  - name: "icl_2"
    address: "10.0.0.2"
    netmask: "255.255.255.0"
    home_port: "a0a-10"
    node: "AFF-02"

create_destination_intercluster_lifs: "yes"
```

```

destination_intercluster_network_port_details:
  using_dedicated_ports: "yes"
  using_ifgrp: "yes"
  using_vlans: "yes"
  failover_for_shared_individual_ports: "yes"
  ifgrp_name: "a0a"
  vlan_id: "10"
  ports:
    - "e0b"
    - "e0g"
  broadcast_domain: "NFS"
  ipspace: "Default"
  failover_group_name: "iclifs"

destination_intercluster_lif_details:
  - name: "icl_1"
    address: "10.0.0.3"
    netmask: "255.255.255.0"
    home_port: "a0a-10"
    node: "DR-AFF-01"
  - name: "icl_2"
    address: "10.0.0.4"
    netmask: "255.255.255.0"
    home_port: "a0a-10"
    node: "DR-AFF-02"

# Variables for SnapMirror Peering
passphrase: "your-passphrase"

# Source & Destination List
dst_cluster_name: "dst-cluster-name"
dst_cluster_ip: "dst-cluster-ip"
dst_vserver: "dst-vserver"
dst_nfs_lif: "dst-nfs-lif"
src_cluster_name: "src-cluster-name"
src_cluster_ip: "src-cluster-ip"
src_vserver: "src-vserver"

# Variable for Oracle Volumes and SnapMirror Details
cg_snapshot_name_prefix: "oracle"
src_orabinary_vols:
  - "binary_vol"
src_db_vols:
  - "db_vol"
src_archivelog_vols:
  - "log_vol"

```

```

snapmirror_policy: "async_policy_oracle"

# Export Policy Details
export_policy_details:
  name: "nfs_export_policy"
  client_match: "0.0.0.0/0"
  ro_rule: "sys"
  rw_rule: "sys"

# Linux env specific config variables
mount_points:
  - "/u01"
  - "/u02"
  - "/u03"
hugepages_nr: "1234"
redhat_sub_username: "xxx"
redhat_sub_password: "xxx"

# DB env specific install and config variables
recovery_type: "scn"
control_files:
  - "/u02/oradata/CDB2/control01.ctl"
  - "/u03/orareco/CDB2/control02.ctl"

```

CVO

```

#####
### Ontap env specific config variables ###
#####

#Inventory group name
#Default inventory group name - "ontap"
#Change only if you are changing the group name either in
inventory/hosts file or in inventory groups in case of AWX/Tower
hosts_group: "ontap"

#CA_signed_certificates (ONLY CHANGE to "true" IF YOU ARE USING CA
SIGNED CERTIFICATES)
ca_signed_certs: "false"

#Names of the Nodes in the Source ONTAP Cluster
src_nodes:
  - "AFF-01"
  - "AFF-02"

#Names of the Nodes in the Destination CVO Cluster

```

```

dst_nodes:
  - "DR-AFF-01"
  - "DR-AFF-02"

#Define whether or not to create intercluster lifs on source cluster
(ONLY CHANGE to "No" IF YOU HAVE ALREADY CREATED THE INTERCLUSTER LIFS)
create_source_intercluster_lifs: "yes"

source_intercluster_network_port_details:
  using_dedicated_ports: "yes"
  using_ifgrp: "yes"
  using_vlans: "yes"
  failover_for_shared_individual_ports: "yes"
  ifgrp_name: "a0a"
  vlan_id: "10"
  ports:
    - "e0b"
    - "e0g"
  broadcast_domain: "NFS"
  ipspace: "Default"
  failover_group_name: "iclifs"

source_intercluster_lif_details:
  - name: "icl_1"
    address: "10.0.0.1"
    netmask: "255.255.255.0"
    home_port: "a0a-10"
    node: "AFF-01"
  - name: "icl_2"
    address: "10.0.0.2"
    netmask: "255.255.255.0"
    home_port: "a0a-10"
    node: "AFF-02"

#####
### CVO Deployment Variables ###
#####

##### Access Keys Variables #####

# Region where your CVO will be deployed.
region_deploy: "us-east-1"

##### CVO and Connector Vars #####

# AWS Managed Policy required to give permission for IAM role creation.

```



```

aws_policy: "arn:aws:iam::1234567:policy/OCCM"

# Specify your aws role name, a new role is created if one already does
not exist.
aws_role_name: "arn:aws:iam::1234567:policy/OCCM"

# Name your connector.
connector_name: "awx_connector"

# Name of the key pair generated in AWS.
key_pair: "key_pair"

# Name of the Subnet that has the range of IP addresses in your VPC.
subnet: "subnet-12345"

# ID of your AWS security group that allows access to on-prem
resources.
security_group: "sg-123123123"

# Your Cloud Manager Account ID.
account: "account-A23123A"

# Name of the your CVO instance
cvo_name: "test_cvo"

# ID of the VPC in AWS.
vpc: "vpc-123123123"

#####
#####
# Variables for - Add on-prem ONTAP to Connector in Cloud Manager
#####
#####

# For Federated users, Client ID from API Authentication Section of
Cloud Central to generate access token.
sso_id: "123123123123123123123"

# For regular access with username and password, please specify "pass"
as the connector_access. For SSO users, use "refresh_token" as the
variable.
connector_access: "pass"

#####
#####
# Variables for SnapMirror Peering
#####

```

```

#####
passphrase: "your-passphrase"

#####
#####
# Source & Destination List
#####
#####
#Please Enter Destination Cluster Name
dst_cluster_name: "dst-cluster-name"

#Please Enter Destination Cluster (Once CVO is Created Add this
Variable to all templates)
dst_cluster_ip: "dst-cluster-ip"

#Please Enter Destination SVM to create mirror relationship
dst_vserver: "dst-vserver"

#Please Enter NFS Lif for dst vserver (Once CVO is Created Add this
Variable to all templates)
dst_nfs_lif: "dst-nfs-lif"

#Please Enter Source Cluster Name
src_cluster_name: "src-cluster-name"

#Please Enter Source Cluster
src_cluster_ip: "src-cluster-ip"

#Please Enter Source SVM
src_vserver: "src-vserver"

#####
#####
# Variable for Oracle Volumes and SnapMirror Details
#####
#####
#Please Enter Source Snapshot Prefix Name
cg_snapshot_name_prefix: "oracle"

#Please Enter Source Oracle Binary Volume(s)
src_orabinary_vols:
  - "binary_vol"
#Please Enter Source Database Volume(s)
src_db_vols:
  - "db_vol"
#Please Enter Source Archive Volume(s)

```

```

src_archivelog_vols:
  - "log_vol"
#Please Enter Destination Snapmirror Policy
snapmirror_policy: "async_policy_oracle"

#####
#####
# Export Policy Details
#####
#####
#Enter the destination export policy details (Once CVO is Created Add
this Variable to all templates)
export_policy_details:
  name: "nfs_export_policy"
  client_match: "0.0.0.0/0"
  ro_rule: "sys"
  rw_rule: "sys"

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

#NFS Mount points for Oracle DB volumes
mount_points:
  - "/u01"
  - "/u02"
  - "/u03"

# Up to 75% of node memory size divided by 2mb. Consider how many
databases to be hosted on the node and how much ram to be allocated to
each DB.
# Leave it blank if hugepage is not configured on the host.
hugepages_nr: "1234"

# RedHat subscription username and password
redhat_sub_username: "xxx"
redhat_sub_password: "xxx"

#####
### DB env specific install and config variables ###
#####
#Recovery Type (leave as scn)
recovery_type: "scn"

```

```
#Oracle Control Files
control_files:
  - "/u02/oradata/CDB2/control01.ctl"
  - "/u03/orareco/CDB2/control02.ctl"
```

Automation Playbooks

There are four separate playbooks that need to be ran.

1. Playbook for Setting up your environment, On-Prem or CVO.
2. Playbook for replicating Oracle Binaries and Databases on a schedule
3. Playbook for replicating Oracle Logs on a schedule
4. Playbook for Recovering your database on a destination host

ONTAP/CVO Setup

ONTAP and CVO Setup

Configure and launch the job template.

1. Create the job template.
 - a. Navigate to Resources → Templates → Add and click Add Job Template.
 - b. Enter the name ONTAP/CVO Setup
 - c. Select the Job type; Run configures the system based on a playbook.
 - d. Select the corresponding inventory, project, playbook, and credentials for the playbook.
 - e. Select the `ontap_setup.yml` playbook for an On-Prem environment or select the `cvo_setup.yml` for replicating to a CVO instance.
 - f. Paste global variables copied from step 4 into the Template Variables field under the YAML tab.
 - g. Click Save.
2. Launch the job template.
 - a. Navigate to Resources → Templates.
 - b. Click the desired template and then click Launch.



We will use this template and copy it out for the other playbooks.

Replication For Binary and Database Volumes

Scheduling the Binary and Database Replication Playbook

Configure and launch the job template.

1. Copy the previously created job template.
 - a. Navigate to Resources → Templates.
 - b. Find the ONTAP/CVO Setup Template, and on the far right click on Copy Template
 - c. Click Edit Template on the copied template, and change the name to Binary and Database Replication Playbook.
 - d. Keep the same inventory, project, credentials for the template.
 - e. Select the `ora_replication_cg.yml` as the playbook to be executed.
 - f. The variables will remain the same, but the CVO cluster IP will need to be set in the variable `dst_cluster_ip`.
 - g. Click Save.
2. Schedule the job template.
 - a. Navigate to Resources → Templates.
 - b. Click the Binary and Database Replication Playbook template and then click Schedules at the top set of options.
 - c. Click Add, add Name Schedule for Binary and Database Replication, choose the Start date/time at the beginning of the hour, choose your Local time zone, and Run frequency. Run frequency will be often the SnapMirror replication will be updated.



A separate schedule will be created for the Log volume replication, so that it can be replicated on a more frequent cadence.

Replication for Log Volumes

Scheduling the Log Replication Playbook

Configure and launch the job template.

1. Copy the previously created job template.
 - a. Navigate to Resources → Templates.
 - b. Find the ONTAP/CVO Setup Template, and on the far right click on Copy Template
 - c. Click Edit Template on the copied template, and change the name to Log Replication Playbook.
 - d. Keep the same inventory, project, credentials for the template.
 - e. Select the ora_replication_logs.yml as the playbook to be executed.
 - f. The variables will remain the same, but the CVO cluster IP will need to be set in the variable `dst_cluster_ip`.
 - g. Click Save.
2. Schedule the job template.
 - a. Navigate to Resources → Templates.
 - b. Click the Log Replication Playbook template and then click Schedules at the top set of options.
 - c. Click Add, add Name Schedule for Log Replication, choose the Start date/time at the beginning of the hour, choose your Local time zone, and Run frequency. Run frequency will be often the SnapMirror replication will be updated.



It is recommended to set the log schedule to update every hour to ensure the recovery to the last hourly update.

Restore and Recover Database

Scheduling the Log Replication Playbook

Configure and launch the job template.

1. Copy the previously created job template.
 - a. Navigate to Resources → Templates.
 - b. Find the ONTAP/CVO Setup Template, and on the far right click on Copy Template
 - c. Click Edit Template on the copied template, and change the name to Restore and Recovery Playbook.
 - d. Keep the same inventory, project, credentials for the template.
 - e. Select the ora_recovery.yml as the playbook to be executed.
 - f. The variables will remain the same, but the CVO cluster IP will need to be set in the variable `dst_cluster_ip`.
 - g. Click Save.



This playbook will not be ran until you are ready to restore your database at the remote site.

Recovering Oracle Database

1. On-premises production Oracle databases data volumes are protected via NetApp SnapMirror replication to either a redundant ONTAP cluster in secondary data center or Cloud Volume ONTAP in public cloud. In a fully configured disaster recovery environment, recovery compute instances in secondary data center or public cloud are standby and ready to recover the production database in the case of a disaster. The standby compute instances are kept in sync with on-prem instances by running parallel updates on OS kernel patch or upgrade in a lockstep.
2. In this solution demonstrated, Oracle binary volume is replicated to target and mounted at target instance to bring up Oracle software stack. This approach to recover Oracle has advantage over a fresh installation of Oracle at last minute when a disaster occurred. It guarantees Oracle installation is fully in sync with current on-prem production software installation and patch levels etc. However, this may or may not have additional software licensing implication for the replicated Oracle binary volume at recovery site depending on how the software licensing is structured with Oracle. User is recommended to check with its software licensing personnel to assess the potential Oracle licensing requirement before deciding to use the same approach.
3. The standby Oracle host at the destination is configured with the Oracle prerequisite configurations.
4. The SnapMirrors are broken and the volumes are made writable and mounted to the standby Oracle host.
5. The Oracle recovery module performs following tasks to recovery and startup Oracle at recovery site after all DB volumes are mounted at standby compute instance.
 - a. Sync the control file: We deployed duplicate Oracle control files on different database volume to protect critical database control file. One is on the data volume and another is on log volume. Since data and log volumes are replicated at different frequency, they will be out of sync at the time of recovery.
 - b. Relink Oracle binary: Since the Oracle binary is relocated to a new host, it needs a relink.
 - c. Recover Oracle database: The recovery mechanism retrieves last System Change Number in last available archived log in Oracle log volume from control file and recovers Oracle database to recoup all business transactions that was able to be replicated to DR site at the time of failure. The database is then started up in a new incarnation to carry on user connections and business transaction at recovery site.



Before running the Recovering playbook make sure you have the following:
Make sure it copy over the /etc/oratab and /etc/oralnst.loc from the source Oracle host to the destination host

TR-4794: Oracle databases on NetApp EF-Series

Mitch Blackburn, Ebin Kadavy, NetApp

TR-4794 is intended to help storage administrators and database administrators successfully deploy Oracle on NetApp EF-Series storage.

[TR-4794: Oracle databases on NetApp EF-Series](#)

Microsoft SQL Server

TR-4951: Backup and Recovery for Microsoft SQL Server on AWS FSx for ONTAP

Author(s): Niyaz Mohammed, Carine Ngwekwe - NetApp Solutions Engineering

This document covers the steps necessary to perform backup and recovery for Microsoft SQL Server on AWS FSx for ONTAP with SnapCenter. This includes the following information:

- NetApp SnapCenter configuration
- SnapCenter backup operations
- Backup operation for an FCI database
- Backup operation for multiple databases
- Restore and recovery


SnapCenter Configuration

The following steps must be performed for SnapCenter configuration and the protection of Microsoft SQL Server resources. Each of the following steps is detailed in the following sections.

1. Configure sysadmin credentials for the SQL Server backup and restore user.
2. Configure storage settings. Provide Amazon Web Services (AWS) management credential to access the Amazon FSx for NetApp ONTAP storage virtual machines (SVMs) from SnapCenter.
3. Add a SQL Server host to SnapCenter. Deploy and install the required SnapCenter Plug-ins.
4. Configure policies. Define the backup operation type, retention, and optional Snapshot backup replication.
5. Configure and protect the Microsoft SQL Server database.

SnapCenter newly installed user interface

Configure credentials for SQL Server backup and restore the user with sysadmin rights.

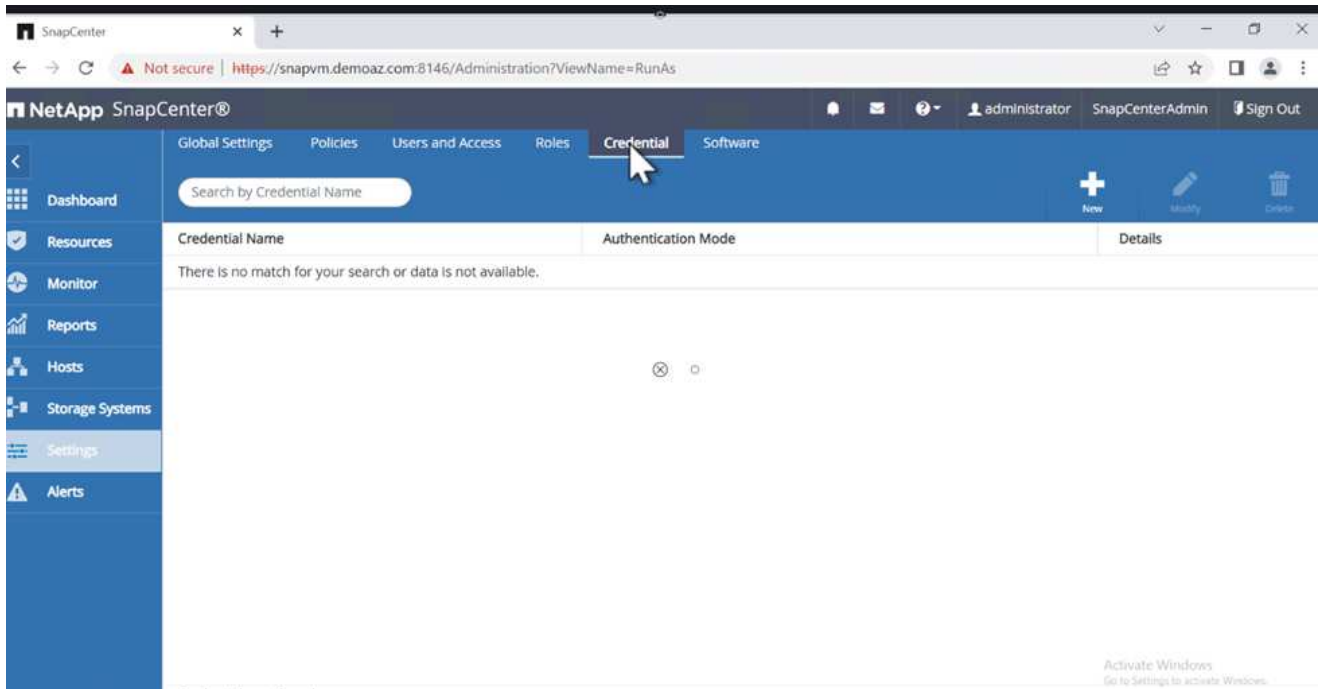


NetApp recommends using role-based access control (RBAC) to delegate data protection and management capabilities to individual users across the SnapCenter and window hosts. The user must have access to the SQL Server hosting the database. For multiple hosts, the username and password must be the same across the various hosts. Furthermore, to enable SnapCenter to deploy the required plug-in on SQL Server hosts, you must register the domain information for SnapCenter to validate your credentials and hosts.

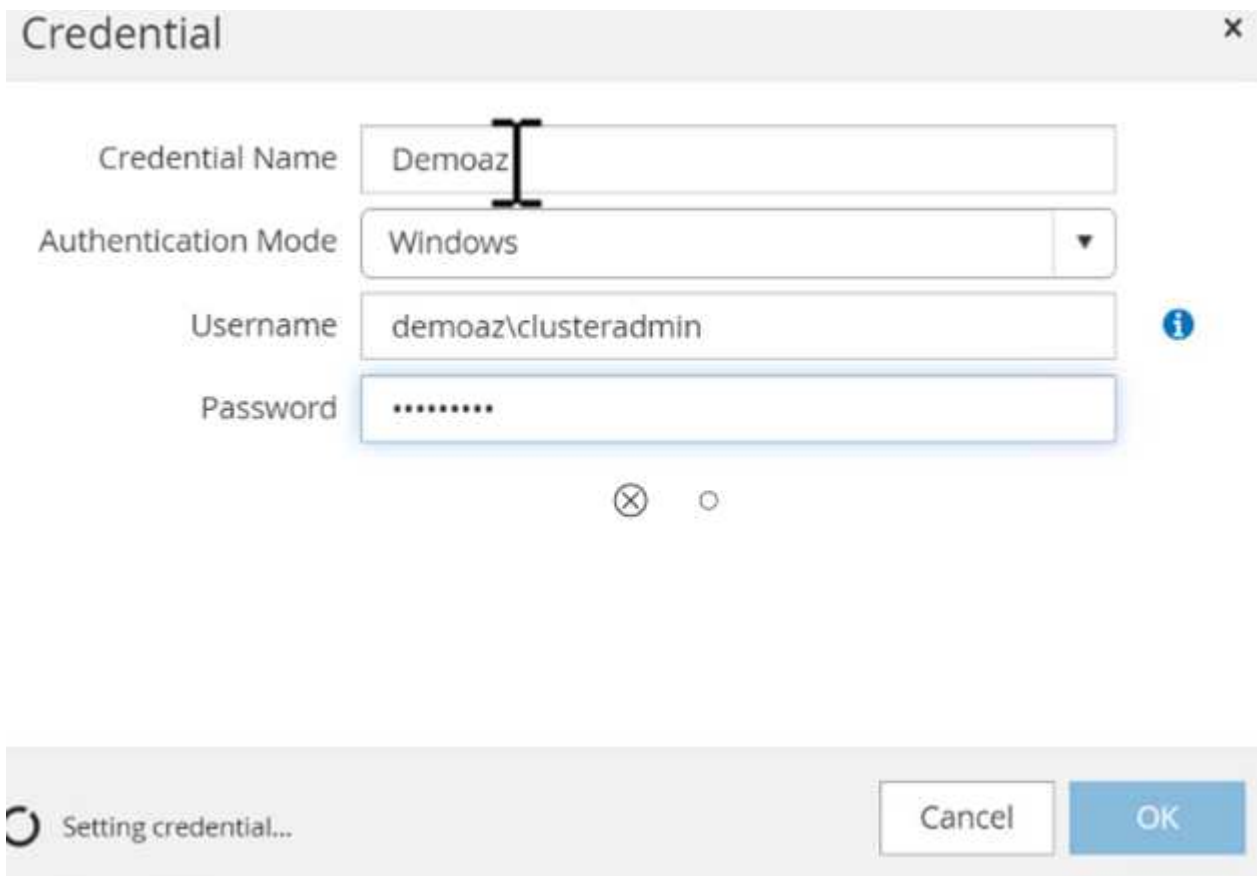
Expand the following sections to see the detailed instructions on how to complete each step.

Add the credentials

Go to **Settings**, select **Credentials**, and click (+).




The new user must have administrator rights on the SQL Server host.

A screenshot of the 'Credential' dialog box in NetApp SnapCenter. The dialog has a title bar 'Credential' with a close button. It contains four input fields: 'Credential Name' with the value 'Demoaz', 'Authentication Mode' with a dropdown menu set to 'Windows', 'Username' with the value 'demoaz\clusteradmin', and 'Password' with a masked field of seven dots. There is an information icon (i) to the right of the Username field. At the bottom, there is a progress indicator 'Setting credential...' with a circular arrow, and two buttons: 'Cancel' and 'OK'.

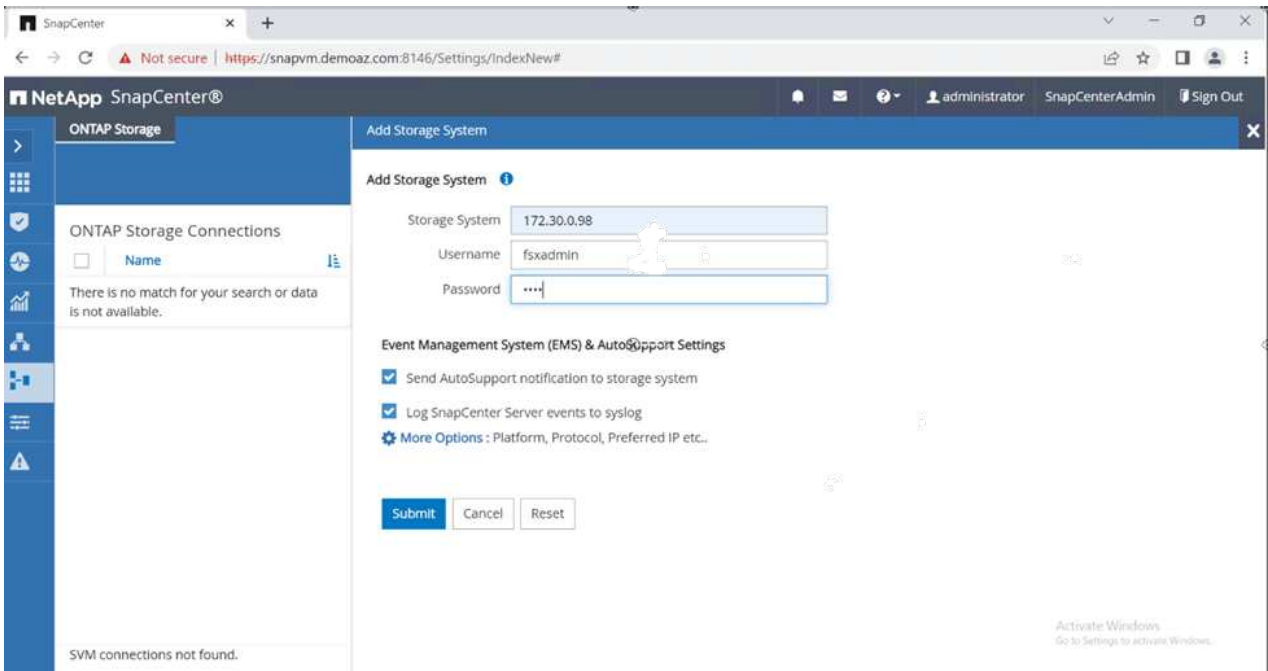
Configure storage

To configure storage in SnapCenter, complete the following steps:

1. In the SnapCenter UI, select **Storage Systems**. There are two storage types, **ONTAP SVM** and **ONTAP Cluster**. By default, the storage type is **ONTAP SVM**.
2. Click (+) to add the storage system information.



3. Provide the **FSx for ONTAP management** endpoint.



4. The SVM is now configured in SnapCenter.

NetApp SnapCenter®

ONTAP Storage

Type: Search by Name

ONTAP Storage Connections

<input type="checkbox"/>	Name	IP	Cluster Name	User Name	Platform	Controller License
<input type="checkbox"/>	ESNSVMTESTRDS		rdsfsxTest01		FSx	Not applicable


Total 1

Activate Windows
Go to Settings to activate Windows.

Add a SQL Server host to SnapCenter

To add a SQL Server host, complete the following steps:

1. From the Host tab, click (+) to add the Microsoft SQL Server host.




2. Provide the fully qualified domain name (FQDN) or IP address of the remote host.




The credentials are populated by default.

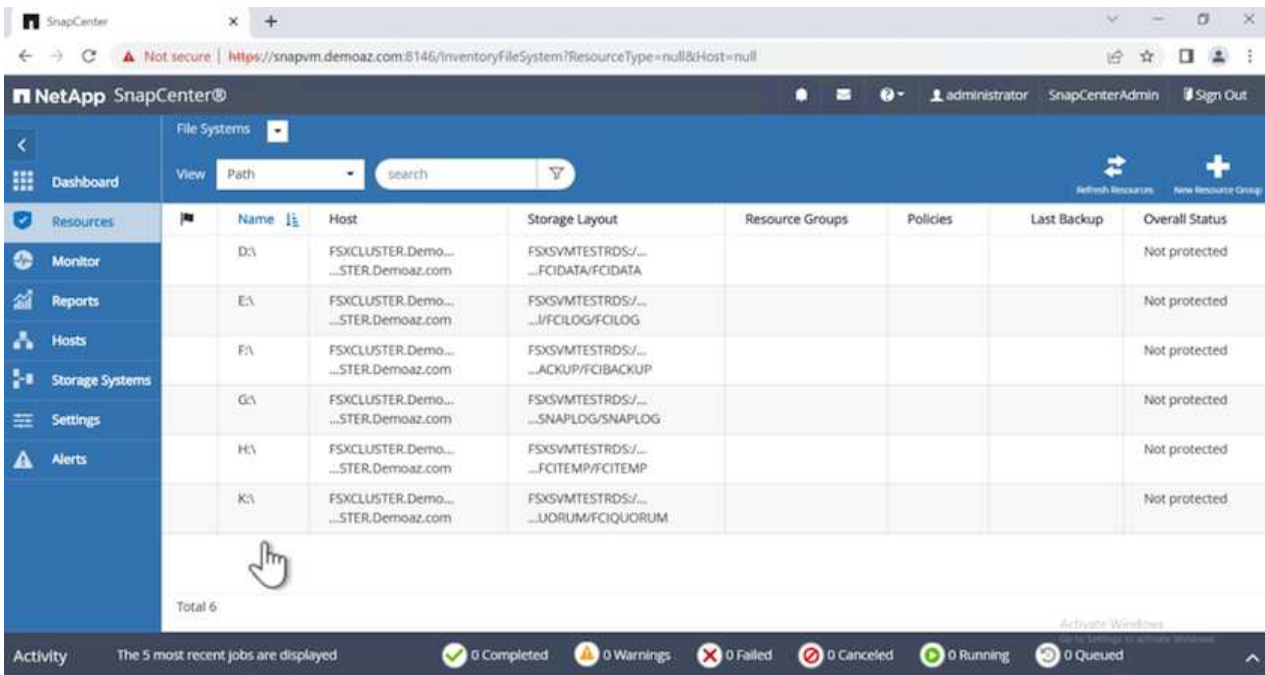
3. Select the option for Microsoft windows and Microsoft SQL Server and then submit.



The SQL Server packages are installed.



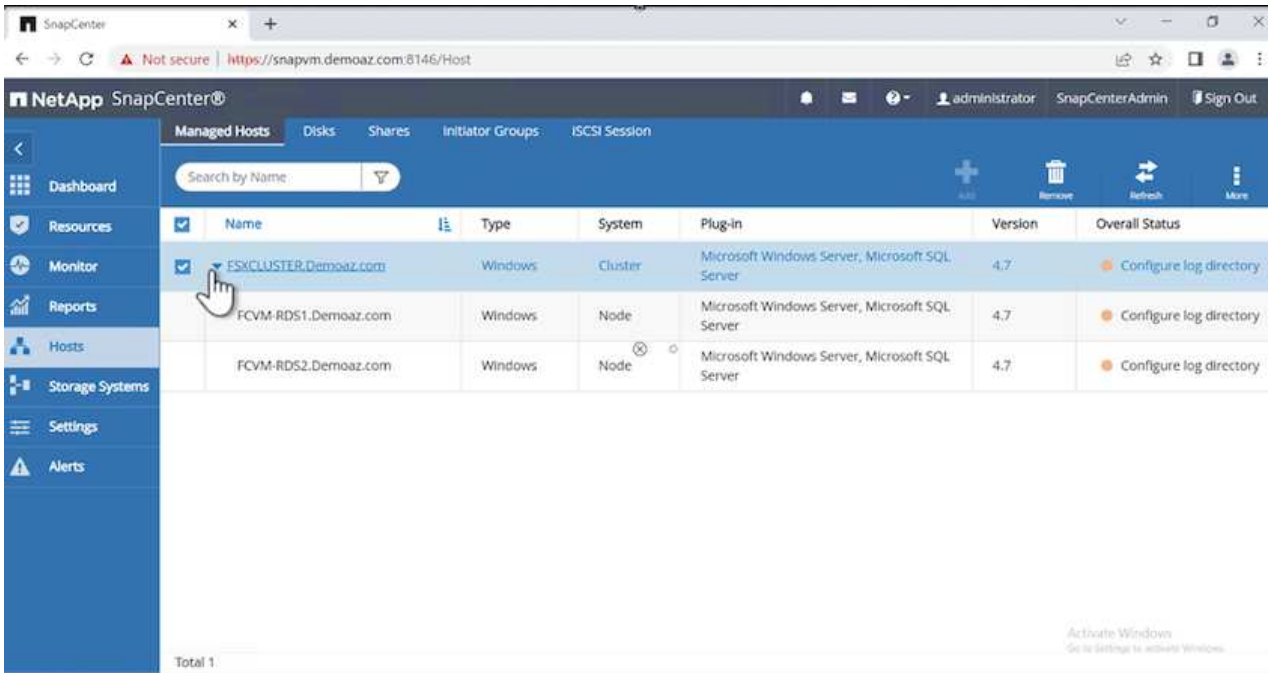
1. After the installation is complete, go to the **Resource** tab to verify whether all FSx for ONTAP iSCSI volumes are present.




Configure log directory

To configure a host log directory, complete the following steps:

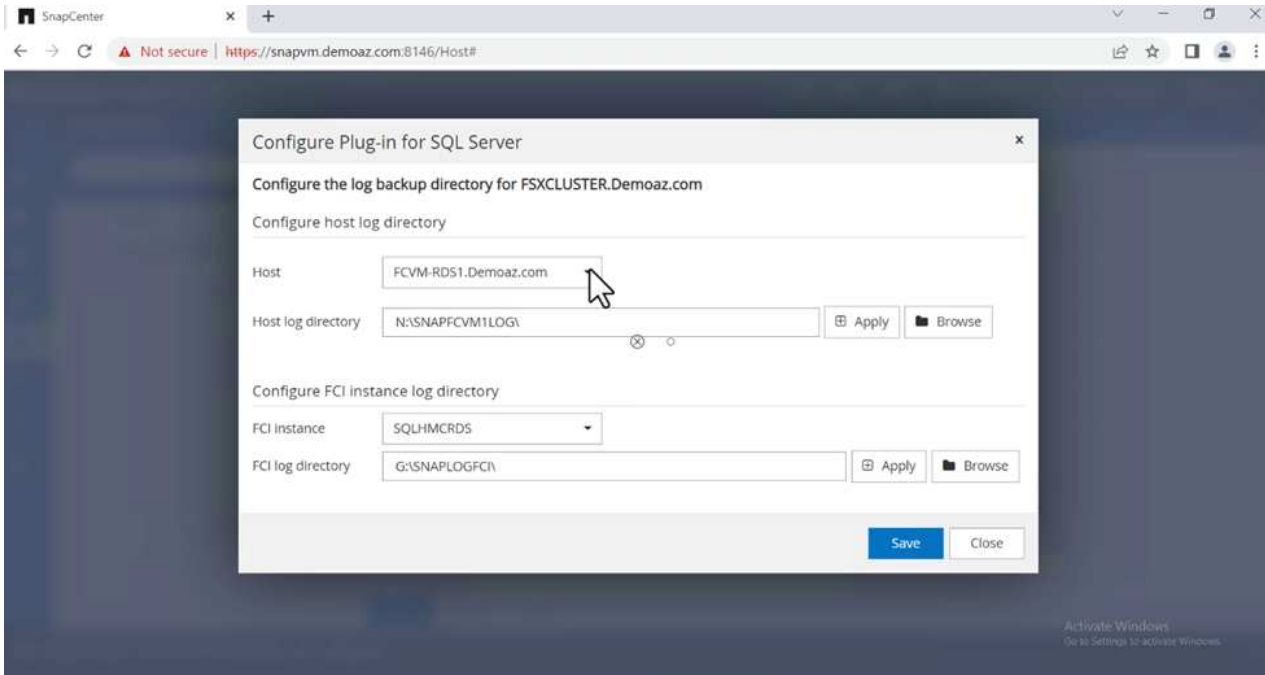
1. Click the check box. A new tab opens.



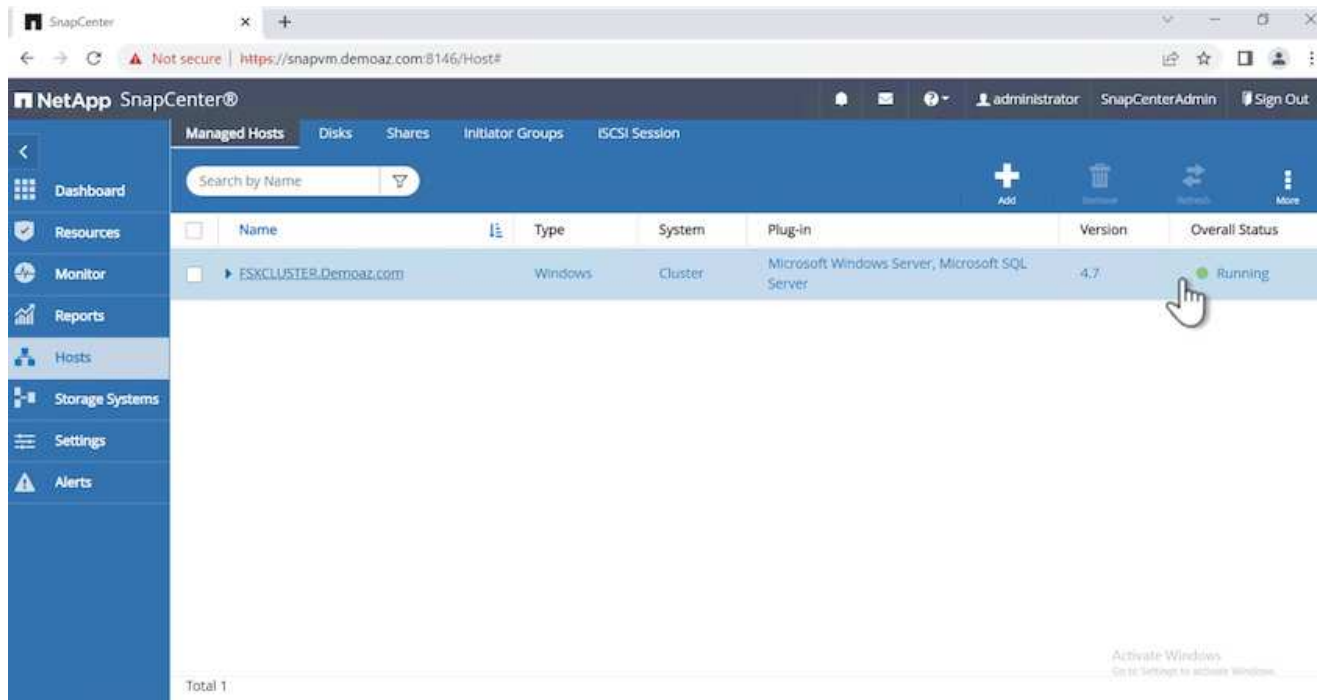
2. Click the **configure log directory** link.



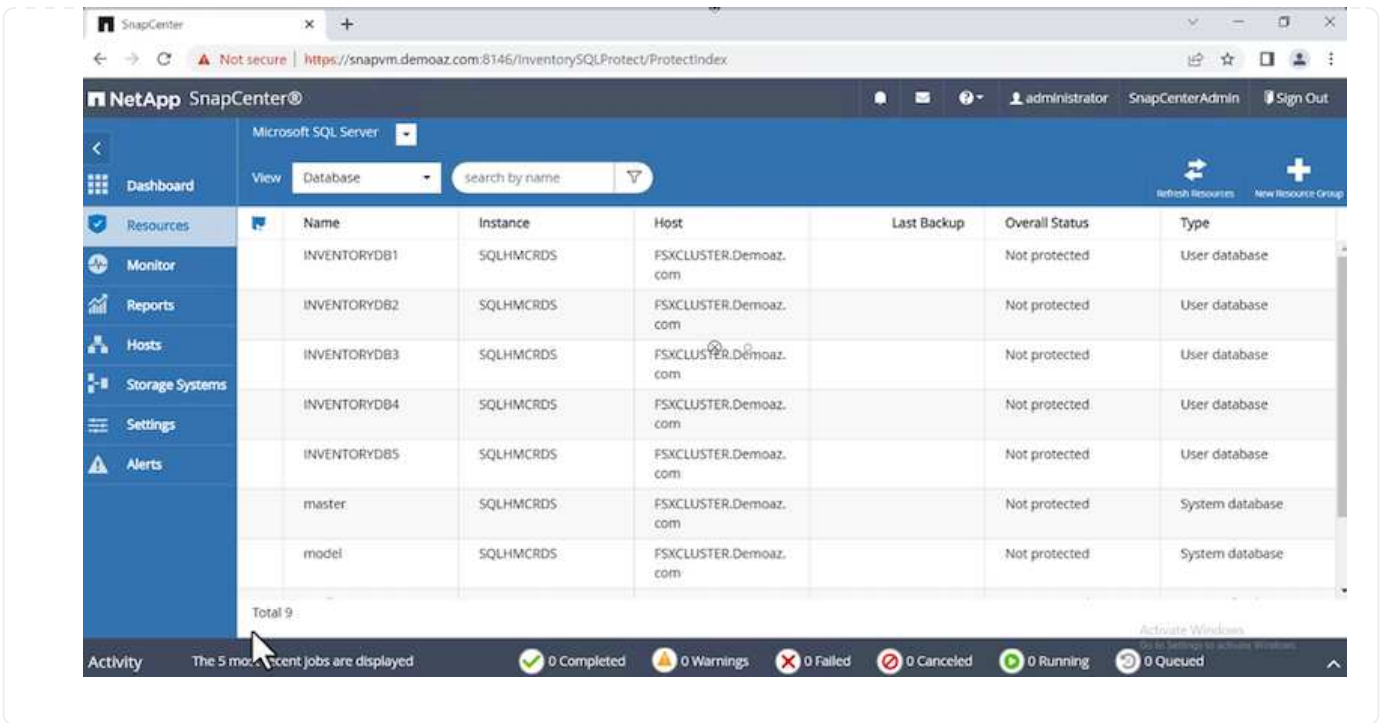
3. Select the drive for the host log directory and the FCI instance log directory. Click **Save**. Repeat the same process for the second node in the cluster. Close the window.



The host is now in a running state.



1. From the **Resources** tab, we have all the servers and databases.



Configure a backup policy


A backup policy is a set of rules that govern how to manage, schedule, and retain backup. It helps with the backup type and frequency based on your company's SLA.

Expand the following sections to see the detailed instructions on how to complete each step.

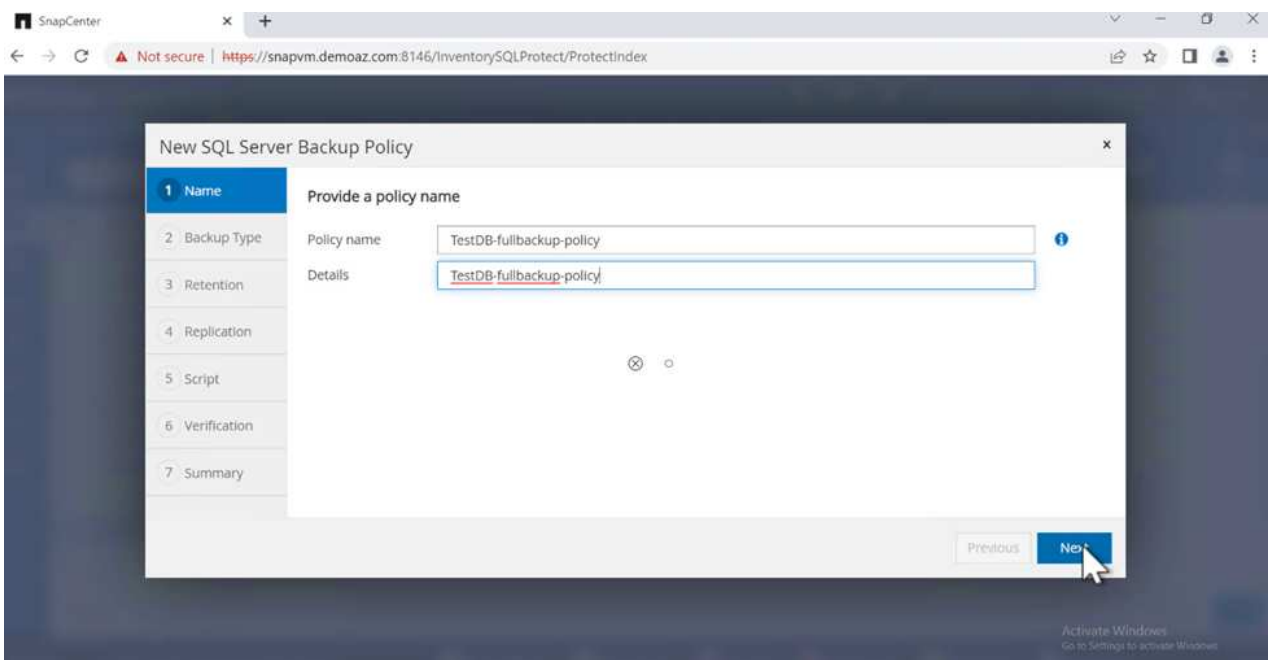
Configure back-up operation for an FCI database

To configure a backup policy for an FCI database, complete the following steps:

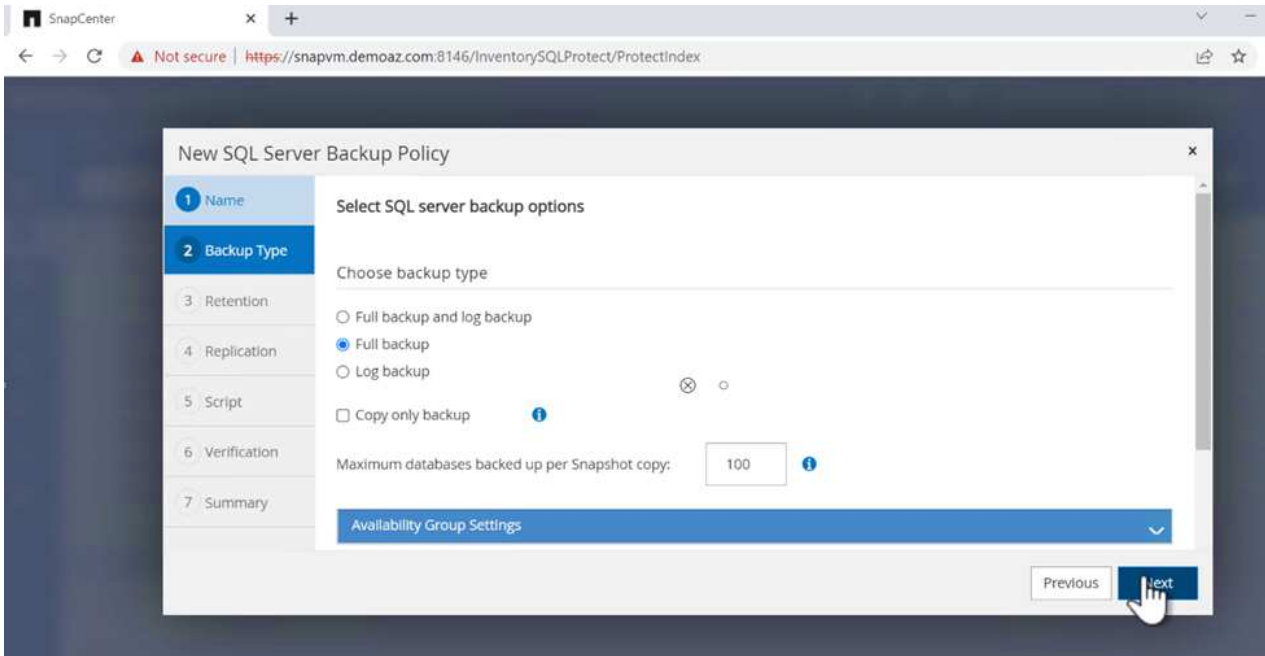
1. Go to **Settings** and select **Policies** on the top left. Then click **New**.




2. Enter the policy name and a description. Click **Next**.



3. Select **Full backup** as the backup type.



4. Select the schedule frequency (this is based on the company SLA). Click **Next**.



5. Configure the retention settings for the backup.

New SQL Server Backup Policy x

- 1 Name
- 2 Backup Type
- 3 Retention**
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Retention settings

Retention settings for up-to-the-minute restore operation ⓘ

Keep log backups applicable to last full backups

Keep log backups applicable to last days

Full backup retention settings ⓘ

Weekly

Total Snapshot copies to keep

Keep Snapshot copies for days

6. Configure the replication options.

New SQL Server Backup Policy ×

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select secondary replication options ⓘ

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label ⓘ

Error retry count ⓘ

7. Specify a run script to run before and after a backup job is run (if any).

New SQL Server Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Specify optional scripts to run before performing a backup job

Prescript full path

Prescript arguments

Specify optional scripts to run after performing a backup job

Postscript full path

Postscript arguments

Script timeout secs

Previous Next

8. Run verification based on the backup schedule.

New SQL Server Backup Policy

- Name
- Backup Type
- Retention
- Replication
- Script
- Verification**
- Summary

Select the options to run backup verification

Run verifications for the following backup schedules

Select how often you want the schedules to occur in the policy. The specific verification times are set at backup job creation enabling you to stagger your verification start times.

Weekly

Database consistency checks options

Limit the integrity structure to physical structure of the database (PHYSICAL_ONLY)

Suppress all information message (NO_INFOMSGS)

Display all reported error messages per object (ALL_ERRORMSG)

Do not check non-clustered indexes (NOINDEX)

Limit the checks and obtain the locks instead of using an internal database Snapshot copy (TABLOCK)

Verification script settings

Script timeout: 60 secs

Prescript full path:

Prescript arguments: Choose optional arguments...

Postscript full path:

Postscript arguments: Choose optional arguments...

Previous **Next**

9. The **Summary** page provides details of the backup policy. Any errors can be corrected here.

SnapCenter

Not secure | https://snapvm.demoaz.com:8146/inventorySQLProtect/ProtectIndex

New SQL Server Backup Policy

- Name
- Backup Type
- Retention
- Replication
- Script
- Verification
- Summary**

Summary

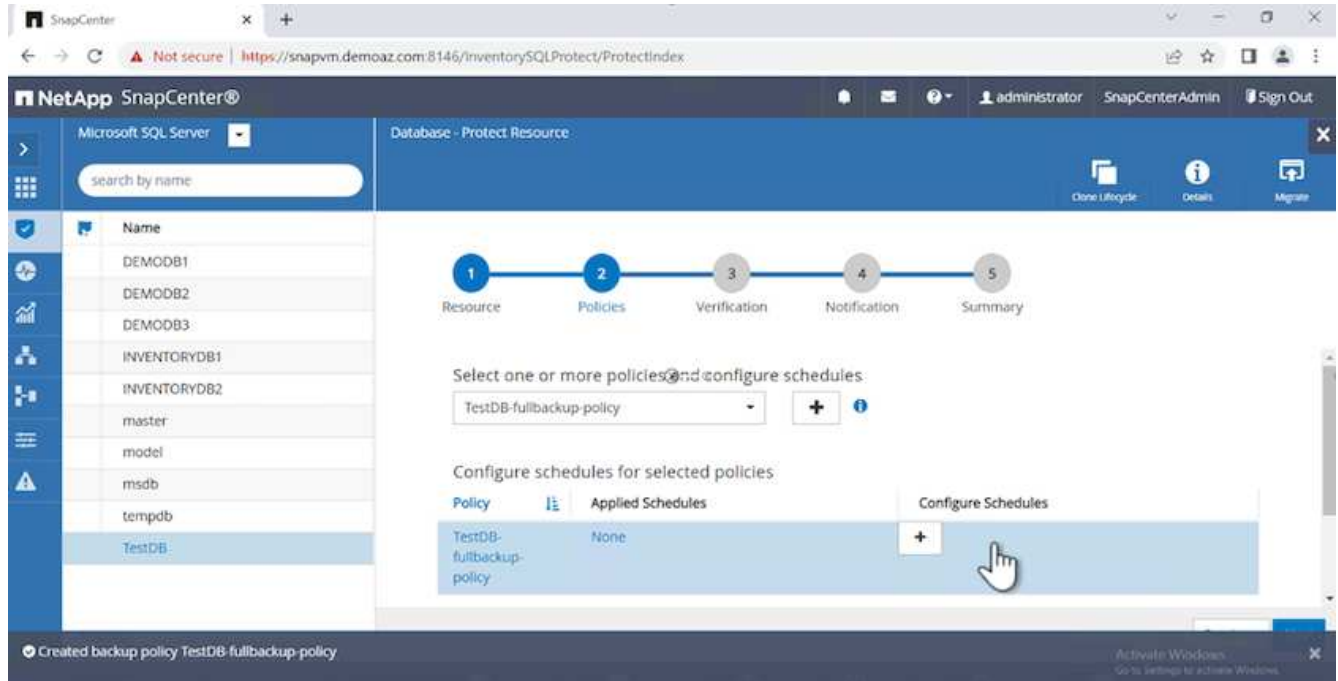
Policy name	TestDB-fullbackup-policy
Details	TestDB-fullbackup-policy
Backup type	Full backup
Availability group settings	Backup only on preferred backup replica
Schedule Type	Hourly <input checked="" type="radio"/> <input type="radio"/>
UTM retention	Total backup copies to retain : 7
Hourly Full backup retention	Total backup copies to retain : 7
Replication	none
Backup prescript settings	undefined
Prescript arguments:	

Previous **Finish**

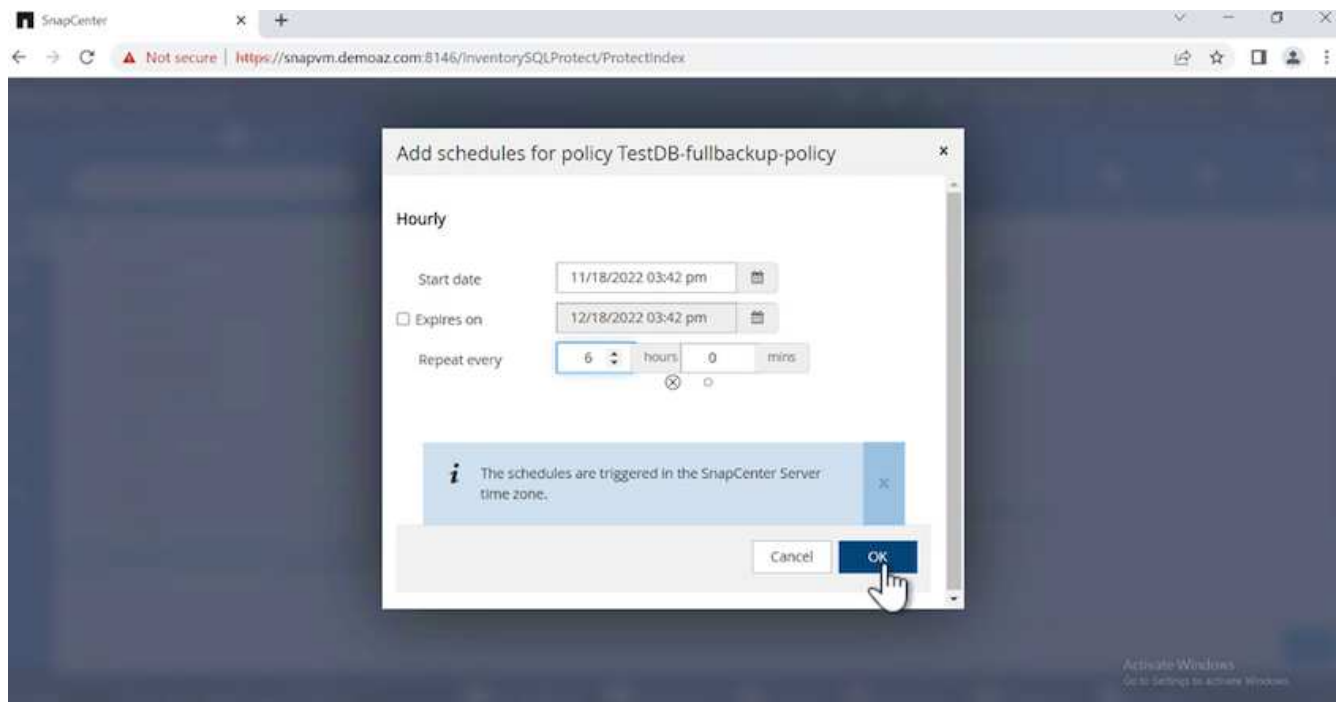
Activate Windows
Go to Settings to activate Windows

Configure and protect MSSQL Server database


1. Set up the starting date and expiration date of the backup policy.




2. Define the schedule for the backup. To do that, click (+) to configure a schedule. Enter the **Start date** and **Expires on** date. Set the time based on the company's SLA.




3. Configure the verification server. From the drop-down menu, select the server.



4. Confirm the configured schedule by clicking the plus sign and confirm.
5. Provide information for email notification. Click **Next**.




The summary of the backup policy for SQL Server database is now configured.




SnapCenter backup operations

To create on-demand SQL Server backups, complete the following steps:


1. From the **Resource** view, select the resource and select **Backup now**.



2. In the **Backup** dialog box, click **Backup**.

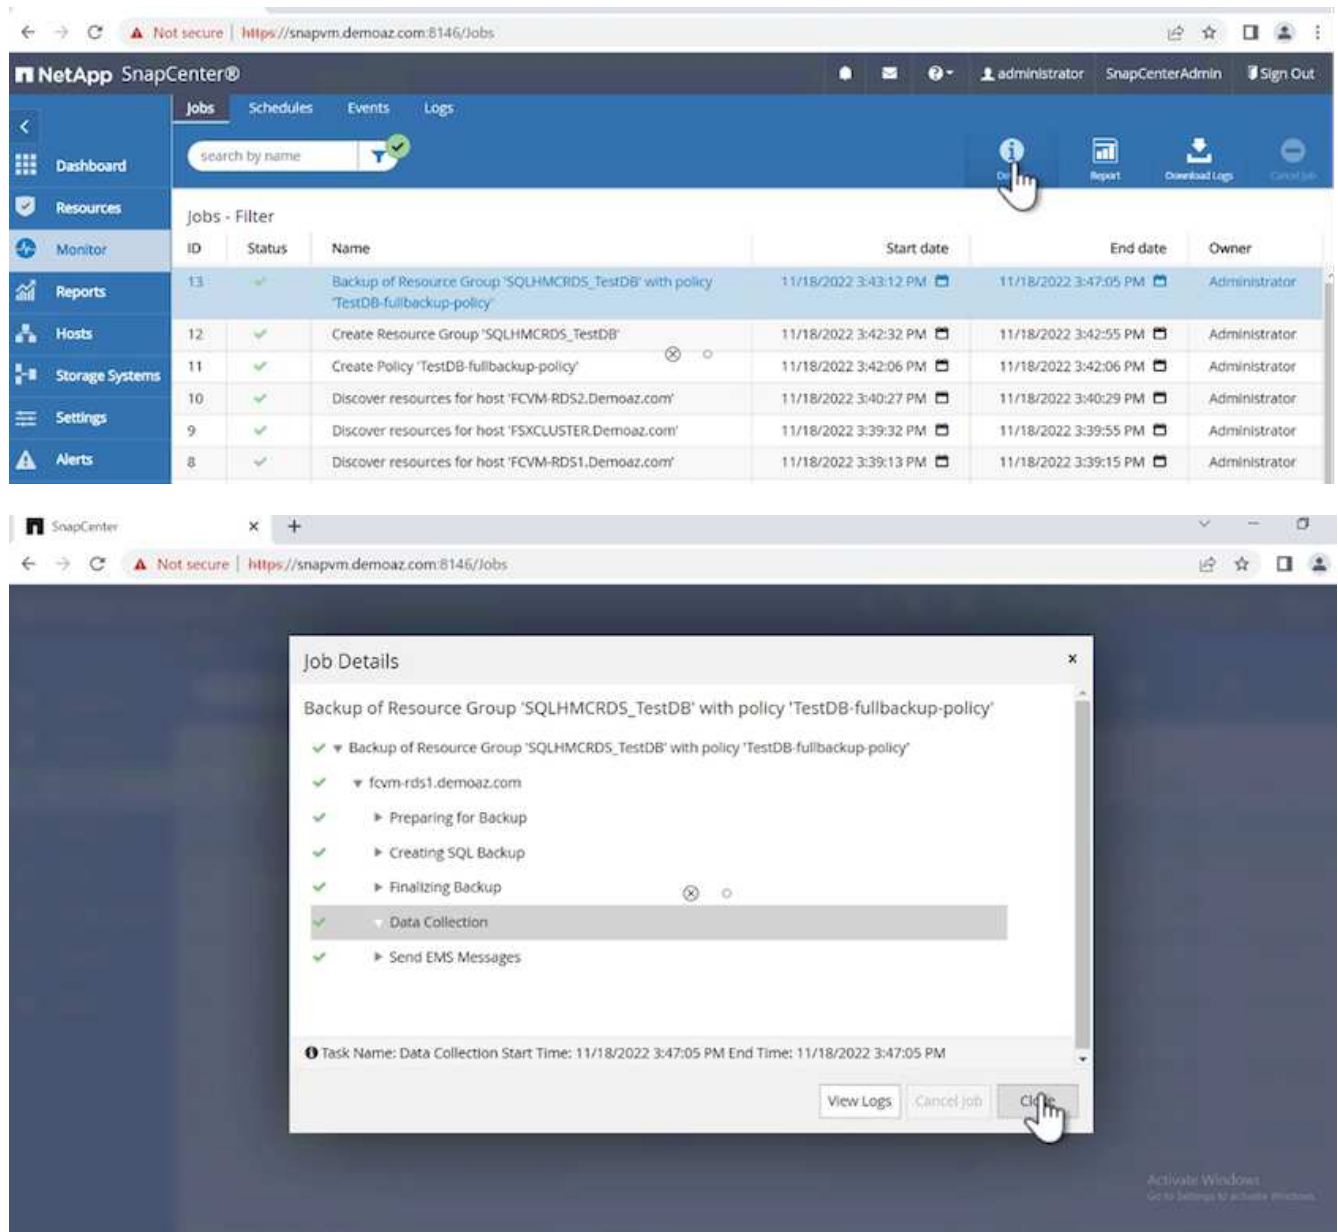


3. A confirmation screen is displayed. Click **Yes** to confirm.



Monitor backup job

1. From the **Monitor** tab, click the job and select **Details** on the right to view the jobs.




When the backup is completed, a new entry is shown in the Topology view.

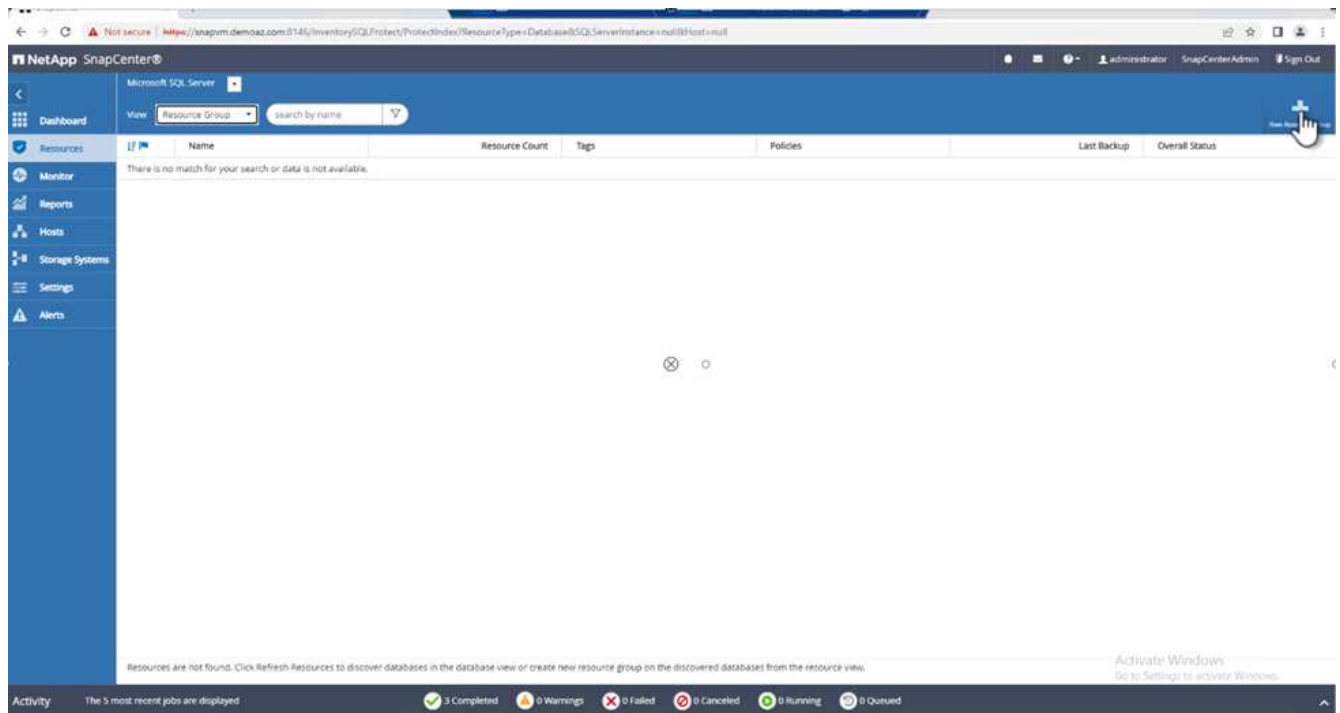
Backup operation for multiple databases

To configure a backup policy for multiple SQL Server databases, create resource group policies by completing the following steps:


1. In the **Resources** tab from the **View** menu, change to a resource group using the drop-down menu.



2. Click (+) for a new resource group.

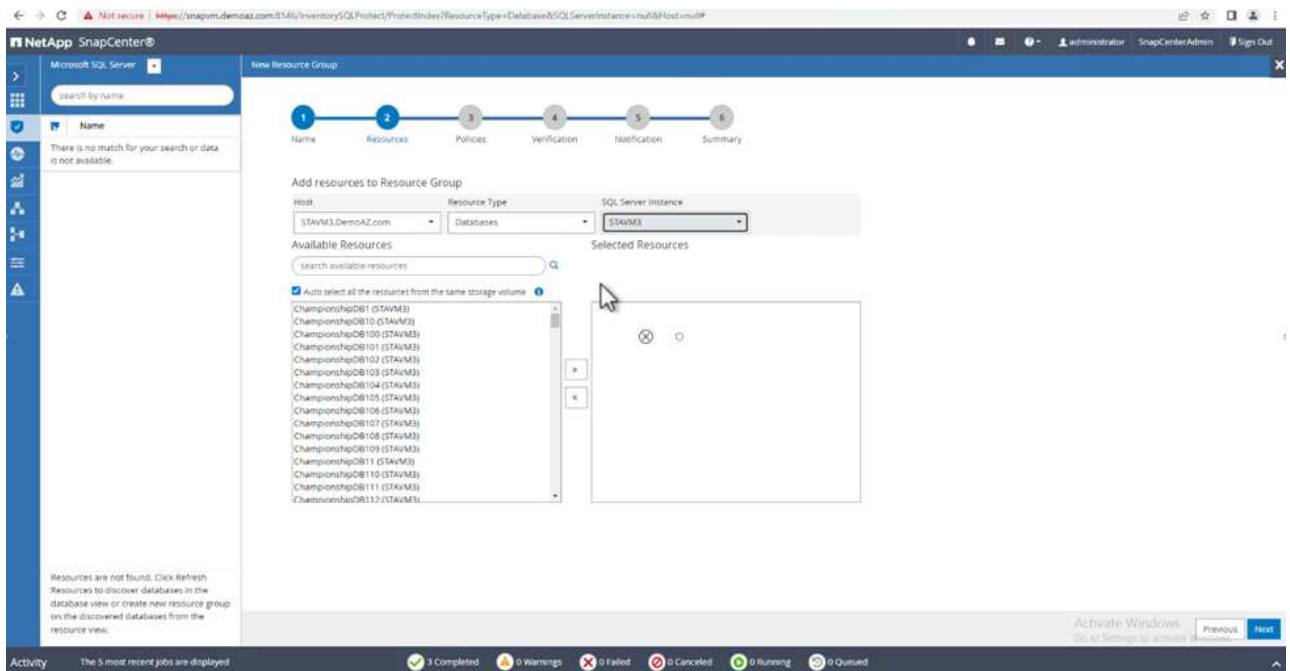


3. Provide a name and tag. Click **Next**.




4. Add resources to the resource group:


- **Host.** Select the server from the drop-down menu hosting the database.
- **Resource type.** From the drop-down menu, select **Database**.
- **SQL Server instance.** Select the server.




The **option** Auto Selects All the Resources from the Same Storage Volume* is selected by default. Clear the option and select only the databases you need to add to the resource group, Click the arrow to add and click **Next**.




5. On the policies, click (+).



6. Enter the resource group policy name.



7. Select **Full backup** and the schedule frequency depending on your company's SLA.



8. Configure the retention settings.

New SQL Server Backup Policy x

- 1 Name
- 2 Backup Type
- 3 Retention**
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Retention settings

Retention settings for up-to-the-minute restore operation ⓘ

Keep log backups applicable to last full backups

Keep log backups applicable to last days

Full backup retention settings ⓘ

Weekly

Total Snapshot copies to keep

Keep Snapshot copies for days

9. Configure the replication options.

New SQL Server Backup Policy x

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select secondary replication options ⓘ

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label ⓘ

Error retry count ⓘ

10. Configure the scripts to run before performing a backup. Click **Next**.

New SQL Server Backup Policy x

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script**
- 6 Verification
- 7 Summary

Specify optional scripts to run before performing a backup job

Prescript full path

Prescript arguments


Specify optional scripts to run after performing a backup job

Postscript full path


Postscript arguments

Script timeout secs

11. Confirm the verification for the following backup schedules.




12. On the **Summary** page, verify the information, and click **Finish**.



Configure and protect multiple SQL Server databases

1. Click the (+) sign to configure the start date and the expire- on date.



2. Set the time.

Add schedules for policy FULLBACKUP_CHAMPIONSHIPDB



Hourly

Start date

11/11/2022 05:30 pm



Expires on

12/11/2022 05:27 pm



Repeat every

1



hours

0

mins




The schedules are triggered in the SnapCenter Server time zone.



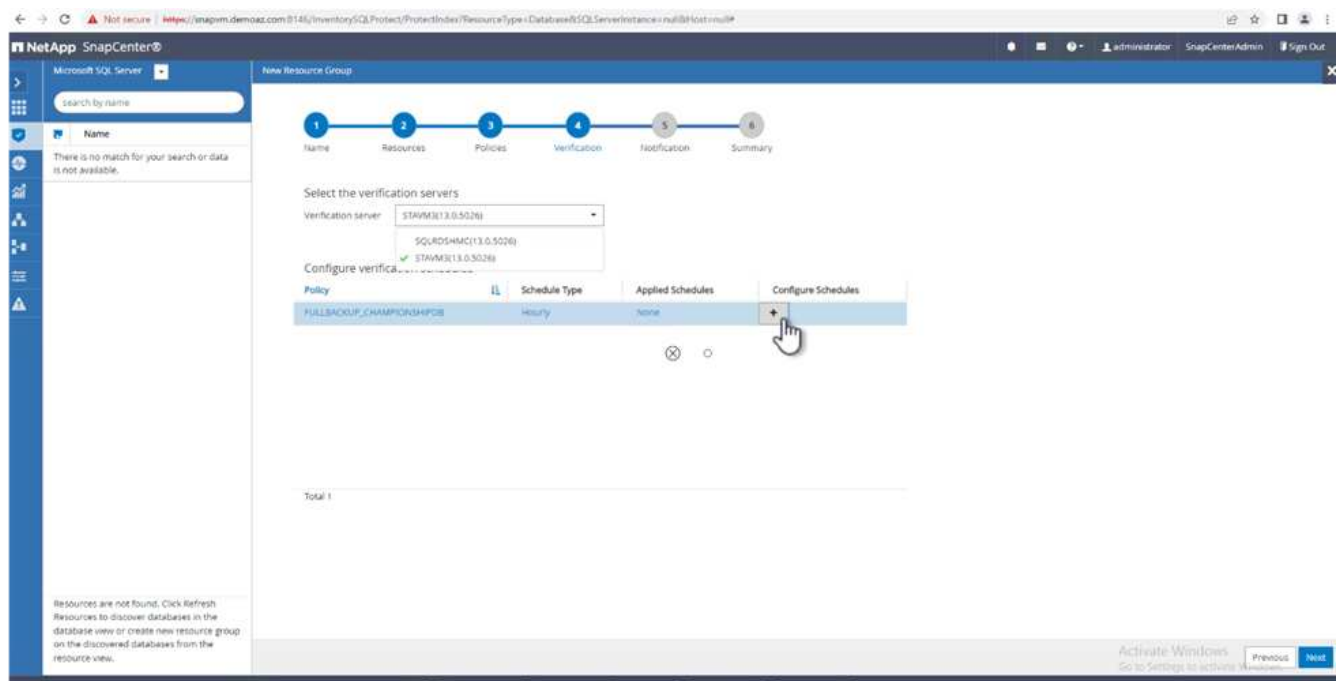
Cancel

OK







3. From the **Verification** tab, select the server, configure the schedule, and click **Next**.



4. Configure notifications to send an email.




The policy is now configured for backing up multiple SQL Server databases.




Trigger on-demand backup for multiple SQL Server databases


1. From the **Resource** tab, select view. From the drop-down menu, select **Resource Group**.




2. Select the resource group name.
3. Click **Backup now** in the upper right.



4. A new window opens. Click the **Verify after backup** checkbox and then click backup.



5. A confirmation message is displayed. Click **Yes**.

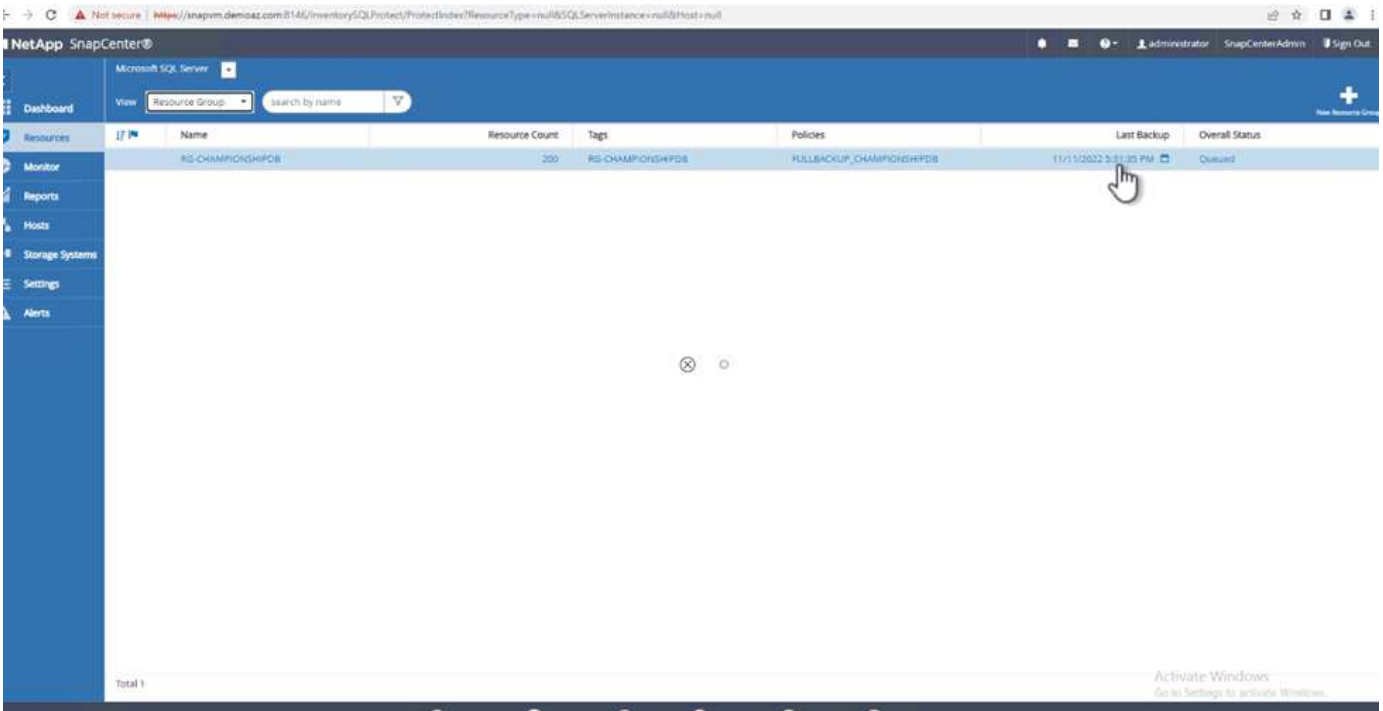


Monitor multiple-database backup jobs

From the left navigation bar, click **Monitor**, select the backup job, and click **Details** to view job progress.



Click the **Resource** tab to see the time it takes for the backup to be completed.




Transaction log backup for multiple database backup

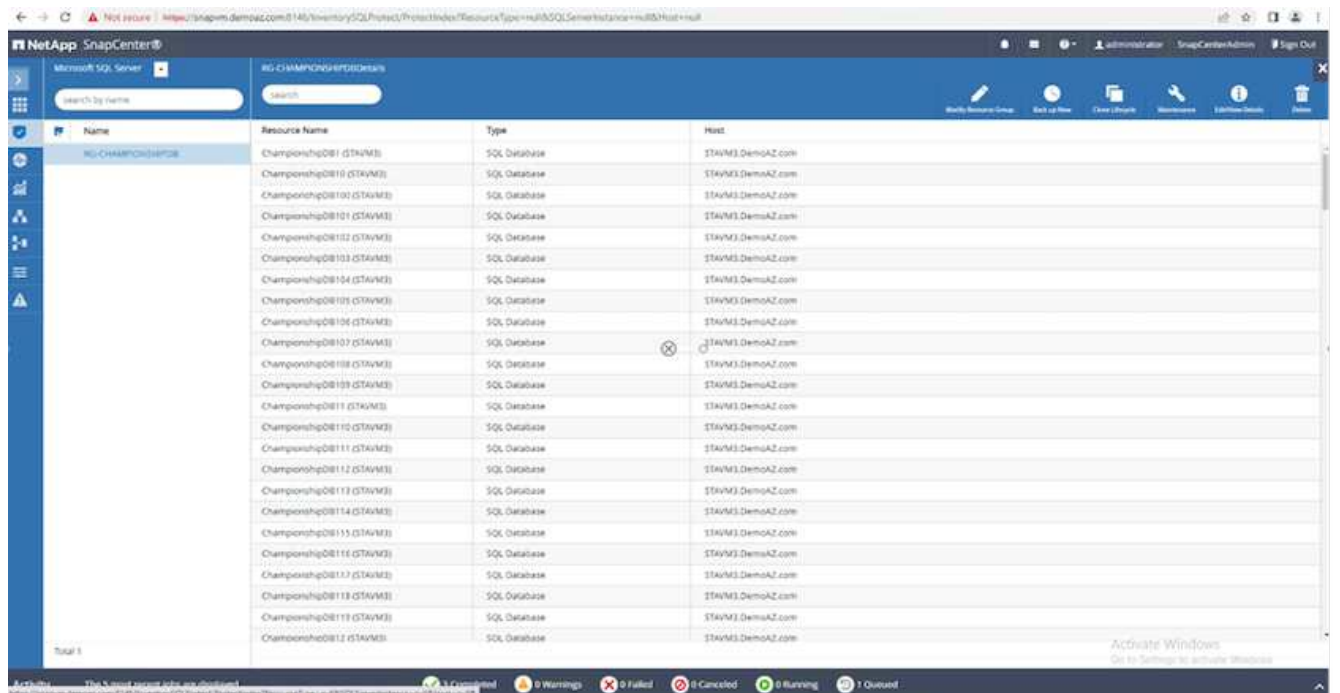
SnapCenter supports the full, bulked logged, and simple recovery models. The simple recovery mode does not support transactional log backup.

To perform a transaction log backup, complete the following steps:


1. From the **Resources** tab, change the view menu from **Database** to **Resource group**.




2. Select the resource group backup policy created.
3. Select **Modify Resource Group** in the upper right.



4. The **Name** section defaults to the backup policy name and tag. Click **Next**.
- The **Resources** tab highlights the bases to which the transaction backup policy is to be configured.



5. Enter the policy name.



6. Select the SQL Server backup options.
7. Select log backup.
8. Set the schedule frequency based on your company's RTO. Click **Next**.

New SQL Server Backup Policy x

- 1 Name
- 2 Backup Type**
- 3 Retention
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Select SQL server backup options

Choose backup type

Full backup and log backup

Full backup

Log backup

Copy only backup i

Maximum databases backed up per Snapshot copy: i

Availability Group Settings v

Schedule frequency

Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling you to stagger your start times.

On demand

Hourly

Daily

Weekly

Monthly

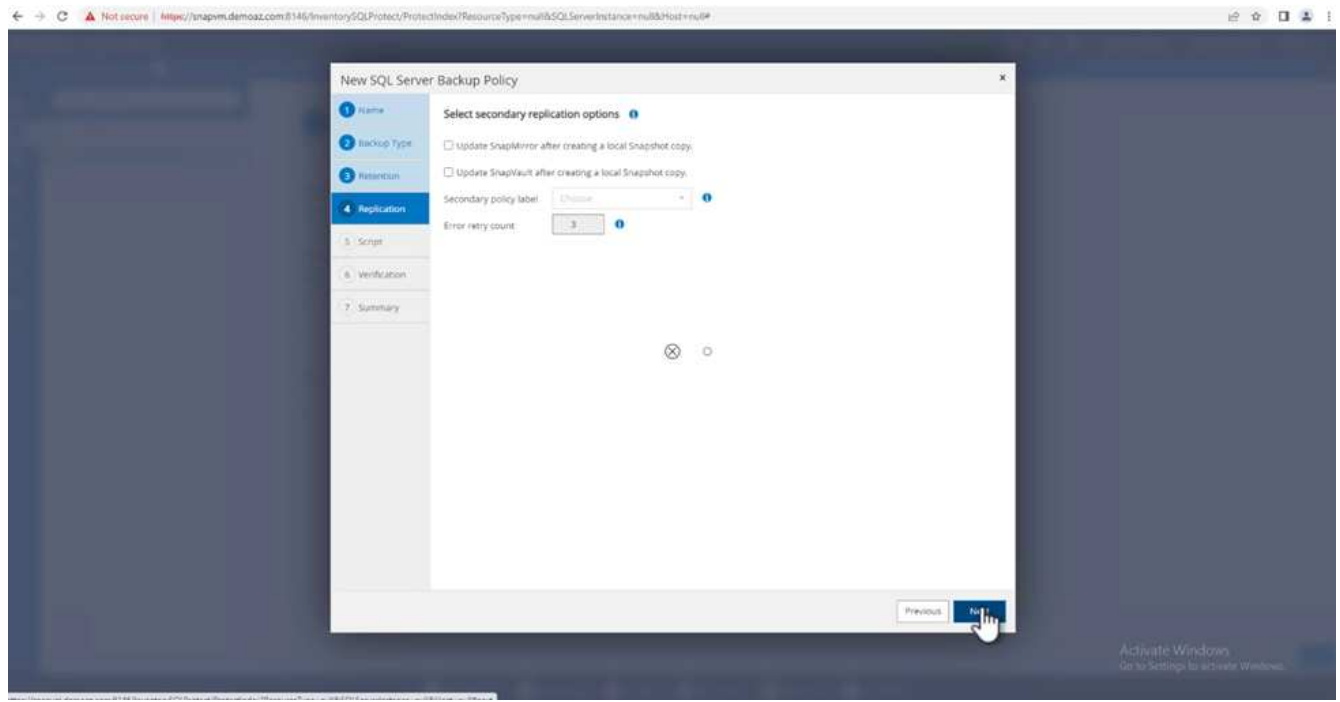
9. Configure the log backup retention settings. Click **Next**.

- 1 Name
- 2 Backup Type
- 3 Retention**
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

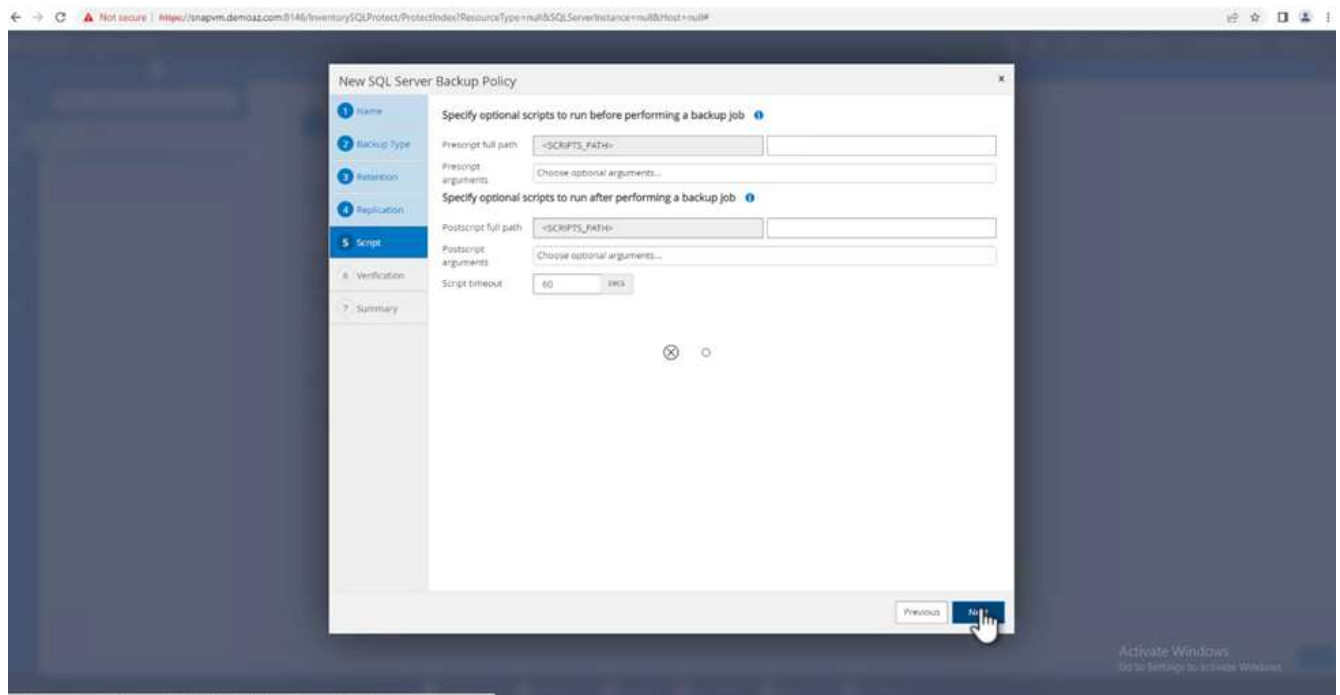
Log backup retention settings

Up-to-the-minute (UTM) retention settings retains log backups created as part of full backup and full and log backup operations. UTM retention settings also decides for how many full backups the log backups are to be retained. For example, if UTM retention settings is configured to retain log backups of the last 5 full backups, then the log backups of the last 5 full backups are retained and the rest are deleted.


10. (Optional) Configure the replication options.




11. (Optional) Configure any scripts to run before performing a backup job.



12. (Optional) Configure backup verification.




13. On the **Summary** page, click **Finish**.




Configure and protect multiple MSSQL Server databases


1. Click the newly created transaction log backup policy.



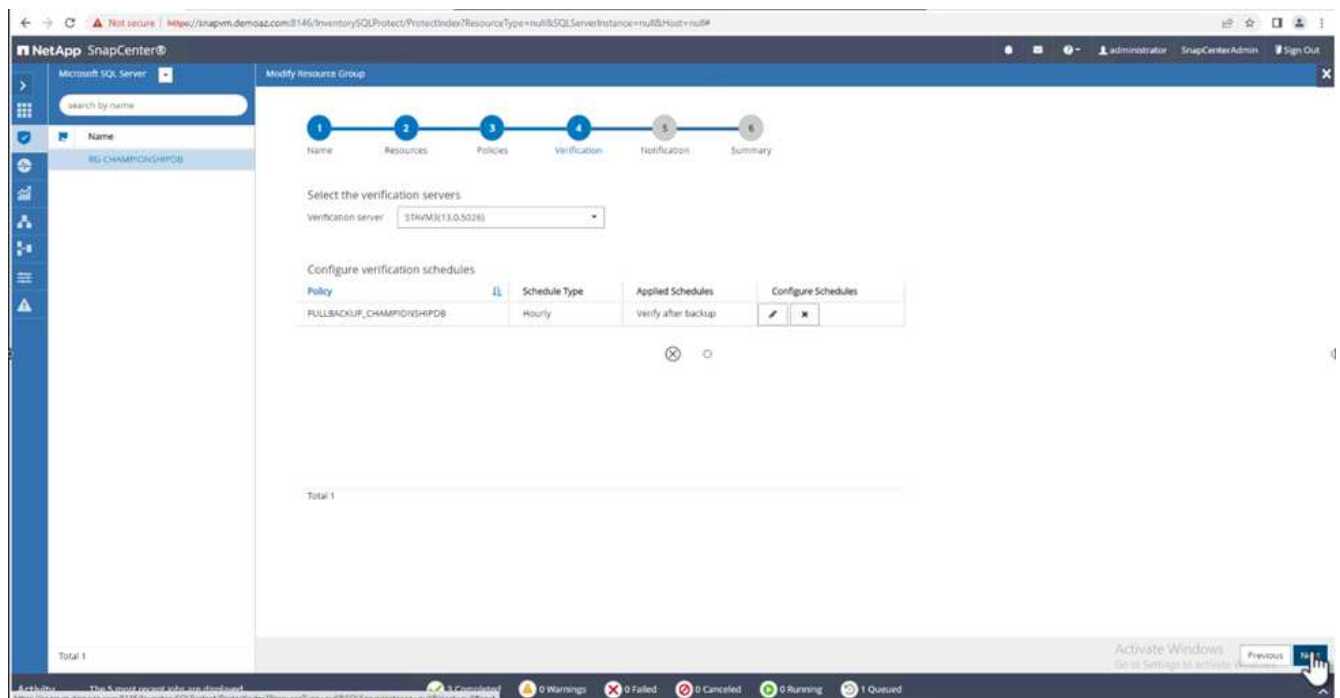
2. Set the **Start date** and **Expires on** date.
3. Enter the frequency of the log backup policy depending on the SLA, RTP, and RPO. Click OK.




4. You can see both policies. Click **Next**.




5. Configure the verification server.



6. Configure email notification.




7. On the **Summary** page, click **Finish**.




Triggering an on-demand transaction log backup for multiple SQL Server databases

To trigger an on-demand backup of the transactional log for multiple SQL server databases, complete the following steps:

1. On the newly created policy page, select **Backup now** at the upper right of the page.




2. From the pop-up on the **Policy** tab, select the drop-down menu, select the backup policy, and configure the transaction log backup.




3. Click **Backup**. A new window is displayed.

4. Click **Yes** to confirm the backup policy.



Monitoring

Move to the **Monitoring** tab and monitor the progress of the backup job.



Restore and recovery

See the following prerequisites necessary for restoring a SQL Server database in SnapCenter.

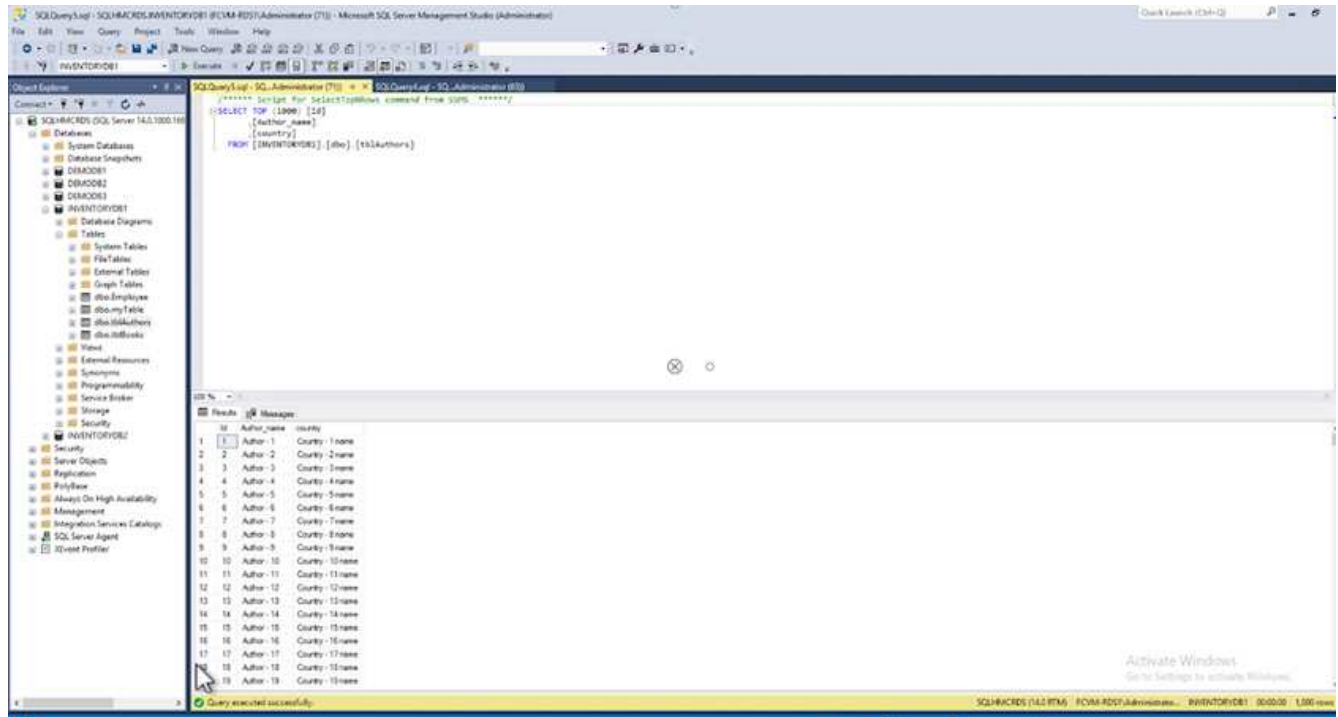
- The target instance must be online and running before a restore job completes.
- SnapCenter operations that are scheduled to run against the SQL Server database must be disabled, including any jobs scheduled on remote management or remote verification servers.

- If you are restoring custom log directory backups to an alternate host, the SnapCenter server and the plugin host must have the same SnapCenter version installed.
- You can restore the system database to an alternate host.
- SnapCenter can restore a database in a Windows cluster without taking the SQL Server cluster group offline.


Restoring deleted tables on a SQL Server database to a point in time

To restore a SQL Server database to a point in time, complete the following steps:


1. The following screenshot shows the initial state of the SQL Server database before the deleted tables.



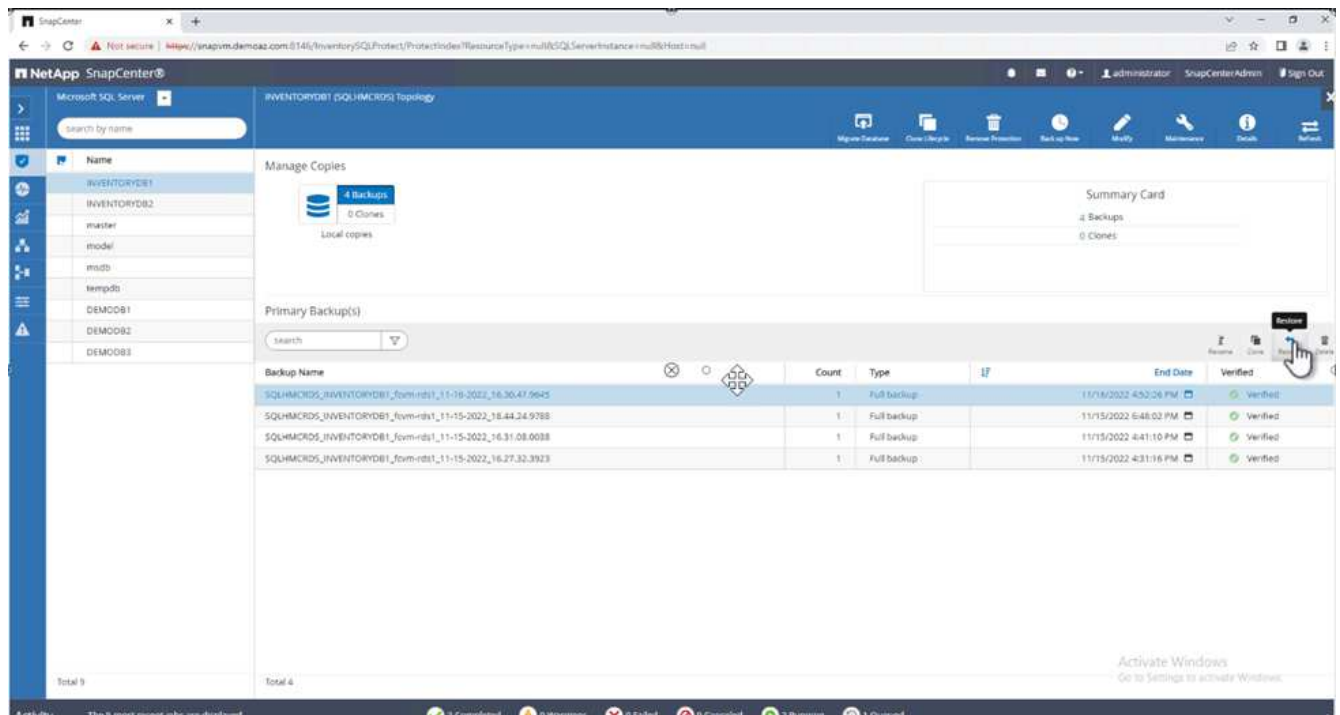
The screenshot shows that 20 rows were deleted from the table.




2. Log into SnapCenter Server. From the **Resources** tab, select the database.



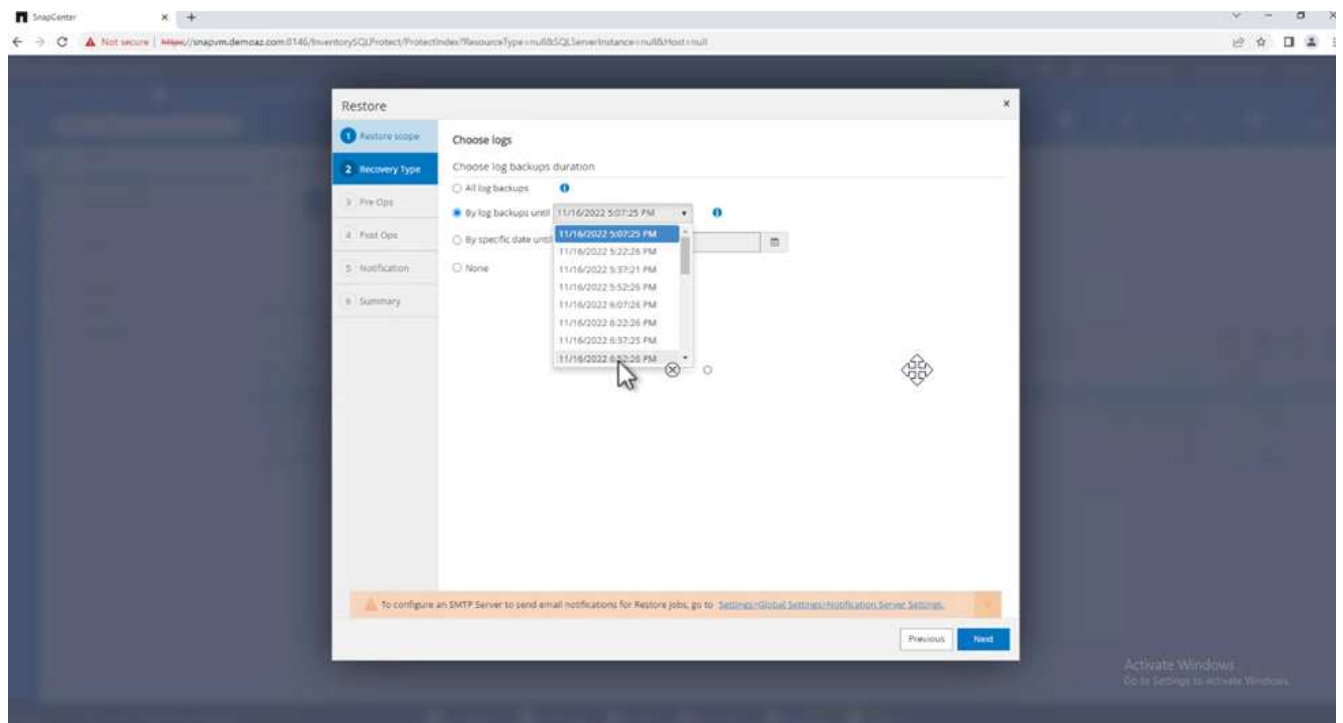

3. Select the most recent backup.
4. On the right, select **Restore**.



5. A new window is displayed. Select the **Restore** option.
6. Restore the database to the same host where the backup was created. Click **Next**.




7. For the **Recovery type**, select **All log backups**. Click **Next**.



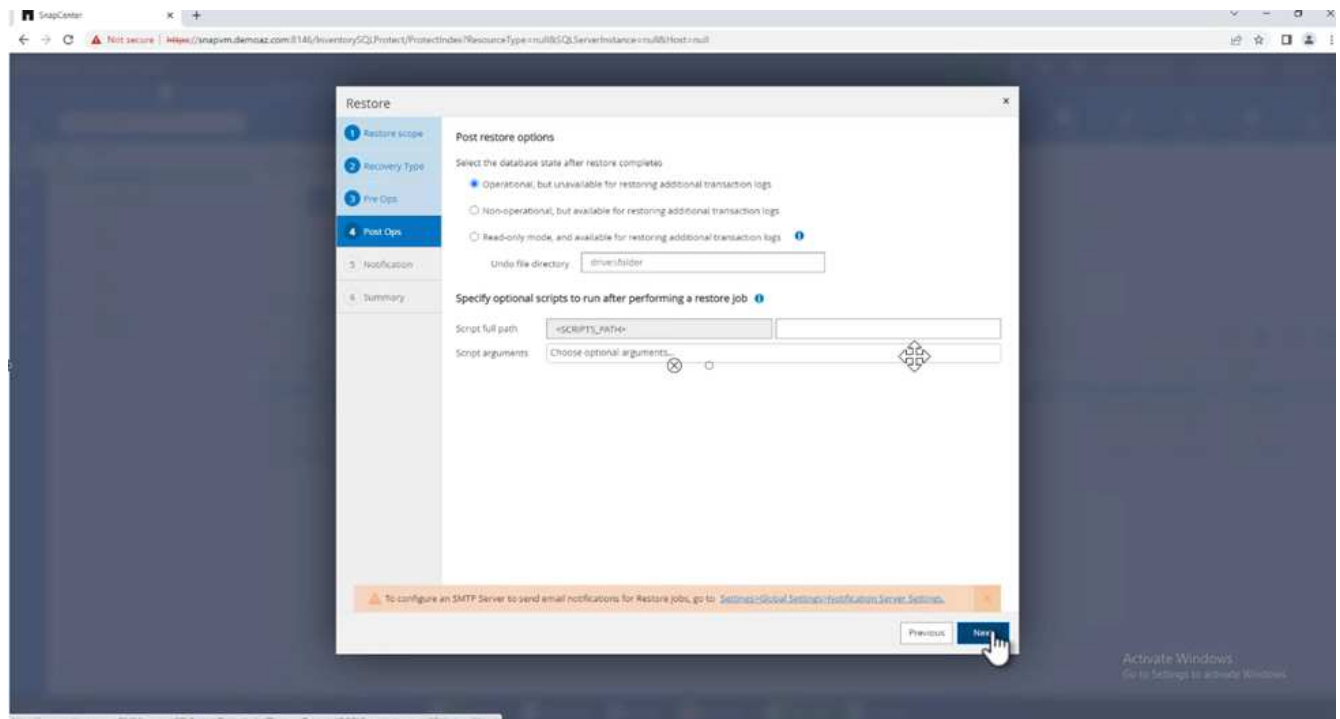
Pre- restore options:

1. Select the option **Overwrite the database with same name during restore**. Click **Next**.




Post-restore options:


1. Select the option **Operational, but unavailable for restoring additional transaction logs**. Click **Next**.



2. Provide the email settings. Click **Next**.



3. On the **Summary** page, click **Finish**.



Monitoring the restore progress

1. From the **Monitoring** tab, click the restore job details to view the progress of the restore job.

The screenshot shows the NetApp SnapCenter interface with the 'Jobs' tab selected. A table lists various backup jobs with columns for ID, Status, Name, Start date, End date, and Owner. A mouse cursor is hovering over the 'Send EMS Messages' task in the job details view.

ID	Status	Name	Start date	End date	Owner
124	✓	Restore 'SQLMCRDSINVENTORYDB1'	11/16/2022 11:11:03 PM		Administrator
130	✓	Backup of Resource Group 'SQLMCRDS_INVENTORYDB1' with policy 'InventoryDB_logbackup_policy'	11/16/2022 11:00:01 PM		Administrator
134	✓	Backup of Resource Group 'RG1-DEMO00' with policy 'demoDB_logbackup_policy'	11/16/2022 10:59:02 PM	11/16/2022 11:10:54 PM	Administrator
133	✓	Backup of Resource Group 'SQLMCRDS_INVENTORYDB2' with policy 'InventoryDB2_fullbackup'	11/16/2022 10:55:01 PM	11/16/2022 10:58:50 PM	Administrator
132	✓	Backup of Resource Group 'SQLMCRDLS_INVENTORYDB1' with policy 'InventoryDB_logbackup_policy'	11/16/2022 10:45:01 PM	11/16/2022 11:10:54 PM	Administrator
131	✓	Backup of Resource Group 'RG1-DEMO00' with policy 'demoDB_logbackup_policy'	11/16/2022 10:44:02 PM	11/16/2022 10:55:53 PM	Administrator
150	✓	Backup of Resource Group 'SQLMCRDS_INVENTORYDB1' with policy 'InventoryDB_logbackup_policy'	11/16/2022 10:30:01 PM	11/16/2022 10:55:54 PM	Administrator
148	✓	Backup of Resource Group 'RG1-DEMO00' with policy 'demoDB_logbackup_policy'	11/16/2022 10:29:02 PM	11/16/2022 10:40:53 PM	Administrator
146	✓	Backup of Resource Group 'SQLMCRDS_INVENTORYDB1' with policy 'InventoryDB_logbackup_policy'	11/16/2022 10:15:01 PM	11/16/2022 10:40:53 PM	Administrator
147	✓	Backup of Resource Group 'RG1-DEMO00' with policy 'demoDB_logbackup_policy'	11/16/2022 10:14:02 PM	11/16/2022 10:25:53 PM	Administrator
146	✓	Backup of Resource Group 'SQLMCRDS_INVENTORYDB1' with policy 'InventoryDB_logbackup_policy'	11/16/2022 10:00:01 PM	11/16/2022 10:25:53 PM	Administrator
145	✓	Backup of Resource Group 'RG1-DEMO00' with policy 'demoDB_logbackup_policy'	11/16/2022 9:58:02 PM	11/16/2022 10:10:53 PM	Administrator
143	✓	Backup of Resource Group 'SQLMCRDS_INVENTORYDB1' with policy 'InventoryDB_logbackup_policy'	11/16/2022 9:45:01 PM	11/16/2022 10:10:53 PM	Administrator
142	✓	Backup of Resource Group 'RG1-DEMO00' with policy 'demoDB_logbackup_policy'	11/16/2022 9:44:02 PM	11/16/2022 9:55:54 PM	Administrator
142	✓	Backup of Resource Group 'SQLMCRDS_INVENTORYDB1' with policy 'InventoryDB_logbackup_policy'	11/16/2022 9:30:01 PM	11/16/2022 9:55:54 PM	Administrator
141	✓	Backup of Resource Group 'RG1-DEMO00' with policy 'demoDB_logbackup_policy'	11/16/2022 9:29:02 PM	11/16/2022 9:40:53 PM	Administrator
140	✓	Backup of Resource Group 'SQLMCRDS_INVENTORYDB1' with policy 'InventoryDB_logbackup_policy'	11/16/2022 9:15:01 PM	11/16/2022 9:40:53 PM	Administrator
139	✓	Backup of Resource Group 'RG1-DEMO00' with policy 'demoDB_logbackup_policy'	11/16/2022 9:14:02 PM	11/16/2022 9:25:54 PM	Administrator
138	✓	Backup of Resource Group 'SQLMCRDS_INVENTORYDB1' with policy 'InventoryDB_logbackup_policy'	11/16/2022 9:00:01 PM	11/16/2022 9:25:54 PM	Administrator
137	✓	Backup of Resource Group 'RG1-DEMO00' with policy 'demoDB_logbackup_policy'	11/16/2022 8:59:02 PM	11/16/2022 9:10:53 PM	Administrator
136	✓	Backup of Resource Group 'SQLMCRDLS_INVENTORYDB1' with policy 'InventoryDB_logbackup_policy'	11/16/2022 8:45:01 PM	11/16/2022 9:10:53 PM	Administrator
135	✓	Backup of Resource Group 'RG1-DEMO00' with policy 'demoDB_logbackup_policy'	11/16/2022 8:44:02 PM	11/16/2022 8:55:54 PM	Administrator
134	✓	Backup of Resource Group 'SQLMCRDS_INVENTORYDB1' with policy 'InventoryDB_logbackup_policy'	11/16/2022 8:30:01 PM	11/16/2022 8:55:54 PM	Administrator
133	✓	Backup of Resource Group 'RG1-DEMO00' with policy 'demoDB_logbackup_policy'	11/16/2022 8:29:02 PM	11/16/2022 8:40:53 PM	Administrator


2. Restore the job details.

The screenshot shows the 'Job Details' dialog box for the job 'Restore 'SQLMCRDSINVENTORYDB1''. The dialog displays a task list with the following items:

- Restore 'SQLMCRDSINVENTORYDB1'
- Job 157 | Tail log backup of 'SQLMCRDSINVENTORYDB1'
- PCVM-RD01-Demoapp.com
 - Preparing for Backup
 - Creating SQL Backup
 - Finalizing Backup
 - Send EMS Messages
- PCVM-RD01-Demoapp.com

At the bottom of the dialog, there are buttons for 'View Logs', 'Cancel', and 'Close'. A mouse cursor is clicking the 'Close' button.

3. Return to SQL Server host > database > table are present.



Where to find additional information

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

- [TR-4714: Best Practices Guide for Microsoft SQL Server using NetApp SnapCenter](#)

<https://www.netapp.com/pdf.html?item=/media/12400-tr4714pdf.pdf>

- [Requirements for restoring a database](#)

https://docs.netapp.com/us-en/snapcenter-45/protect-scsql/concept_requirements_for_restoring_a_database.html

- [Understanding cloned database lifecycles](#)

<https://library.netapp.com/ecmdocs/ECMP1217281/html/GUID-4631AFF4-64FE-4190-931E-690FCADA5963.html>

TR-4923: SQL Server on AWS EC2 using Amazon FSx for NetApp ONTAP

This solution covers the deployment of SQL Server on AWS EC2 using Amazon FSx for NetApp ONTAP.

Authors: Pat Sinthusan and Niyaz Mohamed, NetApp

Introduction

Many companies that would like to migrate applications from on-premises to the cloud find that the effort is hindered by the differences in capabilities offered by on-premises storage systems and cloud storage services. That gap has made migrating enterprise applications such as Microsoft SQL Server much more problematic. In particular, gaps in the services needed to run an enterprise application such as robust snapshots, storage efficiency capabilities, high availability, reliability, and consistent performance have forced customers to make design tradeoffs or forgo application migration. With FSx for NetApp ONTAP, customers no longer need to

compromise. FSx for NetApp ONTAP is a native (1st party) AWS service sold, supported, billed, and fully managed by AWS. It uses the power of NetApp ONTAP to provide the same enterprise grade storage and data management capabilities NetApp has provided on-premises for three decades in AWS as a managed service.

With SQL Server on EC2 instances, database administrators can access and customize their database environment and the underlying operating system. A SQL Server on EC2 instance in combination with [AWS FSx ONTAP](#) to store the database files, enables high performance, data management, and a simple and easy migration path using block-level replication. Therefore, you can run your complex database on AWS VPC with an easy lift-and-shift approach, fewer clicks, and no schema conversions.

Benefits of using Amazon FSx for NetApp ONTAP with SQL Server

Amazon FSx for NetApp ONTAP is the ideal file storage for SQL Server deployments in AWS. Benefits include the following:

- Consistent high performance and throughput with low latency
- Intelligent caching with NVMe cache to improve performance
- Flexible sizing so that you can increase or shrink capacity, throughput, and IOPs on the fly
- Efficient on-premises-to-AWS block replication
- The use of iSCSI, a well-known protocol for the database environment
- Storage efficiency features like thin provisioning and zero-footprint clones
- Backup time reduction from hours to mins, thereby reducing the RTO
- Granular backup and recovery of SQL databases with the intuitive NetApp SnapCenter UI
- The ability to perform multiple test migrations before actual migration
- Shorter downtime during migration and overcoming migration challenges with file-level or I/O-level copy
- Reducing MTTR by finding the root cause after a major release or patch update

Deploying SQL Server databases on FSx ONTAP with the iSCSI protocol, as is commonly used on-premises, provides an ideal database storage environment with superior performance, storage efficiency, and data-management capabilities. Using multiple iSCSI sessions, assuming a 5% working set size, fitting a Flash Cache delivers over 100K IOPs with the FSx ONTAP service. This configuration provides complete control over performance for the most demanding applications. SQL Server running on smaller EC2 instances connected to FSx for ONTAP can perform the same as SQL Server running on a much larger EC2 instance, because only network bandwidth limits are applied against FSx for ONTAP. Reducing the size of instances also reduces the compute cost, which provides a TCO-optimised deployment. The combination of SQL using iSCSI, SMB3.0 with multichannel, continuous availability shares on FSx for ONTAP provides great advantages for SQL workloads.

Before you begin

The combination of Amazon FSx for NetApp ONTAP and SQL Server on EC2 instance enables the creation of enterprise-level database storage designs that can meet today's most demanding application requirements. To optimize both technologies, it is vital to understand SQL Server I/O patterns and characteristics. A well-designed storage layout for a SQL Server database supports the performance of SQL Server and the management of the SQL Server infrastructure. A good storage layout also allows the initial deployment to be successful and the environment to grow smoothly over time as your business grows.

Prerequisites

Before you complete the steps in this document, you should have the following prerequisites:

- An AWS account
- Appropriate IAM roles to provision EC2 and FSx for ONTAP
- A Windows Active Directory domain on EC2
- All SQL Server nodes must be able to communicate with each other
- Make sure DNS resolution works and host names can be resolved. If not, use host file entry.
- General knowledge of SQL Server installation

Also, please refer to the NetApp Best Practices for SQL Server environments to ensure the best storage configuration.

Best practice configurations for SQL Server environments on EC2

With FSx ONTAP, procuring storage is the easiest task and can be performed by updating the file system. This simple process enables dynamic cost and performance optimization as needed, it helps to balance the SQL workload, and it is also a great enabler for thin provisioning. FSx ONTAP thin provisioning is designed to present more logical storage to EC2 instances running SQL Server than what is provisioned in the file system. Instead of allocating space upfront, storage space is dynamically allocated to each volume or LUN as data is written. In most configurations, free space is also released back when data in the volume or LUN is deleted (and is not being held by any Snapshot copies). The following table provides configuration settings for dynamically allocating storage.

Setting	Configuration
Volume guarantee	None (set by default)
LUN reservation	Enabled
fractional_reserve	0% (set by default)
snap_reserve	0%
Autodelete	volume / oldest_first
Autosize	On
try_first	Autogrow
Volume tiering policy	Snapshot only
Snapshot policy	None

With this configuration, the total size of the volumes can be greater than the actual storage available in the file system. If the LUNs or Snapshot copies require more space than is available in the volume, the volumes automatically grow, taking more space from the containing file system. Autogrow allows FSx ONTAP to automatically increase the size of the volume up to a maximum size that you predetermine. There must be space available in the containing file system to support the automatic growth of the volume. Therefore, with autogrow enabled, you should monitor the free space in the containing filesystem and update the file system when needed.

Along with this, set the [space-allocation](#) option on LUN to enabled so that FSx ONTAP notifies the EC2 host when the volume has run out of space and the LUN in the volume cannot accept writes. Also, this option enables FSx for ONTAP to reclaim space automatically when the SQL Server on EC2 host deletes data. The space-allocation option is set to disabled by default.



If a space-reserved LUN is created in a none-guaranteed volume, then the LUN behaves the same as a non-space-reserved LUN. This is because a none-guaranteed volume has no space to allocate to the LUN; the volume itself can only allocate space as it is written to due to its none guarantee.

With this configuration, FSx ONTAP administrators can generally size the volume so that they must manage and monitor the used space in the LUN on the host side and in the file system.



NetApp recommends using a separate file system for SQL server workloads. If the file system is used for multiple applications, monitor the space usage of both the file system and volumes within the file system to make sure that volumes are not competing for available space.



Snapshot copies used to create FlexClone volumes are not deleted by the autodelete option.



Overcommitment of storage must be carefully considered and managed for a mission-critical application such as SQL server for which even a minimal outage cannot be tolerated. In such a case, it is best to monitor storage consumption trends to determine how much, if any, overcommitment is acceptable.

Best Practices

1. For optimal storage performance, provision file-system capacity to 1.35x times the size of total database usage.
2. Appropriate monitoring accompanied by an effective action plan is required when using thin provisioning to avoid application downtime.
3. Make sure to set Cloudwatch and other monitoring tool alerts so that people are contacted with enough time to react as storage is filled.

Configure Storage for SQL Server and deploy Snapcenter for Backup, Restore and clone operations

In order to perform SQL server operations with SnapCenter, you must first create volumes and LUNs for SQL server.

Create volumes and LUNs for SQL Server

To create volumes and LUNs for SQL Server, complete the following steps:

1. Open the Amazon FSx console at <https://console.aws.amazon.com/fsx/>
2. Create an Amazon FSx for the NetApp ONTAP file system using the Standard Create option under Creation Method. This allows you to define FSxadmin and vsadmin credentials.

Creation method

Quick create

Use recommended best-practice configurations. Most configuration options can be changed after the file system is created.

Standard create

You set all of the configuration options, including specifying performance, networking, security, backups, and maintenance.

3. Specify the password for fsxadmin.

File system administrative password

Password for this file system's "fsxadmin" user, which you can use to access the ONTAP CLI or REST API.

- Don't specify a password
- Specify a password

Password

Confirm password

4. Specify the password for SVMs.

SVM administrative password

Password for this SVM's "vsadmin" user, which you can use to access the ONTAP CLI or REST API.

- Don't specify a password
- Specify a password

Password

Confirm password

5. Create volumes by following the step listed in [Creating a volume on FSx for NetApp ONTAP](#).

Best practices

- Disable storage Snapshot copy schedules and retention policies. Instead, use NetApp SnapCenter to coordinate Snapshot copies of the SQL Server data and log volumes.

- Configure databases on individual LUNs on separate volumes to leverage fast and granular restore functionality.
- Place user data files (.mdf) on separate volumes because they are random read/write workloads. It is common to create transaction log backups more frequently than database backups. For this reason, place transaction log files (.ldf) on a separate volume from the data files so that independent backup schedules can be created for each. This separation also isolates the sequential write I/O of the log files from the random read/write I/O of data files and significantly improves SQL Server performance.
- Tempdb is a system database used by Microsoft SQL Server as a temporary workspace, especially for I/O intensive DBCC CHECKDB operations. Therefore, place this database on a dedicated volume. In large environments in which volume count is a challenge, you can consolidate tempdb into fewer volumes and store it in the same volume as other system databases after careful planning. Data protection for tempdb is not a high priority because this database is recreated every time Microsoft SQL Server is restarted.

6. Use the following SSH command to create volumes:

```
vol create -vserver svm001 -volume vol_awssqlprod01_data -aggregate
aggr1 -size 800GB -state online -tiering-policy snapshot-only
-percent-snapshot-space 0 -autosize-mode grow -snapshot-policy none
-security-style ntfs
volume modify -vserver svm001 -volume vol_awssqlprod01_data
-fractional-reserve 0
volume modify -vserver svm001 -volume vol_awssqlprod01_data -space
-mgmt-try-first vol_grow
volume snapshot autodelete modify -vserver svm001 -volume
vol_awssqlprod01_data -delete-order oldest_first
```

7. Start the iSCSI service with PowerShell using elevated privileges in Windows Servers.

```
Start-Service -Name msiscsi
Set-Service -Name msiscsi -StartupType Automatic
```

8. Install Multipath-IO with PowerShell using elevated privileges in Windows Servers.

```
Install-WindowsFeature -name Multipath-IO -Restart
```

9. Find the Windows initiator Name with PowerShell using elevated privileges in Windows Servers.

```
Get-InitiatorPort | select NodeAddress
```

```
PS C:\Users\administrator.CONTOSO> Get-InitiatorPort | select NodeAddress
NodeAddress
-----
iqn.1991-05.com.microsoft:ws2019-sql1.contoso.net
```

10. Connect to Storage virtual machines (SVM) using putty and create an iGroup.

```
igroup create -igroup igrp_ws2019sql1 -protocol iscsi -ostype windows -initiator iqn.1991-05.com.microsoft:ws2019-sql1.contoso.net
```

11. Use the following SSH command to create LUNs:

```
lun create -path /vol/vol_awssqlprod01_data/lun_awssqlprod01_data -size 700GB -ostype windows_2008 -space-allocation enabled
lun create -path /vol/vol_awssqlprod01_log/lun_awssqlprod01_log -size 100GB -ostype windows_2008 -space-allocation enabled
```

```
svmsql:> lun create -path /vol/vol_awssqlprod01_data/lun_awssqlprod01_data -size 700GB -ostype windows_2008
Created a LUN of size 700g (751619276800)

svmsql:> lun create -path /vol/vol_awssqlprod01_log/lun_awssqlprod01_log -size 100GB -ostype windows_2008
Created a LUN of size 100g (107374182400)

svmsql:> lun show
Vserver Path State Mapped Type Size
-----
svmsql /vol/vol_awssqlprod01_data/lun_awssqlprod01_data online unmapped windows_2008 700GB
svmsql /vol/vol_awssqlprod01_log/lun_awssqlprod01_log online unmapped windows_2008 100GB
2 entries were displayed.
```

12. To achieve I/O alignment with the OS partitioning scheme, use windows_2008 as the recommended LUN type. Refer [here](#) for additional information.
13. Use the following SSH command to the map igroup to the LUNs that you just created.

```
lun show
lun map -path /vol/vol_awssqlprod01_data/lun_awssqlprod01_data -igroup igrp_awssqlprod01
lun map -path /vol/vol_awssqlprod01_log/lun_awssqlprod01_log -igroup igrp_awssqlprod01
```

```

svmsql::> lun show
Vserver  Path                                     State  Mapped  Type      Size
-----
svmsql   /vol/vol_awssqlprod01_data/lun_awssqlprod01_data  online unmapped windows_2008 700GB
svmsql   /vol/vol_awssqlprod01_log/lun_awssqlprod01_log    online unmapped windows_2008 100GB
2 entries were displayed.

svmsql::> lun map -path /vol/vol_awssqlprod01_data/lun_awssqlprod01_data -igroup igrp_awssqlprod01
svmsql::> lun map -path /vol/vol_awssqlprod01_log/lun_awssqlprod01_log -igroup igrp_awssqlprod01

svmsql::>
svmsql::> lun show
Vserver  Path                                     State  Mapped  Type      Size
-----
svmsql   /vol/vol_awssqlprod01_data/lun_awssqlprod01_data  online mapped   windows_2008 700GB
svmsql   /vol/vol_awssqlprod01_log/lun_awssqlprod01_log    online mapped   windows_2008 100GB
2 entries were displayed.

```


- For a shared disk that uses the Windows Failover Cluster, run an SSH command to map the same LUN to the igroup that belong to all servers that participate in the Windows Failover Cluster.
- Connect Windows Server to an SVM with an iSCSI target. Find the target IP address from AWS Portal.

svmsql (svm-09e98ab33a31b724a)

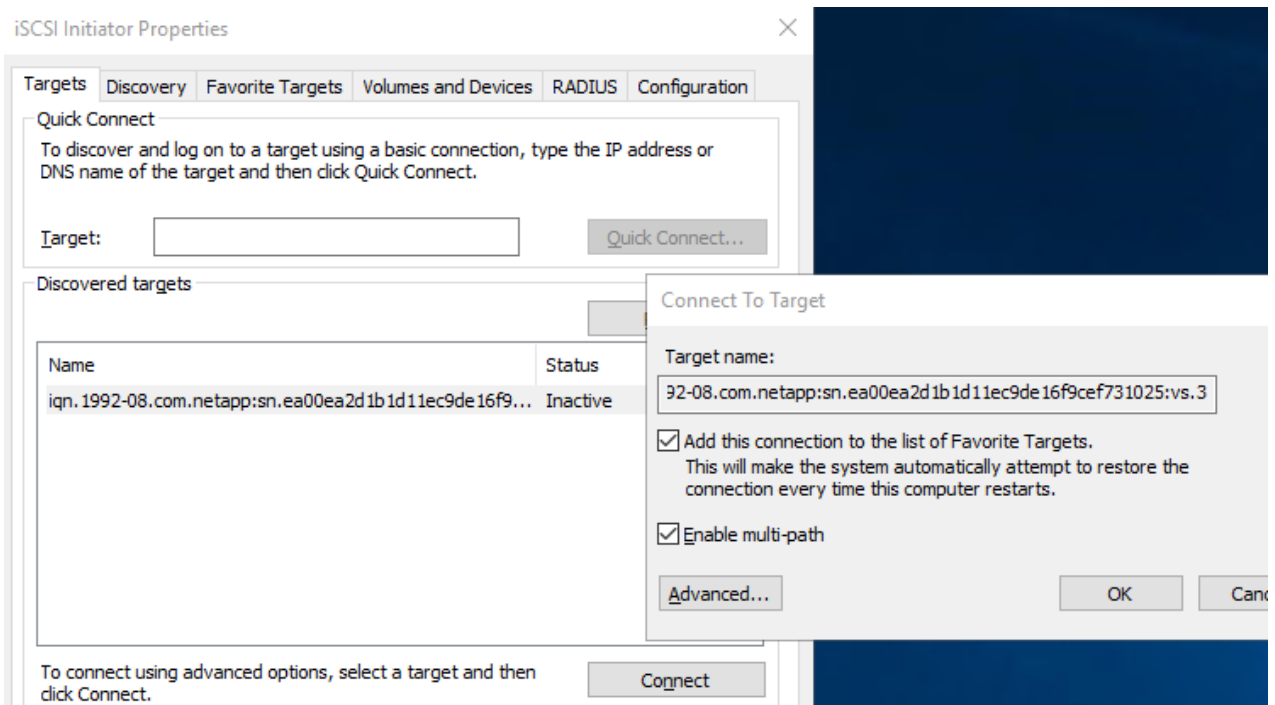
Summary	
SVM ID svm-09e98ab33a31b724a	Creation time 2021-09-21T13:19:34-07:00
SVM name svmsql	Lifecycle state Created
UUID ea00ea2d-1b1d-11ec-9de1-6f9cef731025	Subtype DEFAULT
File system ID fs-0ab4b447ebd6082aa	
Resource ARN arn:aws:fsx:us-west-2:139763910815:storage-virtual-machine/fs-0ab4b447ebd6082aa/svm-09e98ab33a31b724a	

Endpoints	
Management DNS name svm-09e98ab33a31b724a.fs-0ab4b447ebd6082aa.fsx.us-west-2.amazonaws.com	Management IP address 198.19.255.153
NFS DNS name svm-09e98ab33a31b724a.fs-0ab4b447ebd6082aa.fsx.us-west-2.amazonaws.com	NFS IP address 198.19.255.153
iSCSI DNS name iscsi.svm-09e98ab33a31b724a.fs-0ab4b447ebd6082aa.fsx.us-west-2.amazonaws.com	iSCSI IP addresses 10.2.1.167, 10.2.2.12

- From Server Manager and the Tools menu, select the iSCSI Initiator. Select the Discovery tab and then select Discover Portal. Supply the iSCSI IP address from previous step and select Advanced. From Local Adapter, select Microsoft iSCSI Initiator. From Initiator IP, select the IP of the server. Then select OK to close all windows.



17. Repeat step 12 for the second iSCSI IP from the SVM.
18. Select the **Targets** tab, select **Connect**, and select **Enable multi-path**.



19. For best performance, add more sessions; NetApp recommends creating five iSCSI sessions. Select **Properties** > **Add session** > **Advanced** and repeat step 12.

```
$TargetPortals = ('10.2.1.167', '10.2.2.12')
foreach ($TargetPortal in $TargetPortals) {New-IscsiTargetPortal
-TargetPortalAddress $TargetPortal}
```

```
$TargetPortals = ('10.2.1.167', '10.2.2.12')
foreach ($TargetPortal in $TargetPortals) {New-IscsiTargetPortal -TargetPortalAddress $TargetPortal}

InitiatorInstanceName :
InitiatorPortalAddress :
IsDataDigest          : False
IsHeaderDigest        : False
TargetPortalAddress   : 10.2.1.167
TargetPortalPortNumber : 3260
PSComputerName        :

InitiatorInstanceName :
InitiatorPortalAddress :
IsDataDigest          : False
IsHeaderDigest        : False
TargetPortalAddress   : 10.2.2.12
TargetPortalPortNumber : 3260
PSComputerName        :
```

Best practices

- Configure five iSCSI sessions per target interface for optimal performance.
- Configure a round-robin policy for the best overall iSCSI performance.
- Make sure that the allocation unit size is set to 64K for partitions when formatting the LUNs
 1. Run the following PowerShell command to make sure that the iSCSI session is persisted.

```
$targets = Get-IscsiTarget
foreach ($target in $targets)
{
Connect-IscsiTarget -IsMultipathEnabled $true -NodeAddress
$target.NodeAddress -IsPersistent $true
}
```

```
PS C:\windows\system32> Connect-IscsiTarget -NodeAddress (Get-IscsiTarget | select -ExpandProperty NodeAddress)

AuthenticationType      : NONE
InitiatorInstanceName   : ROOT\ISCSIPT\0000_0
InitiatorNodeAddress    : iqn.1991-05.com.microsoft:awssqlprod01.cloudheroes.dom
InitiatorPortalAddress  : 0.0.0.0
InitiatorSideIdentifier : 400001370000
IsConnected             : True
IsDataDigest            : False
IsDiscovered            : True
IsHeaderDigest          : False
IsPersistent            : True
NumberOfConnections     : 1
SessionIdentifier       : ffff9988350ff010-4000013700000012
TargetNodeAddress       : iqn.1992-08.com.netapp:sn.ea00ea2d1b1d11ec9de16f9cef731025:vs.3
TargetSideIdentifier    : 0200
PSComputerName          :
```

2. Initialize disks with the following PowerShell command.


```
$disks = Get-Disk | where PartitionStyle -eq raw
foreach ($disk in $disks) {Initialize-Disk $disk.Number}
```

```
PS C:\Windows\system32> $disks = Get-Disk | where PartitionStyle -eq raw
foreach ($disk in $disks) {Initialize-Disk $disk.Number}
PS C:\Windows\system32> Get-Disk
```

Number	Friendly Name	Serial Number	HealthStatus	OperationalStatus	Total Size	Partition Style
0	AWS PVDISK					
1	NETAPP LUN C-Mode	vo105d1c31fcb4c790ab	Healthy	Online	30 GB	MBR
2	NETAPP LUN C-Mode	1wB0p7RmR2s2	Healthy	Online	700 GB	GPT
		1wB0p7RmR2s3	Healthy	Online	100 GB	GPT

3. Run the Create Partition and Format Disk commands with PowerShell.

```
New-Partition -DiskNumber 1 -DriveLetter F -UseMaximumSize
Format-Volume -DriveLetter F -FileSystem NTFS -AllocationUnitSize
65536
New-Partition -DiskNumber 2 -DriveLetter G -UseMaximumSize
Format-Volume -DriveLetter G -FileSystem NTFS -AllocationUnitSize
65536
```

You can automate volume and LUN creation using the PowerShell script from Appendix B. LUNs can also be created using SnapCenter.

Once the volumes and LUNs are defined, you need to set up SnapCenter to be able to perform the database operations.

SnapCenter overview

NetApp SnapCenter is next-generation data protection software for tier-1 enterprise applications. SnapCenter, with its single-pane-of-glass management interface, automates and simplifies the manual, complex, and time-consuming processes associated with the backup, recovery, and cloning of multiple databases and other application workloads. SnapCenter leverages NetApp technologies, including NetApp Snapshots, NetApp SnapMirror, SnapRestore, and NetApp FlexClone. This integration allows IT organizations to scale their storage infrastructure, meet increasingly stringent SLA commitments, and improve the productivity of administrators across the enterprise.

SnapCenter Server requirements

The following table lists the minimum requirements for installing the SnapCenter Server and plug-in on Microsoft Windows Server.


Components	Requirement
Minimum CPU count	Four cores/vCPUs
Memory	Minimum: 8GB Recommended: 32GB
Storage space	Minimum space for installation: 10GB Minimum space for repository: 10GB
Supported operating system	<ul style="list-style-type: none">• Windows Server 2012• Windows Server 2012 R2• Windows Server 2016• Windows Server 2019
Software packages	<ul style="list-style-type: none">• .NET 4.5.2 or later• Windows Management Framework (WMF) 4.0 or later• PowerShell 4.0 or later

For detailed information, refer to [space and sizing requirements](#).

For version compatibility, see the [NetApp Interoperability Matrix Tool](#).

Database storage layout

The following figure depicts some considerations for creating the Microsoft SQL Server database storage layout when backing up with SnapCenter.



Best practices

1. Place databases with I/O-intensive queries or with large database size (say 500GB or more) on a separate volume for faster recovery. This volume should also be backed up by separate jobs.
2. Consolidate small-to-medium size databases that are less critical or have fewer I/O requirements to a single volume. Backing up a large number of databases residing in the same volume leads to fewer Snapshot copies that need to be maintained. It is also a best practice to consolidate Microsoft SQL Server instances to use the same volumes to control the number of backup Snapshot copies taken.
3. Create separate LUNs to store full text-related files and file-streaming related files.
4. Assign separate LUNs per host to store Microsoft SQL Server log backups.
5. System databases that store database server metadata configuration and job details are not updated frequently. Place system databases/tempdb in separate drives or LUNs. Do not place system databases in the same volume as the user databases. User databases have a different backup policy, and the frequency of user database backup is not same for system databases.
6. For Microsoft SQL Server Availability Group setup, place the data and log files for replicas in an identical folder structure on all nodes.

In addition to the performance benefit of segregating the user database layout into different volumes, the database also significantly affects the time required to back up and restore. Having separate volumes for data and log files significantly improves the restore time as compared to a volume hosting multiple user data files. Similarly, user databases with a high I/O intensive application are prone to an increase in the backup time. A more detailed explanation about backup and restore practices is provided later in this document.



Starting with SQL Server 2012 (11.x), system databases (Master, Model, MSDB, and TempDB), and Database Engine user databases can be installed with an SMB file server as a storage option. This applies to both stand-alone SQL Server and SQL Server failover cluster installations. This enables you to use FSx for ONTAP with all its performance and data management capabilities, including volume capacity, performance scalability, and data protection features, which SQL Server can take advantage of. Shares used by the application servers must be configured with the continuously available property set and the volume should be created with NTFS security style. NetApp Snapcenter cannot be used with databases placed on SMB shares from FSx for ONTAP.



For SQL Server databases that do not use SnapCenter to perform backups, Microsoft recommends placing the data and log files on separate drives. For applications that simultaneously update and request data, the log file is write intensive, and the data file (depending on your application) is read/write intensive. For data retrieval, the log file is not needed. Therefore, requests for data can be satisfied from the data file placed on its own drive.




When you create a new database, Microsoft recommends specifying separate drives for the data and logs. To move files after the database is created, the database must be taken offline. For more Microsoft recommendations, see [Place Data and Log Files on Separate Drives](#).

Installation and setup for SnapCenter

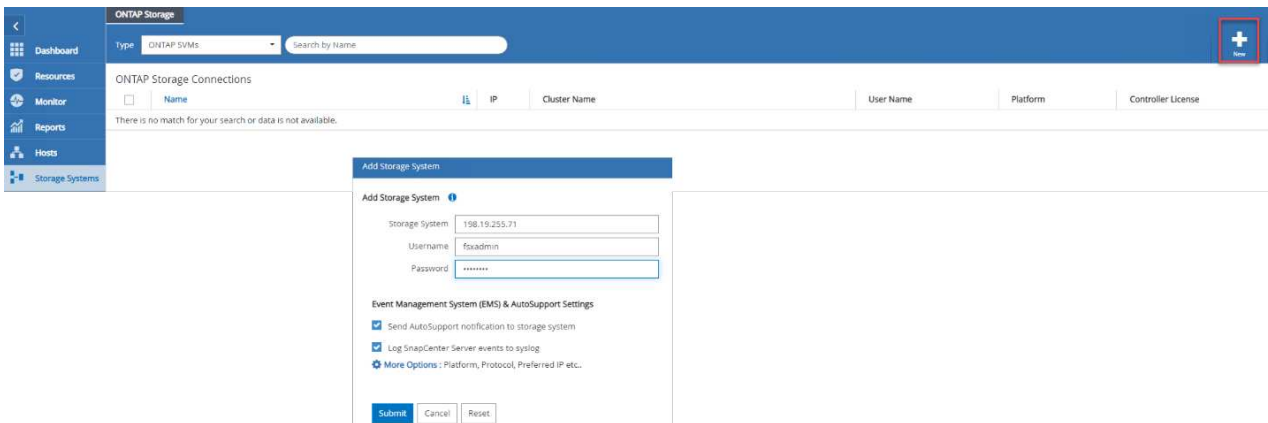
Follow the [Install the SnapCenter Server](#) and [Installing SnapCenter Plug-in for Microsoft SQL Server](#) to install and setup SnapCenter.

After Installing SnapCenter, complete the following steps to set it up.


1. To set up credentials, select **Settings** > **New** and then enter the credential information.



2. Add the storage system by selecting **Storage Systems** > **New** and then provide the appropriate FSx for ONTAP storage information.



3. Add hosts by selecting **Hosts** > **Add**, and then provide the host information. SnapCenter automatically installs the Windows and SQL Server plug-in. This process might take some time.



After all Plug-ins are installed, you must configure the log directory. This is the location where the transaction log backup resides. You can configure the log directory by selecting the host and then select configure the log directory.



SnapCenter uses a host log directory to store transaction log backup data. This is at the host and instance level. Each SQL Server host used by SnapCenter must have a host log directory configured to perform log backups. SnapCenter has a database repository, so metadata related to backup, restore, or cloning operations is stored in a central database repository.

The size of the host log directory is calculated as follows:

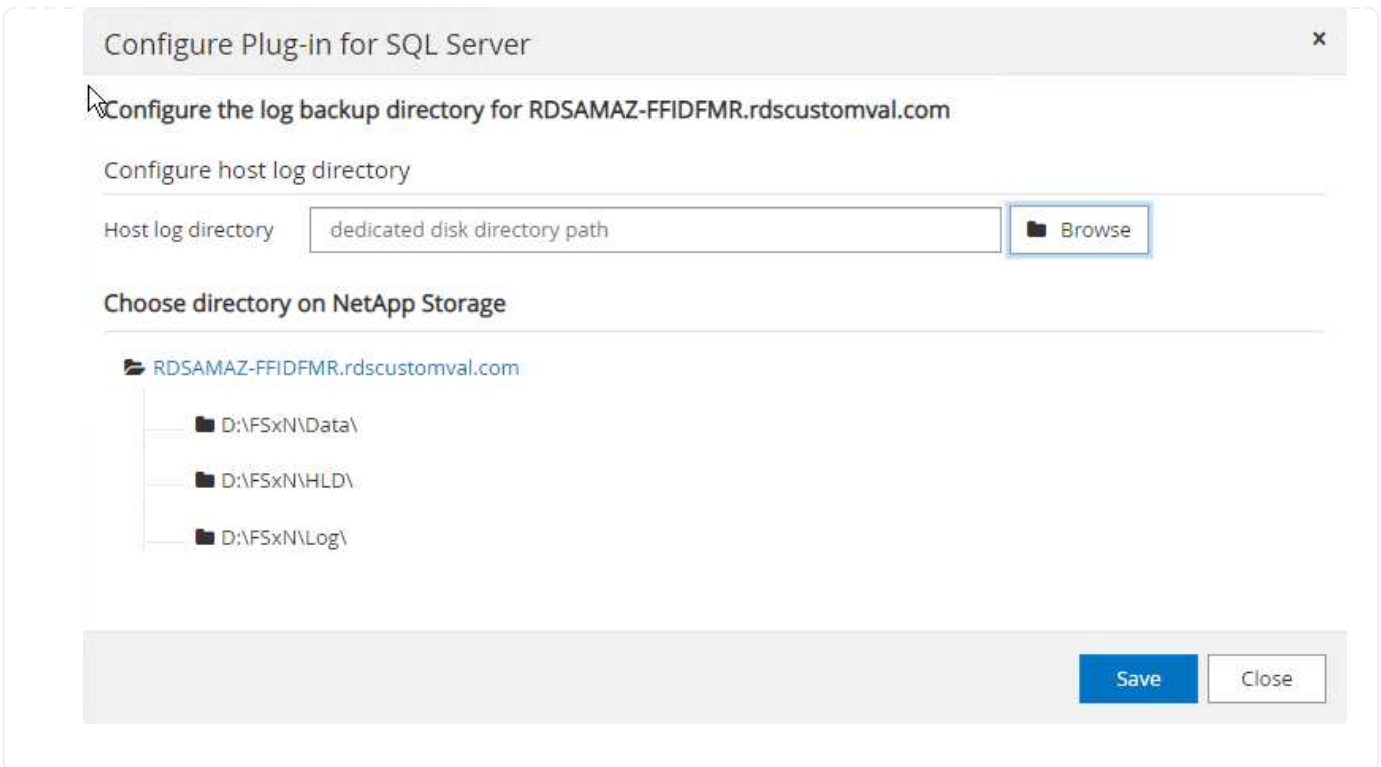
Size of host log directory = system database size + (maximum DB LDF size × daily log change rate % × (Snapshot copy retention) ÷ (1 – LUN overhead space %))

The host log directory sizing formula assumes the following:

- A system database backup that does not include the tempdb database
- A 10% LUN overhead spacePlace the host log directory on a dedicated volume or LUN. The amount of data in the host log directory depends on the size of the backups and the number of days that backups are retained.

The screenshot displays the SnapCenter interface. On the left, under 'Managed Hosts', there is a search bar and a table with one entry: 'RDSAMAZ-FFIDFMR.rdscustomval.com'. On the right, the 'Host Details' panel shows the following information: Host Name: RDSAMAZ-FFIDFMR.rdscustomval.com; Host IP: 10.0.1.56; Overall Status: Configure log directory (indicated by an orange dot); Host Type: Windows; System: Stand-alone; Credentials: SCAdmin; Plug-ins: SnapCenter Plug-ins package 4.6.0.6965 for Windows, with sub-items for Microsoft Windows and Microsoft SQL Server. At the bottom of the details panel, there is a 'More Options' section with a gear icon and a list of options: Port, gMSA, Install Path, Add Plug-Ins... Below this are three buttons: 'Submit' (highlighted in blue), 'Cancel', and 'Reset'.

If the LUNs have already been provisioned, you can select the mount point to represent the host log directory.



Now you are ready to perform backup, restore and clone operations for SQL Server.

Backup database with SnapCenter

After placing the database and log files on the FSx ONTAP LUNs, SnapCenter can be used to back up the databases. The following processes are used to create a full backup.

Best Practices

- In SnapCenter terms, RPO can be identified as the backup frequency, for example, how frequently you want to schedule the backup so that you can reduce the loss of data to up to few minutes. SnapCenter allows you to schedule backups as frequently as every five minutes. However, there might be a few instances in which a backup might not complete within five minutes during peak transaction times or when the rate of change of data is more in the given time. A best practice is to schedule frequent transaction log backups instead of full backups.
- There are numerous approaches to handle the RPO and RTO. One alternative to this backup approach is to have separate backup policies for data and logs with different intervals. For example, from SnapCenter, schedule log backups in 15-minute intervals and data backups in 6-hour intervals.
- Use a resource group for a backup configuration for Snapshot optimization and the number of jobs to be managed.
 1. Select **Resources**, and then select **Microsoft SQL Server** *on the drop-down menu on the top left. Select ***Refresh Resources**.

Name	Instance	Host	Last Backup	Overall Status	Type
DWConfiguration	RDSAMAZ-FFDFMR	RDSAMAZ-FFDFMR.us-east-1.amazonaws.com		Not available for backup	User database
DWdiagnostics	RDSAMAZ-FFDFMR	RDSAMAZ-FFDFMR.us-east-1.amazonaws.com		Not available for backup	User database
DWQueue	RDSAMAZ-FFDFMR	RDSAMAZ-FFDFMR.us-east-1.amazonaws.com		Not available for backup	User database
master	RDSAMAZ-FFDFMR	RDSAMAZ-FFDFMR.us-east-1.amazonaws.com		Not available for backup	System database
model	RDSAMAZ-FFDFMR	RDSAMAZ-FFDFMR.us-east-1.amazonaws.com		Not available for backup	System database
msdb	RDSAMAZ-FFDFMR	RDSAMAZ-FFDFMR.us-east-1.amazonaws.com		Not available for backup	System database
SeattleRetail	RDSAMAZ-FFDFMR	RDSAMAZ-FFDFMR.us-east-1.amazonaws.com		Not prepared	User database
tempdb	RDSAMAZ-FFDFMR	RDSAMAZ-FFDFMR.us-east-1.amazonaws.com		Not available for backup	System database

2. Select the database to be backed up, then select **Next** and **(*)** to add the policy if one has not been created. Follow the ***New SQL Server Backup Policy** to create a new policy.

Name
DWConfiguration
DWdiagnostics
DWQueue
master
model
msdb
SeattleRetail
tempdb



Select one or more policies and configure schedules

Full Backup i

Configure schedules for selected policies

Policy	Applied Schedules	Configure Schedules
Full Backup	None	To schedule operations select a policy that has the appropriate schedule associated, or modify the selected policy to allow schedules.

3. Select the verification server if necessary. This server is the server that SnapCenter runs DBCC CHECKDB after a full backup has been created. Click **Next** for notification, and then select **Summary** to review. After reviewing, click **Finish**.

Name
DWConfiguration
DWDiagnostics
DWQueue
master
model
msdb
SeattleRetail
tempdb

1 Resource 2 Policies 3 Verification 4 Notification 5 Summary

Select the verification servers

Verification server:

Configure verification schedules

Policy	Schedule Type	Applied Schedules	Configure Schedules
There is no match for your search or data is not available.			

4. Click **Back up Now** to test the backup. In the pop-up windows, select **Backup**.

Backup

Create a backup for the selected resource

Resource Name:

Policy: ⓘ

Verify after backup

5. Select **Monitor** to verify that the backup has been completed.

ID	Status	Name	Start date	End date	Owner
58	✓	Backup of Resource Group 'RDSMANZ-FIDMML-SeattleRetail' with policy 'Full Backup'	03/29/2022 1:47:30 AM	03/29/2022 1:47:41 AM	RDS\CUSTOMVAJ\Administrator
59	✓	Create Resource Group 'RDSMANZ-FIDMML-SeattleRetail'	03/29/2022 1:45:24 AM	03/29/2022 1:45:26 AM	RDS\CUSTOMVAJ\Administrator
60	✓	Create Policy 'Full Backup'	03/29/2022 1:41:37 AM	03/29/2022 1:41:40 AM	RDS\CUSTOMVAJ\Administrator
61	✓	Discover resources for all hosts	03/29/2022 1:38:12 AM	03/29/2022 1:38:17 AM	RDS\CUSTOMVAJ\Administrator

Best Practices


- Backup the transaction log backup from SnapCenter so that during the restoration process, SnapCenter can read all the backup files and restore in sequence automatically.
- If third party products are used for backup, select Copy backup in SnapCenter to avoid log sequence issues, and test the restore functionality before rolling into production.

Restore database with SnapCenter

One of the major benefits of using FSx ONTAP with SQL Server on EC2 is its ability to perform fast and granular restore at each database level.

Complete the following steps to restore an individual database to a specific point in time or up to the minute with SnapCenter.

1. Select Resources and then select the database that you would like to restore.



2. Select the backup name that the database needs to be restored from and then select restore.
3. Follow the **Restore** pop-up windows to restore the database.
4. Select **Monitor** to verify that the restore process is successful.



Considerations for an instance with a large number of small-to-large size databases

SnapCenter can back up a large number of sizeable databases in an instance or group of instances within a resource group. The size of a database is not the major factor in backup time. The duration of a backup can vary depending on number of LUNs per volume, the load on Microsoft SQL Server, the total number of databases per instance, and, specifically, the I/O bandwidth and usage. While configuring the policy to back up databases from an instance or resource group, NetApp recommends that you restrict the maximum database backed up per Snapshot copy to 100 per host. Make sure the total number of Snapshot copies does not exceed the 1,023-copy limit.

NetApp also recommends that you limit the backup jobs running in parallel by grouping the number of databases instead of creating multiple jobs for each database or instance. For optimal performance of the backup duration, reduce the number of backup jobs to a number that can back up around 100 or fewer databases at a time.

As previously mentioned, I/O usage is an important factor in the backup process. The backup process must wait to quiesce until all the I/O operations on a database are complete. Databases with highly intensive I/O operations should be deferred to another backup time or should be isolated from other backup jobs to avoid affecting other resources within the same resource group that are to be backed up.

For an environment that has six Microsoft SQL Server hosts hosting 200 databases per instance, assuming four LUNs per host and one LUN per volume created, set the full backup policy with the maximum databases backed up per Snapshot copy to 100. Two hundred databases on each instance are laid out as 200 data files distributed equally on two LUNs, and 200 log files are distributed equally on two LUNs, which is 100 files per LUN per volume.

Schedule three backup jobs by creating three resource groups, each grouping two instances that include a total of 400 databases.

Running all three backup jobs in parallel backs up 1,200 databases simultaneously. Depending on the load on the server and I/O usage, the start and end time on each instance can vary. In this instance, a total of 24 Snapshot copies are created.

In addition to the full backup, NetApp recommends that you configure a transaction log backup for critical databases. Make sure that the database property is set to full recovery model.

Best practices

1. Do not include the tempdb database in a backup because the data it contains is temporary. Place tempdb on a LUN or an SMB share that is in a storage system volume in which Snapshot copies will not be created.
2. A Microsoft SQL Server instance with a high I/O intensive application should be isolated in a different backup job to reduce the overall backup time for other resources.
3. Limit the set of databases to be simultaneously backed up to approximately 100 and stagger the remaining set of database backups to avoid a simultaneous process.
4. Use the Microsoft SQL Server instance name in the resource group instead of multiple databases because whenever new databases are created in Microsoft SQL Server instance, SnapCenter automatically considers a new database for backup.
5. If you change the database configuration, such as changing the database recovery model to the full recovery model, perform a backup immediately to allow up-to-the-minute restore operations.
6. SnapCenter cannot restore transaction log backups created outside of SnapCenter.
7. When cloning FlexVol volumes, make sure that you have sufficient space for the clone metadata.

8. When restoring databases, make sure that sufficient space is available on the volume.
9. Create a separate policy to manage and back up system databases at least once a week.

Cloning databases with SnapCenter

To restore a database onto another location on a dev or test environment or to create a copy for business analysis purposes, the NetApp best practice is to leverage the cloning methodology to create a copy of the database on the same instance or an alternate instance.

The cloning of databases that are 500GB on an iSCSI disk hosted on a FSx for ONTAP environment typically takes less than five minutes. After cloning is complete, the user can then perform all the required read/write operation on the cloned database. Most of the time is consumed for disk scanning (diskpart). The NetApp cloning procedure typically take less than 2 minutes regardless of the size of the databases.

The cloning of a database can be performed with the dual method: you can create a clone from the latest backup or you can use clone life-cycle management through which the latest copy can be made available on the secondary instance.

SnapCenter allows you to mount the clone copy on the required disk to maintain the format of the folder structure on the secondary instance and continue to schedule backup jobs.

Clone databases to the new database name in the same instance

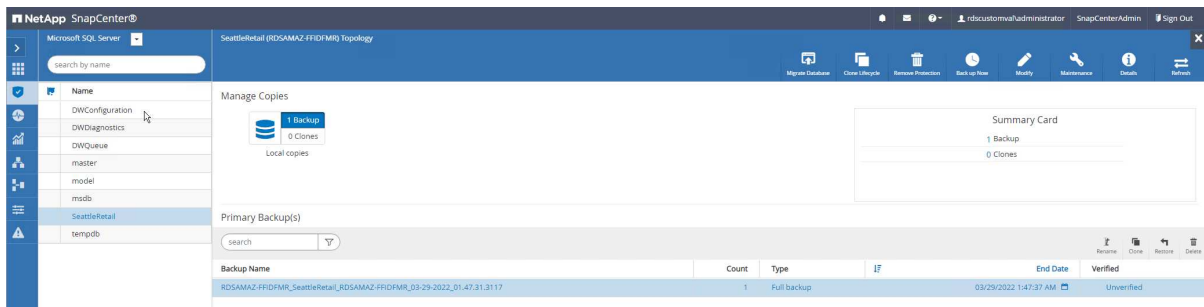
The following steps can be used to clone databases to the new database name in the same SQL server instance running on EC2:

1. Select Resources and then the database that need to be cloned.
2. Select the backup name that you would like to clone and select Clone.
3. Follow the clone instructions from the backup windows to finish the clone process.
4. Select Monitor to make sure that cloning is completed.


Clone databases into the new SQL Server instance running on EC2

The following steps are used to clone databases to the new SQL server instance running on EC2:

1. Create a new SQL Server on EC2 in the same VPC.
2. Enable the iSCSI protocol and MPIO, and then setup the iSCSI connection to FSx for ONTAP by following step 3 and 4 in the section “Create volumes and LUNs for SQL Server.”
3. Add a new SQL Server on EC2 into SnapCenter by following step 3 in the section “Installing and setup for SnapCenter.”
4. Select Resource > View Instance, and then select Refresh Resource.
5. Select Resources, and then the database that you would like to clone.
6. Select the backup name that you would like to clone, and then select Clone.



7. Follow the Clone from Backup instructions by providing the new SQL Server instance on EC2 and instance name to finish the clone process.
8. Select Monitor to make sure that cloning is completed.



To learn more about this process, watch the following video:

[Clone databases into the new SQL Server instance running on EC2](#)

Appendices

Appendix A: YAML file for use in Cloud Formation Template

The following .yaml file can be used with the Cloud Formation Template in AWS Console.

- <https://github.com/NetApp/fsxn-iscsisetup-cft>

To automate iSCSI LUN creation and NetApp SnapCenter installation with PowerShell, clone the repo from [this GitHub link](#).

Appendix B: Powershell scripts for provisioning volumes and LUNs

The following script is used to provision volumes and LUNs and also to set up iSCSI based on the instruction provided above. There are two PowerShell scripts:

- `_EnableMPIO.ps1`

```
Function Install_MPIO_ssh {
    $hostname = $env:COMPUTERNAME
    $hostname = $hostname.Replace('-', '_')

    #Add schedule action for the next step
    $path = Get-Location
    $path = $path.Path + '\2_CreateDisks.ps1'
    $arg = '-NoProfile -WindowStyle Hidden -File ' + $path
    $schAction = New-ScheduledTaskAction -Execute "Powershell.exe"
-Argument $arg
    $schTrigger = New-ScheduledTaskTrigger -AtStartup
    $schPrincipal = New-ScheduledTaskPrincipal -UserId "NT AUTHORITY
\SYSTEM" -LogonType ServiceAccount -RunLevel Highest
    $return = Register-ScheduledTask -Action $schAction -Trigger
$schTrigger -TaskName "Create Vols and LUNs" -Description "Scheduled
Task to run configuration Script At Startup" -Principal $schPrincipal
    #Install -Module PosH-SSH
    Write-host 'Enable MPIO and SSH for PowerShell' -ForegroundColor
Yellow
    $return = Find-PackageProvider -Name 'Nuget' -ForceBootstrap
-IncludeDependencies
    $return = Find-Module PoSH-SSH | Install-Module -Force
    #Install Multipath-IO with PowerShell using elevated privileges in
Windows Servers
    Write-host 'Enable MPIO' -ForegroundColor Yellow
    $return = Install-WindowsFeature -name Multipath-IO -Restart
}
Install_MPIO_ssh
Remove-Item -Path $MyInvocation.MyCommand.Source
```

- `_CreateDisks.ps1`

```
....
#Enable MPIO and Start iSCSI Service
Function PrepISCSI {
    $return = Enable-MSDSMAutomaticClaim -BusType iSCSI
    #Start iSCSI service with PowerShell using elevated privileges in
Windows Servers
```

```

$return = Start-service -Name msiscsi
$return = Set-Service -Name msiscsi -StartupType Automatic
}
Function Create_igroup_vols_luns ($fsxN){
    $hostname = $env:COMPUTERNAME
    $hostname = $hostname.Replace('-', '_')
    $volsluns = @()
    for ($i = 1;$i -lt 10;$i++){
        if ($i -eq 9){
            $volsluns
+=(@{volname=('v_'+$hostname+'_log');volsize=$fsxN.logvolsize;lunname=(
'l_'+$hostname+'_log');lunsize=$fsxN.loglunsize})
        } else {
            $volsluns
+=(@{volname=('v_'+$hostname+'_data'+[string]$i);volsize=$fsxN.datavols
ize;lunname=('l_'+$hostname+'_data'+[string]$i);lunsize=$fsxN.datalunsi
ze})
        }
    }
    $secStringPassword = ConvertTo-SecureString $fsxN.password
-AsPlainText -Force
    $credObject = New-Object System.Management.Automation.PSCredential
($fsxN.login, $secStringPassword)
    $igroup = 'igrp_'+$hostname
    #Connect to FSx N filesystem
    $session = New-SSHSession -ComputerName $fsxN.svmip -Credential
$credObject -AcceptKey:$true
    #Create igroup
    Write-host 'Creating igroup' -ForegroundColor Yellow
    #Find Windows initiator Name with PowerShell using elevated
privileges in Windows Servers
    $initport = Get-InitiatorPort | select -ExpandProperty NodeAddress
    $sshcmd = 'igroup create -igroup ' + $igroup + ' -protocol iscsi
-ostype windows -initiator ' + $initport
    $ret = Invoke-SSHCommand -Command $sshcmd -SSHSession $session
    #Create vols
    Write-host 'Creating Volumes' -ForegroundColor Yellow
    foreach ($vollun in $volsluns){
        $sshcmd = 'vol create ' + $vollun.volname + ' -aggregate agr1
-size ' + $vollun.volsize #+ ' -vserver ' + $vserver
        $return = Invoke-SSHCommand -Command $sshcmd -SSHSession
$session
    }
    #Create LUNs and mapped LUN to igroup
    Write-host 'Creating LUNs and map to igroup' -ForegroundColor
Yellow

```



```

    foreach ($vollun in $volsluns){
        $ssshcmd = "lun create -path /vol/" + $vollun.volname + "/" +
$vollun.lunname + " -size " + $vollun.lunsize + " -ostype Windows_2008
" #-vserver " +$vserver
        $return = Invoke-SSHCommand -Command $ssshcmd -SSHSession
$session
        #map all luns to igroup
        $ssshcmd = "lun map -path /vol/" + $vollun.volname + "/" +
$vollun.lunname + " -igroup " + $igroup
        $return = Invoke-SSHCommand -Command $ssshcmd -SSHSession
$session
    }
}
Function Connect_iSCSI_to_SVM ($TargetPortals){
    Write-host 'Online, Initialize and format disks' -ForegroundColor
Yellow
    #Connect Windows Server to svm with iSCSI target.
    foreach ($TargetPortal in $TargetPortals) {
        New-IscsiTargetPortal -TargetPortalAddress $TargetPortal
        for ($i = 1; $i -lt 5; $i++){
            $return = Connect-IscsiTarget -IsMultipathEnabled $true
-IsPersistent $true -NodeAddress (Get-iscsiTarget | select
-ExpandProperty NodeAddress)
        }
    }
}
Function Create_Partition_Format_Disks{

    #Create Partion and format disk
    $disks = Get-Disk | where PartitionStyle -eq raw
    foreach ($disk in $disks) {
        $return = Initialize-Disk $disk.Number
        $partition = New-Partition -DiskNumber $disk.Number
-AssignDriveLetter -UseMaximumSize | Format-Volume -FileSystem NTFS
-AllocationUnitSize 65536 -Confirm:$false -Force
        # $return = Format-Volume -DriveLetter $partition.DriveLetter
-FileSystem NTFS -AllocationUnitSize 65536
    }
}
Function UnregisterTask {
    Unregister-ScheduledTask -TaskName "Create Vols and LUNs"
-Confirm:$false
}
Start-Sleep -s 30
$fsxN = @{svmip ='198.19.255.153';login =
'vsadmin';password='net@pp11';datavolsize='10GB';datalunsize='8GB';logv

```

```
olsize='8GB';loglunsize='6GB'}
$TargetPortals = ('10.2.1.167', '10.2.2.12')
PrepISCSI
Create_igroup_vols_luns $fsxN
Connect_iSCSI_to_SVM $TargetPortals
Create_Partition_Format_Disks
UnregisterTask
Remove-Item -Path $MyInvocation.MyCommand.Source
....
```

Run the file `EnableEMPIO.ps1` first and the second script executes automatically after the server has been rebooted. These PowerShell scripts can be removed after they have been executed due to credential access to the SVM.

Where to find additional information

- Amazon FSx for NetApp ONTAP

<https://docs.aws.amazon.com/fsx/latest/ONTAPGuide/what-is-fsx-ontap.html>

- Getting Started with FSx for NetApp ONTAP

<https://docs.aws.amazon.com/fsx/latest/ONTAPGuide/getting-started.html>

- Overview of the SnapCenter interface

<https://www.youtube.com/watch?v=IVEBF4kV6Ag&t=0s>

- Tour through SnapCenter navigation pane options

https://www.youtube.com/watch?v=_IDKt-koySQ

- Setup SnapCenter 4.0 for SQL Server plug-in

<https://www.youtube.com/watch?v=MopbUFSdHKE>

- How to back up and restore databases using SnapCenter with SQL Server plug-in

https://www.youtube.com/watch?v=K343qPD5_Ys

- How to clone a database using SnapCenter with SQL Server plug-in

<https://www.youtube.com/watch?v=ogEc4DkGv1E>

TR-4897: SQL Server on Azure NetApp Files - Real Deployment View

This document covers a real-time deployment of SQL Server Always On availability group (AOAG) on Azure NetApp Files leveraging Azure Virtual Machines.

Niyaz Mohamed, NetApp

IT organizations face constant change. Gartner reports nearly 75% of all databases will require cloud-based storage by 2022. As a leading relational database management system (RDBMS), Microsoft SQL Server is the go-to choice for Windows platform-designed applications and organizations that rely on SQL Server for everything from enterprise resource planning (ERP) to analytics to content management. SQL Server has helped to revolutionize the way enterprises manage massive data sets and power their applications to meet the schema and query performance demands.

Most IT organizations follow a cloud-first approach. Customers in a transformation phase evaluate their current IT landscape and then migrate their database workloads to the cloud based on an assessment and discovery exercise. Some factors driving customers toward cloud migration include elasticity/burst, data center exit, data center consolidation, end-of-life scenarios, mergers, acquisitions, and so on. The reason for migration can vary based on each organization and their respective business priorities. When moving to the cloud, choosing the right cloud storage is very important in order to unleash the power of SQL Server database cloud deployment.

Use case

Moving the SQL Server estate to Azure and integrating SQL Server with Azure's vast array of platform-as-a-service (PaaS) features such as Azure Data Factory, Azure IoT Hub, and Azure Machine Learning creates tremendous business value to support digital transformation. Adopting the cloud also enables the respective business unit to focus on productivity and delivering new features and enhancements faster (DevTest use case) than relying on the CAPEX model or traditional private cloud models. This document covers a real-time deployment of SQL Server Always On availability group (AOAG) on Azure NetApp Files leveraging Azure Virtual Machines.

Azure NetApp Files provides enterprise-grade storage with continuously available file shares. Continuously available shares are required by SQL Server production databases on SMB file share to make sure that the node always has access to the database storage, including during disruptive scenarios such as controller upgrades or failures. Continuously available file shares eliminate the need to replicate data between storage nodes. Azure NetApp Files uses SMB 3.0 scale-out, persistent handles, and transparent failover to support nondisruptive operations (NDOs) for planned and unplanned downtime events, including many administrative tasks.

When planning cloud migrations, you should always evaluate the best approach to use. The most common and easiest approach for application migration is rehosting (also known as lift and shift). The example scenario provided in this document uses the rehosting method. SQL Server on Azure virtual machines with Azure NetApp Files allows you to use full versions of SQL Server in the cloud without having to manage on-premises hardware. SQL Server virtual machines (VMs) also simplify licensing costs when you pay as you go and provides elasticity and bursting capabilities for development, test, and estate refresh scenarios.

Factors to consider

This section describes the different issues you should consider when Azure NetApp Files with SQL Server in the cloud.

VM performance

Selecting the right VM size is important for optimal performance of a relational database in a public cloud. Microsoft recommends that you continue using the same database performance-tuning options that are applicable to SQL Server in on-premises server environments. Use [memory-optimized](#) VM sizes for the best performance of SQL Server workloads. Collect the performance data of existing deployment to identify the RAM and CPU utilization while choosing the right instances. Most deployments choose between the D, E, or M series.

Notes:

- For the best performance of SQL Server workloads, use memory-optimized VM sizes.
- NetApp and Microsoft recommend that you identify the storage performance requirements before choosing the instance type with the appropriate memory-to-vCore ratio. This also helps select a lower-instance type with the right network bandwidth to overcome storage throughput limits of the VM.

VM redundancy

To increase redundancy and high availability, SQL Server VMs should either be in the same [availability set](#) or different [availability zones](#). When creating Azure VMs, you must choose between configuring availability sets versus availability zones; an Azure VM cannot participate in both.

High availability

For high availability, configuring SQL Server AOAG or Always On Failover Cluster Instance (FCI) is the best option. For AOAG, this involves multiple instances of SQL Server on Azure Virtual Machines in a virtual network. If high availability is required at the database level, consider configuring SQL Server availability groups.

Storage configuration

Microsoft SQL Server can be deployed with an SMB file share as the storage option. Starting with SQL Server 2012, system databases (master, model, msdb, or tempdb), and user databases can be installed with Server Message Block (SMB) file server as a storage option. This applies to both SQL Server stand-alone and SQL Server FCI.



File share storage for SQL Server databases should support continuously available property. This provides uninterrupted access to the file-share data.

Azure NetApp Files provides high performing file storage to meet any demanding workload, and it reduces SQL Server TCO as compared to block storage solutions. With block storage, VMs have imposed limits on I/O and bandwidth for disk operations; network bandwidth limits alone are applied against Azure NetApp Files. In other words, no VM-level I/O limits are applied to Azure NetApp Files. Without these I/O limits, SQL Server running on smaller VMs connected to Azure NetApp Files can perform as well as SQL Server running on much larger VMs. Azure NetApp Files reduce SQL Server deployment costs by reducing compute and software licensing costs. For detailed cost analysis and performance benefits of using Azure NetApp Files for SQL Server deployment, see the [Benefits of using Azure NetApp Files for SQL Server deployment](#).

Benefits

The benefits of using Azure NetApp Files for SQL Server include the following:

- Using Azure NetApp Files allows you to use smaller instances, thus reducing compute cost.
- Azure NetApp Files also reduces software licensing costs, which reduce the overall TCO.
- Volume reshaping and dynamic service level capability optimizes cost by sizing for steady-state workloads and avoiding overprovisioning.

Notes:

- To increase redundancy and high availability, SQL Server VMs should either be in the same [availability set](#) or in different [availability zones](#). Consider file path requirements if user-defined data files are required; in which case, select SQL FCI over SQL AOAG.
- The following UNC path is supported: `\\ANFSMB-b4ca.anf.test\SQLDB` and `\\ANFSMB-b4ca.anf.test\SQLDB\`.

- The loopback UNC path is not supported.
- For sizing, use historic data from your on-premises environment. For OLTP workloads, match the target IOPS with performance requirements using workloads at average and peak times along with the disk reads/sec and disk writes/sec performance counters. For data warehouse and reporting workloads, match the target throughput using workloads at average and peak times and the disk read bytes/sec and disk write bytes/sec. Average values can be used in conjunction with volume reshaping capabilities.

Create continuously available shares

Create continuously available shares with the Azure portal or Azure CLI. In the portal, select the Enable Continuous Availability property option. For the Azure CLI, specify the share as a continuously available share by using the `az netappfiles volume create` with the `smb-continuously-avl` option set to `$True`. To learn more about creating a new, continuous availability-enabled volume, see [Creating a Continuously Available Share](#).

Notes:

- Enable continuous availability for the SMB volume as shown in the following image.
- If a non-administrator domain account is used, make sure the account has the required security privilege assigned.
- Set the appropriate permissions at the share level and proper file-level permissions.
- A continuously available property cannot be enabled on existing SMB volumes. To convert an existing volume to use a continuously available share, use NetApp Snapshot technology. For more information, see [Convert existing SMB volumes to use Continuous Availability](#).

Create a volume

...

×

Basics **Protocol** Tags Review + create

Configure access to your volume.

Access

Protocol type NFS SMB Dual-protocol (NFSv3 and SMB)

Configuration

Active Directory * ⓘ ▼

Share name * ⓘ

Enable Continuous Availability ⓘ

Review + create

< Previous


Next : Tags >

Performance

Azure NetApp Files supports three service levels: Standard (16MBps per terabyte), Premium (64MBps per terabyte), and Ultra (128MBps per terabyte). Provisioning the right volume size is important for optimal performance of the database workload. With Azure NetApp Files, volume performance and the throughput limit are based on a combination of the following factors:

- The service level of the capacity pool to which the volume belongs
- The quota assigned to the volume
- The quality of service (QoS) type (auto or manual) of the capacity pool

For more information, see [Service levels for Azure NetApp Files](#).




Performance validation

As with any deployment, testing the VM and storage is critical. For storage validation, tools such as HammerDB, Apploader, the [SQL Server storage benchmark \(SB\) tool](#), or any custom script or FIO with the appropriate read/write mix should be used. Keep in mind however that most SQL Server workloads, even busy OLTP workloads, are closer to 80%–90% read and 10%–20% write.

To showcase performance, a quick test was performed against a volume using premium service levels. In this test, the volume size was increased from 100GB to 2TB on the fly without any disruption to application access and zero data migration.

ANF Premium Tier Quotas




Here is another example of real time performance testing with HammerDB performed for the deployment covered in this paper. For this testing, we used a small instance with eight vCPUs, a 500GB Premium SSD, and a 500GB SMB Azure NetApp Files volume. HammerDB was configured with 80 warehouses and eight users.

The following chart shows that Azure NetApp Files was able to deliver 2.6x the number of transactions per minute at 4x lower latency when using a comparable sized volume (500GB).

An additional test was performed by resizing to a larger instance with 32x vCPUs and a 16TB Azure NetApp Files volume. There was a significant increase in transactions per minute with consistent 1ms latency. HammerDB was configured with 80 warehouses and 64 users for this test.

SQL Hammer DB Results



Cost optimization

Azure NetApp Files allows nondisruptive, transparent volume resizing and the ability to change the service levels with zero downtime and no effect on applications. This is a unique capability allowing dynamic cost management that avoids the need to perform database sizing with peak metrics. Rather, you can use steady state workloads, which avoids upfront costs. The volume reshaping and dynamic service-level change allows you to adjust the bandwidth and service level of Azure NetApp Files volumes on demand almost instantaneously without pausing I/O, while retaining data access.

Azure PaaS offerings such as LogicApp or Functions can be used to easily resize the volume based on a specific webhook or alert rule trigger to meet the workload demands while dynamically handling the cost.

For example, consider a database that needs 250MBps for steady state operation; however, it also requires a peak throughput of 400MBps. In this case, the deployment should be performed with a 4TB volume within the Premium service level to meet the steady-state performance requirements. To handle the peak workload, increase the volume size using Azure functions to 7TB for that specific period, and then downsize the volume to make the deployment cost effective. This configuration avoids overprovisioning of the storage.


Real-time, high-level reference design

This section covers a real-time deployment of a SQL database estate in an AOAG configuration using an Azure NetApp Files SMB volume.


- Number of nodes: 4
- Number of databases: 21
- Number of availability groups: 4
- Backup retention: 7 days
- Backup archive: 365 days



Deploying FCI with SQL Server on Azure virtual machines with an Azure NetApp Files share provides a cost-efficient model with a single copy of the data. This solution can prevent add-file operation issues if the file path differs from the secondary replica.



The following image shows the databases within AOAG spread across the nodes.



Data layout

The user database files (.mdf) and user database transaction log files (.ldf) along with tempDB are stored on the same volume. The service level is Ultra.

The configuration consists of four nodes and four AGs. All 21 databases (part of Dynamic AX, SharePoint, RDS connection broker, and indexing services) are stored on the Azure NetApp Files volumes. The databases are balanced between the AOAG nodes to use the resources on the nodes effectively. Four D32 v3 instances are added in the WSFC, which participates in the AOAG configuration. These four nodes are provisioned in the Azure virtual network and are not migrated from on-premises.

Notes:

- If the logs require more performance and throughput depending on the nature of the application and the queries executed, the database files can be placed on the Premium service level, and the logs can be stored at the Ultra service level.
- If the tempdb files have been placed on Azure NetApp Files, then the Azure NetApp Files volume should be separated from the user database files. Here is an example distribution of the database files in AOAG.

Notes:

- To retain the benefits of Snapshot copy-based data protection, NetApp recommends not combining data and log data into the same volume.
- An add-file operation performed on the primary replica might fail on the secondary databases if the file path of a secondary database differs from the path of the corresponding primary database. This can happen if the share path is different on primary and secondary nodes (due to different computer accounts). This failure could cause the secondary databases to be suspended. If the growth or performance pattern cannot be predicted and the plan is to add files later, a SQL Server failover cluster with Azure NetApp Files is an acceptable solution. For most deployments, Azure NetApp Files meets the performance requirements.

Migration

There are several ways to migrate an on-premises SQL Server user database to SQL Server in an Azure virtual machine. The migration can be either online or offline. The options chosen depend on the SQL Server version, business requirements, and the SLAs defined within the organization. To minimize downtime during the database migration process, NetApp recommends using either the AlwaysOn option or the transactional replication option. If it is not possible to use these methods, you can migrate the database manually.

The simplest and most thoroughly tested approach for moving databases across machines is backup and restore. Typically, you can start with a database backup followed by a copy of the database backup into Azure. You can then restore the database. For the best data transfer performance, migrate the database files into the Azure VM using a compressed backup file. The high-level design referenced in this document uses the backup approach to Azure file storage with Azure file sync and then restore to Azure NetApp files.



Azure Migrate can be used to discover, assess, and migrate SQL Server workloads.

To perform a migration, complete the following high-level steps:

1. Based on your requirements, set up connectivity.
2. Perform a full database backup to an on-premises file-share location.
3. Copy the backup files to an Azure file share with Azure file sync.
4. Provision the VM with the desired version of SQL Server.
5. Copy the backup files to the VM by using the `copy` command from a command prompt.
6. Restore the full databases to SQL Server on Azure virtual machines.



To restore 21 databases, it took approximately nine hours. This approach is specific to this scenario. However, other migration techniques listed below can be used based on your situation and requirements.

Other migration options to move data from an on-premises SQL Server to Azure NetApp Files include the following:

- Detach the data and log files, copy them to Azure Blob storage, and then attach them to SQL Server in the Azure VM with an ANF file share mounted from the URL.
- If you are using Always On availability group deployment on-premises, use the [Add Azure Replica Wizard](#) to create a replica in Azure and then perform failover.
- Use SQL Server [transactional replication](#) to configure the Azure SQL Server instance as a subscriber, disable replication, and point users to the Azure database instance.
- Ship the hard drive using the Windows Import/Export Service.

Backup and recovery

Backup and recovery are an important aspect of any SQL Server deployment. It is mandatory to have the appropriate safety net to quickly recover from various data failure and loss scenarios in conjunction with high availability solutions such as AOAG. SQL Server Database Quiesce Tool, Azure Backup (streaming), or any third-party backup tool such as Commvault can be used to perform an application-consistent backup of the databases,

Azure NetApp Files Snapshot technology allows you to easily create a point-in-time (PiT) copy of the user databases without affecting performance or network utilization. This technology also allows you to restore a Snapshot copy to a new volume or quickly revert the affected volume to the state it was in when that Snapshot copy was created by using the revert volume function. The Azure NetApp Files snapshot process is very quick and efficient, which allows for multiple daily backups, unlike the streaming backup offered by Azure backup. With multiple Snapshot copies possible in a given day, the RPO and RTO times can be significantly reduced. To add application consistency so that data is intact and properly flushed to the disk before the Snapshot copy is taken, use the SQL Server database quiesce tool ([SCSQLAPI tool](#); access to this link requires NetApp SSO login credentials). This tool can be executed from within PowerShell, which quiesces the SQL Server database and in turn can take the application-consistent storage Snapshot copy for backups.

*Notes: *

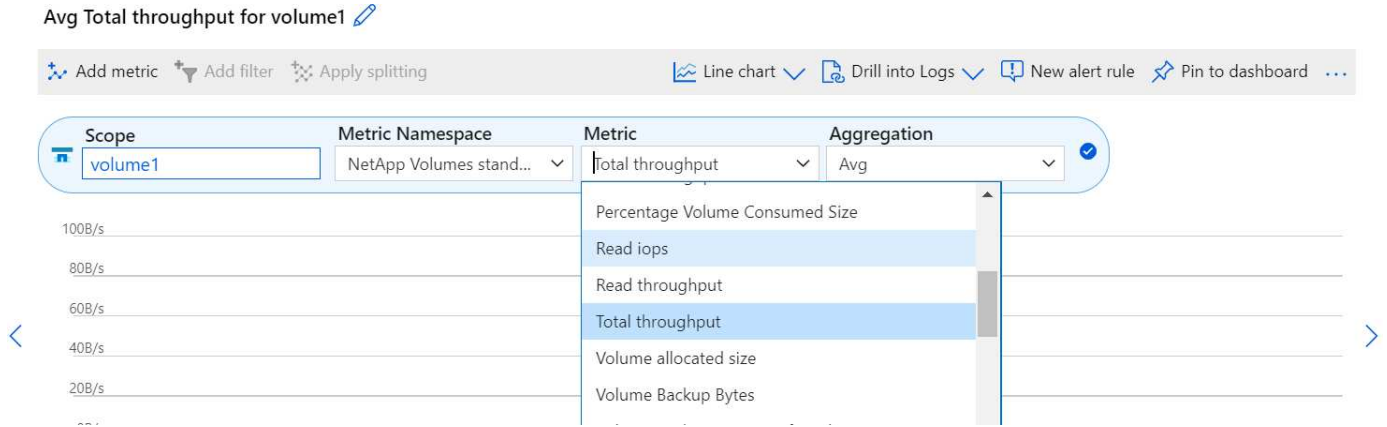
- The SCSQLAPI tool only supports the 2016 and 2017 versions of SQL Server.
- The SCSQLAPI tool only works with one database at a time.
- Isolate the files from each database by placing them onto a separate Azure NetApp Files volume.

Because of SCSQL API's vast limitations, [Azure Backup](#) was used for data protection in order to meet the SLA requirements. It offers a stream-based backup of SQL Server running in Azure Virtual Machines and Azure NetApp Files. Azure Backup allows a 15-minute RPO with frequent log backups and PiT recovery up to one second.

Monitoring

Azure NetApp Files is integrated with Azure Monitor for the time series data and provides metrics on allocated storage, actual storage usage, volume IOPS, throughput, disk read bytes/sec, disk write bytes/sec, disk reads/sec and disk writes/sec, and associated latency. This data can be used to identify bottlenecks with alerting and to perform health checks to verify that your SQL Server deployment is running in an optimal configuration.

In this HLD, ScienceLogic is used to monitor Azure NetApp Files by exposing the metrics using the appropriate service principal. The following image is an example of the Azure NetApp Files Metric option.



DevTest using thick clones

With Azure NetApp Files, you can create instantaneous copies of databases to test functionality that should be implemented by using the current database structure and content during the application development cycles, to use the data extraction and manipulation tools when populating data warehouses, or to even recover data that was mistakenly deleted or changed. This process does not involve copying data from Azure Blob containers, which makes it very efficient. After the volume is restored, it can be used for read/write operations, which significantly reduces validation and time to market. This needs to be used in conjunction with SCSQLAPI for application consistency. This approach provides yet another continuous cost optimization technique along with Azure NetApp Files leveraging the Restore to New volume option.

Notes:

- The volume created from the Snapshot copy using the Restore New Volume option consumes capacity from the capacity pool.
- You can delete the cloned volumes by using REST or Azure CLI to avoid additional costs (in case the capacity pool must be increased).

Hybrid storage options

Although NetApp recommends using the same storage for all the nodes in SQL Server availability groups, there are scenarios in which multiple storage options can be used. This scenario is possible for Azure NetApp Files in which a node in AOAG is connected with an Azure NetApp Files SMB file share and the second node is connected with an Azure Premium disk. In these instances, make sure that the Azure NetApp Files SMB share is holding the primary copy of the user databases and the Premium disk is used as the secondary copy.

Notes:

- In such deployments, to avoid any failover issues, make sure that continuous availability is enabled on the SMB volume. With no continuously available attribute, the database can fail if there is any background maintenance at the storage layer.
- Keep the primary copy of the database on the Azure NetApp Files SMB file share.

Business continuity

Disaster recovery is generally an afterthought in any deployment. However, disaster recovery must be addressed during the initial design and deployment phase to avoid any impact to your business. With Azure NetApp Files, the cross-region replication (CRR) functionality can be used to replicate the volume data at the

block level to the paired region to handle any unexpected regional outage. The CRR-enabled destination volume can be used for read operations, which makes it an ideal candidate for disaster recovery simulations. In addition, the CRR destination can be assigned with the lowest service level (for instance, Standard) to reduce the overall TCO. In the event of a failover, replication can be broken, which makes the respective volume read/write capable. Also, the service level of the volume can be changed by using the dynamic service level functionality to significantly reduce disaster recovery cost. This is another unique feature of Azure NetApp Files with block replication within Azure.

Long-term Snapshot copy archive

Many organizations must perform long-term retention of snapshot data from database files as a mandatory compliance requirement. Although this process is not used in this HLD, it can be easily accomplished by using a simple batch script using [AzCopy](#) to copy the snapshot directory to the Azure Blob container. The batch script can be triggered based on a specific schedule by using scheduled tasks. The process is straightforward—it includes the following steps:

1. Download the AzCopy V10 executable file. There is nothing to install because it is an `exe` file.
2. Authorize AzCopy by using a SAS token at the container level with the appropriate permissions.
3. After AzCopy is authorized, the data transfer begins.

Notes:

- In batch files, make sure to escape the `%` characters that appear in SAS tokens. This can be done by adding an additional `%` character next to existing `%` characters in the SAS token string.
- The [Secure Transfer Required](#) setting of a storage account determines whether the connection to a storage account is secured with Transport Layer Security (TLS). This setting is enabled by default. The following batch script example recursively copies data from the Snapshot copy directory to a designated Blob container:

```
SET source="Z:\~snapshot"  
echo %source%  
SET  
dest="https://testanfacct.blob.core.windows.net/azcoptst?sp=racwdl&st=2020-10-21T18:41:35Z&se=2021-10-22T18:41:00Z&sv=2019-12-12&sr=c&sig=ZxRUJwFlLXgHS8As7HzXJOaDXXVJ7PxxIX3ACpx56XY%%3D"  
echo %dest%
```

The following example cmd is executed in PowerShell:

```
-recursive
```

```
INFO: Scanning...
INFO: Any empty folders will not be processed, because source and/or
destination doesn't have full folder support
Job b3731dd8-da61-9441-7281-17a4db09ce30 has started
Log file is located at: C:\Users\niyaz\.azcopy\b3731dd8-da61-9441-7281-
17a4db09ce30.log
0.0 %, 0 Done, 0 Failed, 2 Pending, 0 Skipped, 2 Total,
INFO: azcopy.exe: A newer version 10.10.0 is available to download
0.0 %, 0 Done, 0 Failed, 2 Pending, 0 Skipped, 2 Total,
Job b3731dd8-da61-9441-7281-17a4db09ce30 summary
Elapsed Time (Minutes): 0.0333
Number of File Transfers: 2
Number of Folder Property Transfers: 0
Total Number of Transfers: 2
Number of Transfers Completed: 2
Number of Transfers Failed: 0
Number of Transfers Skipped: 0
TotalBytesTransferred: 5
Final Job Status: Completed
```

Notes:

- A similar backup feature for long-term retention will soon be available in Azure NetApp Files.
- The batch script can be used in any scenario that requires data to be copied to a Blob container in any region.

Cost optimization

With volume reshaping and dynamic service level change, which is completely transparent to the database, Azure NetApp Files allows continuous cost optimizations in Azure. This capability is used in this HLD extensively to avoid overprovisioning of additional storage to handle workload spikes.

Resizing the volume can be easily accomplished by creating an Azure function in conjunction with the Azure alert logs.

Conclusion

Whether you are targeting an all-cloud or hybrid cloud with stretch databases, Azure NetApp Files provides excellent options to deploy and manage the database workloads while reducing your TCO by making data requirements seamless to the application layer.

This document covers recommendations for planning, designing, optimizing, and scaling Microsoft SQL Server deployments with Azure NetApp Files, which can vary greatly between implementations. The right solution depends on both the technical details of the implementation and the business requirements driving the project.

Takeaways

The key points of this document include:

- You can now use Azure NetApp Files to host the database and file share witness for SQL Server cluster.
- You can boost the application response times and deliver 99.9999% availability to provide access to SQL Server data when and where it is needed.
- You can simplify the overall complexity of the SQL Server deployment and ongoing management, such as raid striping, with simple and instant resizing.
- You can rely on intelligent operations features to help you deploy SQL Server databases in minutes and speed development cycles.
- If Azure Cloud is the destination, Azure NetApp Files is the right storage solution for optimized deployment.

Where to find additional information

To learn more about the information described in this document, refer to the following website links:

- Solution architectures using Azure NetApp Files

<https://docs.microsoft.com/en-us/azure/azure-netapp-files/azure-netapp-files-solution-architectures>

- Benefits of using Azure NetApp Files for SQL Server deployment

<https://docs.microsoft.com/en-us/azure/azure-netapp-files/solutions-benefits-azure-netapp-files-sql-server>

- SQL Server on Azure Deployment Guide Using Azure NetApp Files

<https://www.netapp.com/pdf.html?item=/media/27154-tr-4888.pdf>

- Fault tolerance, high availability, and resilience with Azure NetApp Files

<https://cloud.netapp.com/blog/azure-anf-blg-fault-tolerance-high-availability-and-resilience-with-azure-netapp-files>

TR-4467: SAP with Microsoft SQL Server on Windows - Best practices using NetApp Clustered Data ONTAP and SnapCenter

Marco Schoen, NetApp

TR-4467 provides customers and partners with best practices for deploying clustered NetApp Data ONTAP in support of SAP Business Suite solutions running in a Microsoft SQL Server on Windows environment.

[TR-4467: SAP with Microsoft SQL Server on Windows - Best practices using NetApp Clustered Data ONTAP and SnapCenter](#)

Modernizing your Microsoft SQL Server environment

Optimize operations and unleash the power of your data - on the premises or in the cloud.

[Modernizing your Microsoft SQL Server environment](#)

TR-4590: Best practice guide for Microsoft SQL Server with ONTAP

Manohar Kulkarni and Pat Sinthusan, NetApp

This document describes best practices and offers insight into design considerations for deploying SQL Server on NetApp storage systems running NetApp ONTAP® software, with the goal of achieving effective and efficient storage deployment and end-to-end data protection and retention planning.

[TR-4590: Best practices guide for Microsoft SQL Server with ONTAP](#)

TR-4764: Best practices for Microsoft SQL Server with NetApp EF-Series

Mitch Blackburn, Pat Sinthusan, NetApp

This best practices guide is intended to help storage administrators and database administrators successfully deploy Microsoft SQL Server on NetApp EF-Series storage.

[TR-4764: Best practices for Microsoft SQL Server with NetApp EF-Series](#)

Open Source Databases

TR-4956: Automated PostgreSQL High Availability Deployment and Disaster Recovery in AWS FSx/EC2

Allen Cao, Niyaz Mohamed, NetApp

This solution provides overview and details for PostgreSQL database deployment and HA/DR setup, failover, resync based on NetApp SnapMirror technology built into FSx ONTAP storage offering and NetApp Ansible automation toolkit in AWS.

Purpose

PostgreSQL is a widely used open-source database that is ranked number four among the top ten most popular database engines by [DB-Engines](#). On one hand, PostgreSQL derives its popularity from its license-free, open-source model while still possessing sophisticated features. On the other hand, because it is open sourced, there is shortage of detailed guidance on production-grade database deployment in the area of high availability and disaster recovery (HA/DR), particularly in the public cloud. In general, it can be difficult to set up a typical PostgreSQL HA/DR system with hot and warm standby, streaming replication, and so on. Testing the HA/DR environment by promoting the standby site and then switching back to the primary can be disruptive to production. There are well documented performance issues on the primary when read workloads are deployed on streaming hot standby.

In this documentation, we demonstrate how you can do away with an application-level PostgreSQL streaming HA/DR solution and build a PostgreSQL HA/DR solution based on AWS FSx ONTAP storage and EC2 compute instances using storage-level replication. The solution creates a simpler and comparable system and delivers equivalent results when compared with traditional PostgreSQL application-level streaming replication for HA/DR.

This solution is built on proven and mature NetApp SnapMirror storage-level replication technology that is available in AWS-native FSX ONTAP cloud storage for PostgreSQL HA/DR. It is simple to implement with an

automation toolkit provided by the NetApp Solutions team. It provides similar functionality while eliminating the complexity and performance drag on the primary site with the application-level streaming-based HA/DR solution. The solution can be easily deployed and tested without affecting the active primary site.

This solution addresses the following use cases:

- Production grade HA/DR deployment for PostgreSQL in the public AWS cloud
- Testing and validating a PostgreSQL workload in the public AWS cloud
- Testing and validating a PostgreSQL HA/DR strategy based on NetApp SnapMirror replication technology

Audience


This solution is intended for the following people:

- The DBA who is interested in deploying PostgreSQL with HA/DR in the public AWS cloud.
- The database solution architect who is interested in testing PostgreSQL workloads in the public AWS cloud.
- The storage administrator who is interested in deploying and managing PostgreSQL instances deployed to AWS FSx storage.
- The application owner who is interested in standing up a PostgreSQL environment in AWS FSx/EC2.

Solution test and validation environment

The testing and validation of this solution was performed in an AWS FSx and EC2 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		
FSx ONTAP storage	Current version	Two FSx HA pairs in the same VPC and availability zone as primary and standby HA clusters
EC2 instance for compute	t2.xlarge/4vCPU/16G	Two EC2 T2 xlarge as primary and standby compute instances
Ansible controller	on-prem Centos VM/4vCPU/8G	A VM to host Ansible automation controller either on-premise or in the cloud
Software		
RedHat Linux	RHEL-8.6.0_HVM-20220503-x86_64-2-Hourly2-GP2	Deployed RedHat subscription for testing
Centos Linux	CentOS Linux release 8.2.2004 (Core)	Hosting Ansible controller deployed in on-premises lab
PostgreSQL	Version 14.5	Automation pulls the latest available version of PostgreSQL from the postgresql.ora yum repo
Ansible	Version 2.10.3	Prerequisites for required collections and libraries installed with requirements playbook

Key factors for deployment consideration

- **PostgreSQL database backup, restore, and recovery.** A PostgreSQL database supports a number of backup methods, such as a logical backup using `pg_dump`, a physical online backup with `pg_basebackup` or a lower-level OS backup command, and storage-level-consistent snapshots. This solution uses NetApp consistency-group snapshots for PostgreSQL database data and WAL volumes backup, restore, and recovery at the standby site. The NetApp consistency-group volume snapshots sequence I/O as it is written to storage and protect the integrity of database data files.
- **EC2 compute instances.** In these tests and validations, we used the AWS EC2 t2.xlarge instance type for the PostgreSQL database compute instance. NetApp recommends using an M5 type EC2 instance as the compute instance for PostgreSQL in deployment because it is optimized for database workloads. The standby compute instance should always be deployed in the same zone as the passive (standby) file system deployed for the FSx HA cluster.
- **FSx storage HA clusters single- or multi-zone deployment.** In these tests and validations, we deployed an FSx HA cluster in a single AWS availability zone. For production deployment, NetApp recommends deploying an FSx HA pair in two different availability zones. A disaster-recovery standby HA pair for business continuity can be set up in a different region if a specific distance is required between the primary and standby. An FSx HA cluster is always provisioned in a HA pair that is sync mirrored in a pair of active-passive file systems to provide storage-level redundancy.
- **PostgreSQL data and log placement.** Typical PostgreSQL deployments share the same root directory or volumes for data and log files. In our tests and validations, we have separated PostgreSQL data and logs into two separate volumes for performance. A soft link is used in the data directory to point to the log directory or volume that hosts PostgreSQL WAL logs and archived WAL logs.
- **PostgreSQL service startup delay timer.** This solution uses NFS mounted volumes to store the PostgreSQL database file and WAL log files. During a database host reboot, PostgreSQL service might try

to start while the volume is not mounted. This results in database service startup failure. A 10 to 15 seconds timer delay is needed for the PostgreSQL database to start up correctly.

- **RPO/RTO for business continuity.** FSx data replication from primary to standby for DR is based on ASYNC, which means that the RPO depends on the frequency of Snapshot backups and SnapMirror replication. A higher frequency of Snapshot copy and SnapMirror replication reduces the RPO. Therefore, there is a balance between potential data loss in the event of a disaster and incremental storage cost. We have determined that Snapshot copy and SnapMirror replication can be implemented in as low as 5 minute intervals for RPO, and PostgreSQL can generally be recovered at the DR standby site in under a minute for the RTO.
- **Database backup.** After a PostgreSQL database is implemented or migrated into AWS FSx storage from an on-premises data center, the data is auto-sync mirrored in the FSx HA pair for protection. Data is further protected with a replicated standby site in case of a disaster. For longer-term backup retention or data protection, NetApp recommends using the built-in PostgreSQL `pg_basebackup` utility to run a full database backup that can be ported to S3 blob storage.

Solution Deployment

The deployment of this solution can be completed automatically using the NetApp Ansible-based automation toolkit by following the detailed instructions outlined below.

1. Read the instructions in the automation toolkit `README.md` [na_postgresql_aws_deploy_hadr](#).
2. Watch the following video walk through.

Automated PostgreSQL Deployment and Protection

1. Configure the required parameters files (`hosts`, `host_vars/host_name.yml`, `fsx_vars.yml`) by entering user-specific parameters into the template in the relevant sections. Then use the copy button to copy files to the Ansible controller host.

Prerequisites for automated deployment

Deployment requires the following prerequisites.

1. An AWS account has been set up, and the necessary VPC and network segments have been created within your AWS account.
2. From the AWS EC2 console, you must deploy two EC2 Linux instances, one as the primary PostgreSQL DB server at the primary and one at the standby DR site. For compute redundancy at the primary and standby DR sites, deploy two additional EC2 Linux instances as standby PostgreSQL DB servers. See the architecture diagram in the previous section for more details about the environment setup. Also review the [User Guide for Linux instances](#) for more information.
3. From the AWS EC2 console, deploy two FSx ONTAP storage HA clusters to host the PostgreSQL database volumes. If you are not familiar with the deployment of FSx storage, see the documentation [Creating FSx for ONTAP file systems](#) for step-by-step instructions.
4. Build a Centos Linux VM to host the Ansible controller. The Ansible controller can be located either on-premises or in the AWS cloud. If it is located on-premises, you must have SSH connectivity to the VPC, EC2 Linux instances, and FSx storage clusters.
5. Set up the Ansible controller as described in the section "Set up the Ansible Control Node for CLI deployments on RHEL/CentOS" from the resource [Getting Started with NetApp solution automation](#).
6. Clone a copy of the automation toolkit from the public NetApp GitHub site.

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

1. From the toolkit root directory, execute the prerequisite playbooks to install the required collections and libraries for the Ansible controller.

```
ansible-playbook -i hosts requirements.yml
```

```
ansible-galaxy collection install -r collections/requirements.yml --force
--force-with-deps
```

1. Retrieve the required EC2 FSx instance parameters for the DB host variables file `host_vars/*` and the global variables file `fsx_vars.yml` configuration.

Configure the hosts file

Input the primary FSx ONTAP cluster management IP and EC2 instances hosts names into the hosts file.

```
# Primary FSx cluster management IP address
[fsx_ontap]
172.30.15.33
```

```
# Primary PostgreSQL DB server at primary site where database is
initialized at deployment time
[postgresql]
psql_01p ansible_ssh_private_key_file=psql_01p.pem
```

```
# Primary PostgreSQL DB server at standby site where postgresql service is
installed but disabled at deployment
# Standby DB server at primary site, to setup this server comment out
other servers in [dr_postgresql]
# Standby DB server at standby site, to setup this server comment out
other servers in [dr_postgresql]
[dr_postgresql] --
psql_01s ansible_ssh_private_key_file=psql_01s.pem
#psql_01ps ansible_ssh_private_key_file=psql_01ps.pem
#psql_01ss ansible_ssh_private_key_file=psql_01ss.pem
```

Configure the host_name.yml file in the host_vars folder

```
# Add your AWS EC2 instance IP address for the respective PostgreSQL
server host
ansible_host: "10.61.180.15"

# "{{groups.postgresql[0]}}" represents first PostgreSQL DB server as
defined in PostgreSQL hosts group [postgresql]. For concurrent multiple
PostgreSQL DB servers deployment, [0] will be incremented for each
additional DB server. For example, "{{groups.postgresql[1]}}" represents
DB server 2, "{{groups.postgresql[2]}}" represents DB server 3 ... As a
good practice and the default, two volumes are allocated to a PostgreSQL
DB server with corresponding /pgdata, /pglogs mount points, which store
PostgreSQL data, and PostgreSQL log files respectively. The number and
naming of DB volumes allocated to a DB server must match with what is
defined in global fsx_vars.yml file by src_db_vols, src_archive_log_vols
parameters, which dictates how many volumes are to be created for each DB
server. aggr_name is aggr1 by default. Do not change. lif address is the
NFS IP address for the SVM where PostgreSQL server is expected to mount
its database volumes. Primary site servers from primary SVM and standby
servers from standby SVM.
host_datastores_nfs:
  - {vol_name: "{{groups.postgresql[0]}}_pgdata", aggr_name: "aggr1", lif:
"172.21.94.200", size: "100"}
  - {vol_name: "{{groups.postgresql[0]}}_pglogs", aggr_name: "aggr1", lif:
"172.21.94.200", size: "100"}

# Add swap space to EC2 instance, that is equal to size of RAM up to 16G
max. Determine the number of blocks by dividing swap size in MB by 128.
swap_blocks: "128"

# Postgresql user configurable parameters
psql_port: "5432"
buffer_cache: "8192MB"
archive_mode: "on"
max_wal_size: "5GB"
client_address: "172.30.15.0/24"
```

Configure the global fsx_vars.yml file in the vars folder

```
#####
##### PostgreSQL HADR global user configuration variables #####
##### Consolidate all variables from FSx, Linux, and postgresql #####
#####
```

```

#####
### Ontap env specific config variables ###
#####

#####
#####
# Variables for SnapMirror Peering
#####
#####

#Passphrase for cluster peering authentication
passphrase: "xxxxxxx"

#Please enter destination or standby FSx cluster name
dst_cluster_name: "FsxId0cf8e0bccb14805e8"

#Please enter destination or standby FSx cluster management IP
dst_cluster_ip: "172.30.15.90"

#Please enter destination or standby FSx cluster inter-cluster IP
dst_inter_ip: "172.30.15.13"

#Please enter destination or standby SVM name to create mirror
relationship
dst_vserver: "dr"

#Please enter destination or standby SVM management IP
dst_vserver_mgmt_lif: "172.30.15.88"

#Please enter destination or standby SVM NFS lif
dst_nfs_lif: "172.30.15.88"

#Please enter source or primary FSx cluster name
src_cluster_name: "FsxId0cf8e0bccb14805e8"

#Please enter source or primary FSx cluster management IP
src_cluster_ip: "172.30.15.20"

#Please enter source or primary FSx cluster inter-cluster IP
src_inter_ip: "172.30.15.5"

#Please enter source or primary SVM name to create mirror relationship
src_vserver: "prod"

#Please enter source or primary SVM management IP
src_vserver_mgmt_lif: "172.30.15.115"

```

```
#####
#####
# Variable for PostgreSQL Volumes, lif - source or primary FSx NFS lif
address
#####
#####

src_db_vols:
- {vol_name: "{{groups.postgresql[0]}}_pgdata", aggr_name: "aggr1", lif:
"172.21.94.200", size: "100"}

src_archivelog_vols:
- {vol_name: "{{groups.postgresql[0]}}_pglogs", aggr_name: "aggr1", lif:
"172.21.94.200", size: "100"}

#Names of the Nodes in the ONTAP Cluster
nfs_export_policy: "default"

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

#NFS Mount points for PostgreSQL DB volumes
mount_points:
- "/pgdata"
- "/pglogs"

#RedHat subscription username and password
redhat_sub_username: "xxxxx"
redhat_sub_password: "xxxxx"

#####
### DB env specific install and config variables ###
#####
#The latest version of PostgreSQL RPM is pulled/installed and config file
is deployed from a preconfigured template
#Recovery type and point: default as all logs and promote and leave all
PITR parameters blank
```

PostgreSQL deployment and HA/DR setup

The following tasks deploy the PostgreSQL DB server service and initialize the database at the primary site on the primary EC2 DB server host. A standby primary EC2 DB server host is then set up at the standby site. Finally, DB volume replication is set up from the primary-site FSx cluster to the standby-site FSx cluster for disaster recovery.

1. Create DB volumes on the primary FSx cluster, and set up postgresql on the primary EC2 instance host.

```
ansible-playbook -i hosts postgresql_deploy.yml -u ec2-user --private-key psql_01p.pem -e @vars/fsx_vars.yml
```

2. Set up the standby DR EC2 instance host.

```
ansible-playbook -i hosts postgresql_standby_setup.yml -u ec2-user --private-key psql_01s.pem -e @vars/fsx_vars.yml
```

3. Set up FSx ONTAP cluster peering and database volume replication.

```
ansible-playbook -i hosts fsx_replication_setup.yml -e @vars/fsx_vars.yml
```

4. Consolidate the previous steps into a single-step PostgreSQL deployment and HA/DR setup.

```
ansible-playbook -i hosts postgresql_hadr_setup.yml -u ec2-user -e @vars/fsx_vars.yml
```

5. For setting up a standby PostgreSQL DB host at either the primary or standby sites, comment out all other servers in the hosts file [dr_postgresql] section and then execute the postgresql_standby_setup.yml playbook with the respective target host (such as psql_01ps or standby EC2 compute instance at primary site). Make sure that a host parameters file such as psql_01ps.yml is configured under the host_vars directory.

```
[dr_postgresql] --  
#psql_01s ansible_ssh_private_key_file=psql_01s.pem  
psql_01ps ansible_ssh_private_key_file=psql_01ps.pem  
#psql_01ss ansible_ssh_private_key_file=psql_01ss.pem
```

```
ansible-playbook -i hosts postgresql_standby_setup.yml -u ec2-user --private-key psql_01ps.pem -e @vars/fsx_vars.yml
```

PostgreSQL database snapshot backup and replication to standby site

PostgreSQL database snapshot backup and replication to the standby site can be controlled and executed on the Ansible controller with a user-defined interval. We have validated that the interval can be as low as 5 minutes. Therefore, in the case of failure at the primary site, there is 5 minutes of potential data loss if failure occurs right before the next scheduled snapshot backup.


```
*/15 * * * * /home/admin/na_postgresql_aws_deploy_hadr/data_log_snap.sh
```

Failover to Standby Site for DR

For testing the PostgreSQL HA/DR system as a DR exercise, execute failover and PostgreSQL database recovery on the primary standby EC2 DB instance on standby site by executing following playbook. In an actually DR scenario, execute the same for an actually failover to DR site.

```
ansible-playbook -i hosts postgresql_failover.yml -u ec2-user --private-key psql_01s.pem -e @vars/fsx_vars.yml
```

Resync Replicated DB volumes after Failover Test

Run resync after the failover test to reestablish database-volume SnapMirror replication.

```
ansible-playbook -i hosts postgresql_standby_resync.yml -u ec2-user --private-key psql_01s.pem -e @vars/fsx_vars.yml
```

Failover from primary EC2 DB server to standby EC2 DB server due to EC2 compute instance failure

NetApp recommends running manual failover or using well-established OS cluster-ware that might require a license.

Where to find additional information

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

- Amazon FSx for NetApp ONTAP

<https://aws.amazon.com/fsx/netapp-ontap/>

- Amazon EC2

https://aws.amazon.com/pm/ec2/?trk=36c6da98-7b20-48fa-8225-4784bced9843&sc_channel=ps&s_kwid=AL!4422!3!467723097970!e!!g!!aws%20ec2&ef_id=Cj0KCQiA54KfBhCKARIsAJzSrdqwQrghn6l71jiWzSeaT9Uh1-vY-VfhJixF-xnv5rWwn2S7RqZOTQ0aAh7eEALw_wcB:G:s&s_kwid=AL!4422!3!467723097970!e!!g!!aws%20ec2

- NetApp Solution Automation

Introduction

TR-4722: MySQL Database on NetApp ONTAP Best Practices

Anup Bharti, Manohar Kulkarni, Jeffrey Steiner NetApp

MySQL and its variants, including MariaDB and Percona, are widely used for many enterprise applications. These applications range from global social networking sites and

massive e-commerce systems to SMB hosting systems containing thousands of database instances. This document describes the configuration requirements and provides guidance on tuning and storage configuration for deploying MySQL on NetApp® ONTAP® data management software. To determine whether the environment, configurations, and versions specified in this report support your environment, consult the Interoperability Matrix Tool (IMT).

[TR-4722: MySQL Database on NetApp ONTAP Best Practices](#)

SnapCenter for Databases

SnapCenter Oracle Clone Lifecycle Automation

Allen Cao, Niyaz Mohamed, NetApp

This solution provides an Ansible based automation toolkit for configuring Oracle database High Availability and Disaster Recovery (HA/DR) with AWS FSx ONTAP as Oracle database storage and EC2 instances as the compute instances in AWS.

Purpose

Customers love the FlexClone feature of NetApp ONTAP storage for databases with significant storage cost savings. This Ansible based toolkit automates the setup, cloning, and refreshing of cloned Oracle databases on schedule using the NetApp SnapCenter command line utilities for streamlined lifecycle management. The toolkit is applicable to Oracle databases deployed to ONTAP storage either on-premises or public cloud and managed by NetApp SnapCenter UI tool.

This solution addresses the following use cases:

- Setup Oracle database clone-specification configuration file.
- Create and refresh clone Oracle database on user defined schedule.

Audience

This solution is intended for the following people:

- A DBA who manages Oracle databases with SnapCenter.
- A storage administrator who manages ONTAP storage with SnapCenter.
- An application owner who has access to SnapCenter UI.

License

By accessing, downloading, installing or using the content in this GitHub repository, you agree the terms of the License laid out in [License file](#).



There are certain restrictions around producing and/or sharing any derivative works with the content in this GitHub repository. Please make sure you read the terms of the License before using the content. If you do not agree to all of the terms, do not access, download or use the content in this repository.

Solution deployment

Prerequisites for deployment

Deployment requires the following prerequisites.

Ansible controller:

```
Ansible v.2.10 and higher
ONTAP collection 21.19.1
Python 3
Python libraries:
  netapp-lib
  xmltodict
  jmespath
```

SnapCenter server:

```
version 5.0
backup policy configured
Source database protected with a backup policy
```

Oracle servers:

```
Source server managed by SnapCenter
Target server managed by SnapCenter
Target server with identical Oracle software stack as source server
installed and configured
```

Download the toolkit

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

Ansible target hosts file configuration

The toolkit includes a hosts file which define the targets that an Ansible playbook running against. Usually, it is the target Oracle clone hosts. Following is an example file. A host entry includes target host IP address as well as ssh key for an admin user access to the host to execute clone or refresh command.

#Oracle clone hosts

```
[clone_1]
ora_04.cie.netapp.com ansible_host=10.61.180.29
ansible_ssh_private_key_file=ora_04.pem
```

```
[clone_2]
```

```
[clone_3]
```

Global variables configuration

The Ansible playbooks take variable inputs from several variable files. Below is an example global variable file vars.yml.

```
# ONTAP specific config variables
```

```
# SnapCtr specific config variables
```

```
snapctr_usr: xxxxxxxx
snapctr_pwd: 'xxxxxxxx'
```

```
backup_policy: 'Oracle Full offline Backup'
```

```
# Linux specific config variables
```

```
# Oracle specific config variables
```

Host variables configuration

Host variables are defined in `host_vars` directory named as `{{ host_name }}`.yml. Below is an example of target Oracle host variable file `ora_04.cie.netapp.com.yml` that shows typical configuration.

```
# User configurable Oracle clone db host specific parameters
```

```
# Source database to clone from
source_db_sid: NTAP1
source_db_host: ora_03.cie.netapp.com
```

```
# Clone database
clone_db_sid: NTAP1DEV
```

```
snapctr_obj_id: '{{ source_db_host }}\{{ source_db_sid }}'
```

Additional clone target Oracle server configuration

Clone target Oracle server should have the same Oracle software stack as source Oracle server installed and patched. Oracle user `.bash_profile` has `$ORACLE_BASE`, and `$ORACLE_HOME` configured. Also, `$ORACLE_HOME` variable should match with source Oracle server setting. Following is an example.

```
# .bash_profile
```

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

```
# User specific environment and startup programs
export ORACLE_BASE=/u01/app/oracle
export ORACLE_HOME=/u01/app/oracle/product/19.0.0/NTAP1
```

Playbook execution

There are total of three playbooks to execute Oracle database clone lifecycle with SnapCenter CLI utilities.

1. Install Ansible controller prerequisites - one time only.

```
ansible-playbook -i hosts ansible_requirements.yml
```

2. Setup clone specification file - one time only.

```
ansible-playbook -i hosts clone_1_setup.yml -u admin -e  
@vars/vars.yml
```

3. Create and refresh clone database regularly from crontab with a shell script to call a refresh playbook.

```
0 */4 * * * /home/admin/na_oracle_clone_lifecycle/clone_1_refresh.sh
```

For an additional clone database, create a separate clone_n_setup.yml and clone_n_refresh.yml, and clone_n_refresh.sh. Configure the Ansible target hosts and hostname.yml file in host_vars directory accordingly.

Where to find additional information

To learn more about the NetApp solution automation, review the following website [NetApp Solution Automation](#)

TR-4988: Oracle Database Backup, Recovery, and Clone on ANF with SnapCenter

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. Database deployed in Azure is protected using SnapCenter UI tool for simplified database management.

Purpose

NetApp SnapCenter software is an easy-to-use enterprise platform to securely coordinate and manage data protection across applications, databases, and file systems. It simplifies backup, restore, and clone lifecycle management by offloading these tasks to application owners without sacrificing the ability to oversee and regulate activity on the storage systems. By leveraging storage-based data management, it enables increased performance and availability, as well as reduced testing and development times.

In TR-4987, [Simplified, Automated Oracle Deployment on Azure NetApp Files with NFS](#), we demonstrate automated Oracle deployment on Azure NetApp Files (ANF) in Azure cloud. In this documentation, we showcase Oracle database protection and management on ANF in Azure cloud with a very user-friendly SnapCenter UI tool.

This solution addresses the following use cases:

- Backup and recovery of Oracle database deployed on ANF in Azure cloud with SnapCenter.
- Manage database snapshots and clone copies to accelerate application development and improve data lifecycle management.

Audience


This solution is intended for the following people:

- A DBA who would like to deploy Oracle databases 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 Oracle databases 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	A capacity pool with Premium service level

Azure VM for DB server	Standard_B4ms - 4 vCPUs, 16GiB	Two Linux virtual machine instances
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; AE Hotpatch - x64 Gen2	Hosting SnapCenter server
Oracle Database	Version 19.18	Patch p34765931_190000_Linux-x86-64.zip
Oracle OPatch	Version 12.2.0.1.36	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_PDB2,NTAP1_PDB3)	/u01, /u02, /u03 NFS mounts on ANF capacity pool
ora-02	NTAP2(NTAP2_PDB1,NTAP2_PDB2,NTAP2_PDB3)	/u01, /u02, /u03 NFS mounts on ANF capacity pool

Key factors for deployment consideration

- **SnapCenter deployment.** SnapCenter can deploy in a Windows domain or Workgroup environment. For domain-based deployment, the domain user account should be a domain administrator account, or the domain user belongs to the local administrator's group on the SnapCenter hosting server.
- **Name resolution.** SnapCenter server needs to resolve the name to the IP address for each managed target database server host. Each target database server host must resolve the SnapCenter server name to the IP address. If a DNS server is unavailable, add naming to local host files for resolution.
- **Resource group configuration.** Resource group in SnapCenter is a logical grouping of similar resources that can be backed up together. Thus, it simplifies and reduces the number of backup jobs in a large database environment.
- **Separate full database and archive log backup.** Full database backup includes data volumes and log volumes consistent group snapshots. A frequent full database snapshot incurs higher storage consumption but improves RTO. An alternative is less frequent full database snapshots and more frequent archive logs backup, which consumes less storage and improves RPO but may extend RTO. Consider your RTO and RPO objectives when setting up the backup scheme. There is also a limit (1023) of the number of snapshot backups on a volume.

- **Privileges delegation.** Leverage role based access control that is built-in within SnapCenter UI to delegate privileges to application and database teams if desired.

Solution deployment

The following sections provide step-by-step procedures for SnapCenter deployment, configuration, and Oracle database backup, recovery, and clone on Azure NetApp Files in the Azure cloud.

Prerequisites for deployment

Deployment requires existing Oracle databases running on ANF in Azure. If not, follow the steps below to create two Oracle databases for solution validation. For details of Oracle database deployment on ANF in Azure cloud with automation, referred to TR-4987: [Simplified, Automated Oracle Deployment on Azure NetApp Files with NFS](#)

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 /tmp and /root 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. Follow instructions in [TR-4887](#) to execute the playbooks.

```
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:


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

8. Review the `Get Started` online menu.

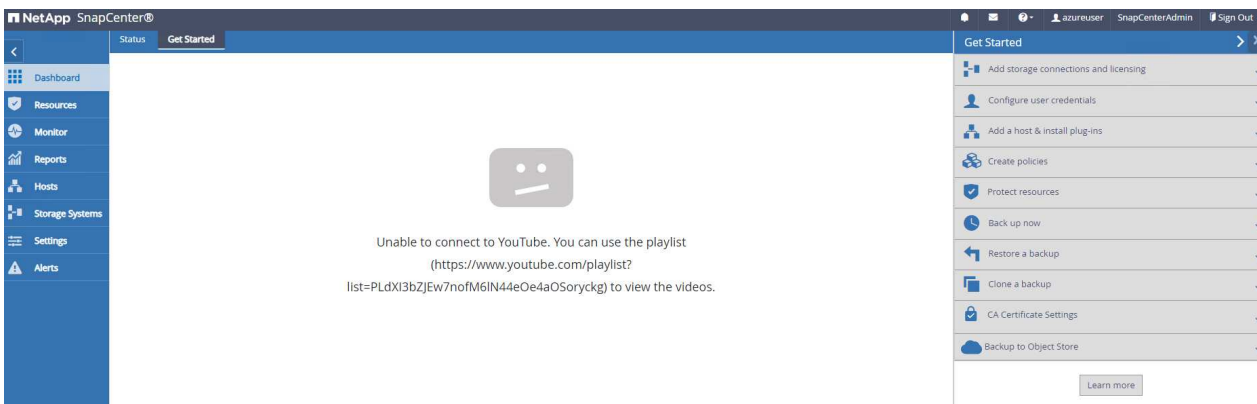
SnapCenter installation and setup

We recommend to go through online [SnapCenter Software documentation](#) before proceeding to SnapCenter installation and configuration: . Following provides a high level summary of steps for installation and setup of SnapCenter software for Oracle on Azure ANF.


1. From SnapCenter Windows server, download and install latest java JDK from [Get Java for desktop applications](#).
2. From SnapCenter Windows server, download and install latest version (currently 5.0) of SnapCenter installation executable from NetApp support site: [NetApp | Support](#).
3. After SnapCenter server installation, launch browser to login to SnapCenter with Windows local admin user or domain user credential via port 8146.



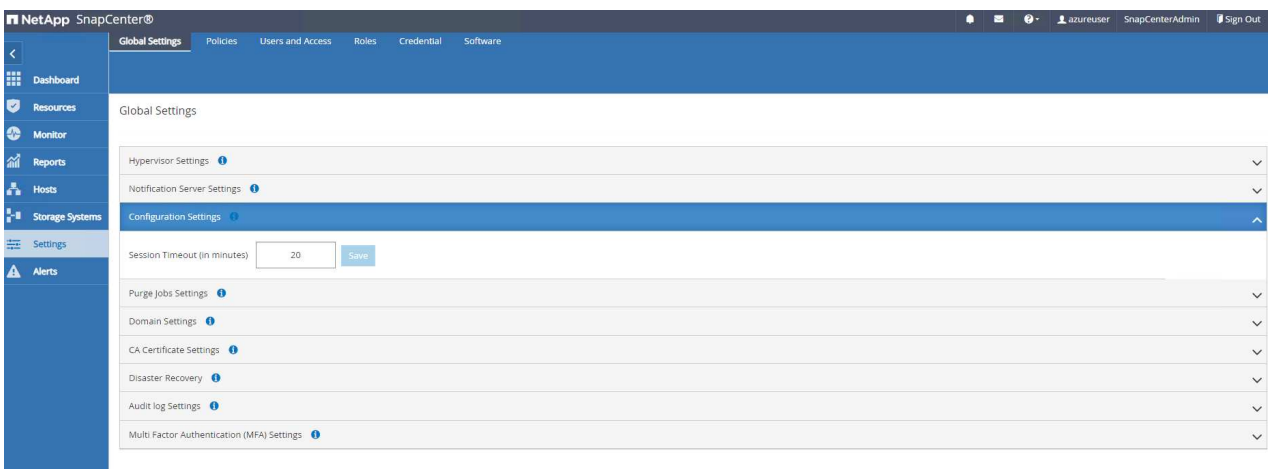
4. Review Get Started online menu.



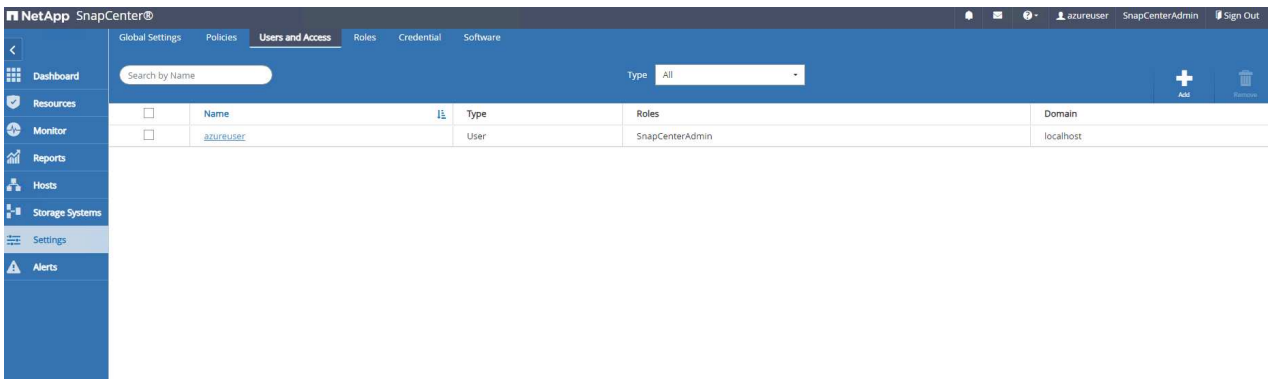
5. In Settings-Global Settings, check Hypervisor Settings and click on Update.



6. If needed, adjust `Session Timeout` for SnapCenter UI to the desired interval.



7. Add additional users to SnapCenter if needed.



8. The `Roles` tab list the built-in roles that can be assigned to different SnapCenter users. Custom roles also can be created by admin user with desired privileges.

Name	Details	Members
<input type="checkbox"/> SnapCenterAdmin	Overall administrator of SnapCenter system	1 User, No Groups
<input type="checkbox"/> App Backup and Clone Admin	App Backup and Clone Admin	No Members
<input type="checkbox"/> Backup and Clone Viewer	Backup and Clone Viewer	No Members
<input type="checkbox"/> Infrastructure Admin	Infrastructure Admin	No Members

9. From Settings-Credential, create credentials for SnapCenter management targets. In this demo use case, they are linux user for login to Azure VM and ANF credential for capacity pool access.

Credential Name	Authentication Mode	Details
azure_anf	AzureCredential	
azureuser	Linux	userid:azureuser

Credential ✕

Credential Name

Authentication Mode

Authentication Type Password Based SSH Key Based i

Username i

SSH Private Key

XRlrK1QCae0Hg==
-----END RSA PRIVATE KEY-----

i

Use sudo privileges i

Credential ✕

Credential Name

Authentication Mode

Azure Details ⓘ

Tenant ID

Client ID

Client Secret Key

10. From Storage Systems tab, add Azure NetApp Files with credential created above.

NetApp SnapCenter | ONTAP Storage | Azure NetApp Files

Search by NetApp Account

NetApp Account	Resource Group	Credential
<input type="checkbox"/> ANFAVSAcct	ANFAVSRG	azure_anf

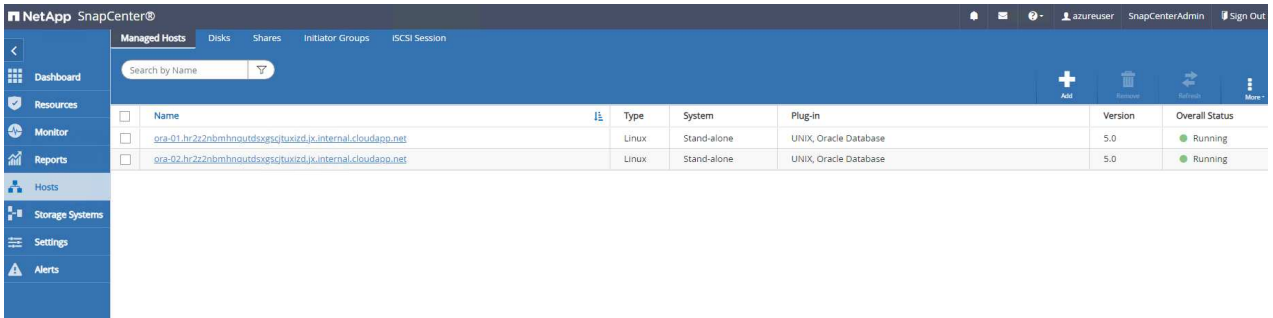
Add Azure NetApp Account ✕

Credential ⓘ

Subscription ⓘ

NetApp Account ⓘ

11. From Hosts tab, add Azure DB VMs, which installs SnapCenter plugin for Oracle on Linux.



Name	Type	System	Plug-in	Version	Overall Status
ora-01.hr2z2nbmhnouds9sqtuozd.jx.internal.cloudapp.net	Linux	Stand-alone	UNIX, Oracle Database	5.0	Running
ora-02.hr2z2nbmhnouds9sqtuozd.jx.internal.cloudapp.net	Linux	Stand-alone	UNIX, Oracle Database	5.0	Running

Add Host

Host Type: Linux

Host Name: ora-01

Credentials: azureuser

Select Plug-ins to Install SnapCenter Plug-ins Package 5.0 for Linux

- Oracle Database
- SAP HANA
- Unix File Systems

 [More Options](#): Port, Install Path, Custom Plug-Ins...

More Options
✕

Port

Installation Path

Skip optional preinstall checks

Add all hosts in the oracle RAC

Custom Plug-ins

Choose a File

No plug-ins found.

12. Once host plugin is installed on DB server VM, databases on the host are auto discovered and visible in Resources tab. Back to Settings-Policies, create backup policies for full Oracle database online backup and archive logs only backup. Refer to this document [Create backup policies for Oracle databases](#) for detailed step by step procedures.


The screenshot shows the NetApp SnapCenter interface. The 'Policies' tab is selected under 'Oracle Database'. A table lists two backup policies:

Name	Backup Type	Schedule Type	Replication	Verification
Oracle archive/logs backup	LOG, ONLINE	Hourly		
Oracle full online backup	FULL, ONLINE	Hourly		

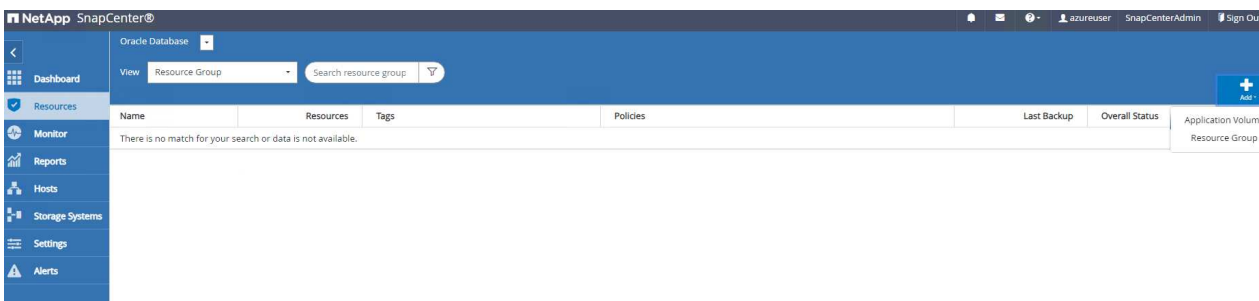
Database backup

A NetApp snapshot backup creates a point-in-time image of the database volumes that you can use to restore in case of a system failure or data loss. Snapshot backups take very little time, usually less than a minute. The backup image consumes minimal storage space and incurs negligible performance overhead because it records only changes to files since the last snapshot copy was made. Following section demonstrates the implementation of snapshots for Oracle database backup in SnapCenter.

1. Navigating to Resources tab, which lists the databases discovered once SnapCenter plugin installed on database VM. Initially, the Overall Status of database shows as Not protected.



2. Click on View drop-down to change to Resource Group. Click on Add sign on the right to add a Resource Group.



3. Name your resource group, tags, and any custom naming.

New Resource Group

1 Name 2 Resources 3 Policies 4 Verification 5 Notification 6 Summary

Provide a name and tags for the resource group

Name

Tags

Use custom name format for Snapshot copy

Backup settings

Exclude archive log destinations from backup

Previous Next

4. Add resources to your Resource Group. Grouping of similar resources can simplify database management in a large environment.

New Resource Group

1 Name 2 Resources 3 Policies 4 Verification 5 Notification 6 Summary

Add resources to Resource Group

Host

Available Resources

Selected Resources

NTAP1 (ora-01.hr22nbnmhnqtdsxsgjtuxizd.jk.internal.cloudapp.i
 NTAP2 (ora-02.hr22nbnmhnqtdsxsgjtuxizd.jk.internal.cloudapp.i

><

Previous Next

5. Select the backup policy and set a schedule by click on '+' sign under Configure Schedules.



Select one or more policies and configure schedules

Oracle full online backup + ⓘ

Configure schedules for selected policies

Policy	Applied Schedules	Configure Schedules
Oracle full online backup	None	+

Total 1

Previous Next

Add schedules for policy Oracle full online backup

Hourly

Start date 02/06/2024 05:55 pm

Expires on 03/06/2024 05:51 pm

Repeat every 2 hours 0 mins

i The schedules are triggered in the SnapCenter Server time zone.

Cancel

OK

6. If backup verification is not configured in policy, leave verification page as is.

New Resource Group

1 Name 2 Resources 3 Policies 4 Verification 5 Notification 6 Summary

Configure verification schedules

Policy Schedule Type Applied Schedules Configure Schedules

There is no match for your search or data is not available.

Total 0

Previous Next

7. In order to email a backup report and notification, a SMTP mail server is needed in the environment. Or leave it blank if a mail server is not setup.

New Resource Group

1 Name 2 Resources 3 Policies 4 Verification 5 Notification 6 Summary

Provide email settings ⓘ

Select the service accounts or people to notify regarding protection issues.

Email preference: Never

From: From email

To: Email to

Subject: Notification

Attach job report

Previous Next

8. Summary of new resource group.

New Resource Group

1 Name 2 Resources 3 Policies 4 Verification 5 Notification 6 Summary

Resource group name: full_online_bkup

Tags: oradata

Policy: Oracle full online backup: Hourly

Plug-in: SnapCenter Plug-in for Oracle Database

Verification enabled for policy: None

Send email: No

Previous Finish

9. Repeat the above procedures to create a database archive log only backup with corresponding backup policy.

NetApp SnapCenter

Oracle Database

View: Resource Group Search resource group

Name	Resources	Tags	Policies	Last Backup	Overall Status
full_online_bkup	2	oradata	Oracle full online backup	02/06/2024 6:00:44 PM	Completed
archivelog_bkup	2	oralog	Oracle archivelogs backup	02/06/2024 5:59:25 PM	Completed

10. Click on a resource group to reveal the resources it includes. Besides the scheduled backup job, an one-off backup can be triggered by clicking on Backup Now.

NetApp SnapCenter

Oracle Database

full_online_bkup Details

Search resource groups search

Name	Resource Name	Type	Host
full_online_bkup	NTAP1	Oracle Database	ora-01.hr222nbmhnqustdxgscjtuxidz.jx.internal.cloudapp.net
archivelog_bkup	NTAP2	Oracle Database	ora-02.hr222nbmhnqustdxgscjtuxidz.jx.internal.cloudapp.net

Modify Resource Group Backup Now Maintenance Delete

Backup



Create a backup for the selected resource group

Resource Group

full_online_bkup

Policy

Oracle full online backup



Verify after backup

Cancel

Backup

11. Click on the running job to open a monitoring window, which allows the operator to track the job progress in real-time.

Job Details



Backup of Resource Group 'full_online_bkup' with policy 'Oracle full online backup'

- ✓ Backup of Resource Group 'full_online_bkup' with policy 'Oracle full online backup'
- ✓ ▶ ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net
- ✓ ▶ ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net

i Task Name: Backup of Resource Group 'full_online_bkup' with policy 'Oracle full online backup' Start Time: 02/06/2024 6:00:05 PM End Time: 02/06/2024 6:00:44 PM

View Logs

Cancel Job

Close

12. A snapshot backup set appears under database topology once a successful backup job finishes. A full database backup set includes a snapshot of the database data volumes and a snapshot of the database log volumes. A log-only backup contains only a snapshot of the database log volumes.

The screenshot shows the NetApp SnapCenter interface for an Oracle Database. The main content area displays backup details for 'full_online_bkup' and 'archivelog_bkup'. A table lists the primary backups, and a summary card provides a high-level overview of the backup status.

Backup Name	Snapshot Lock Expiration	Count	Type	End Date	Verified	Mounted	RMAN Cataloged	SCN
ora-01_02-06-2024_18_00_06_0582_1		1	Log	02/06/2024 6:00:41 PM	Not Applicable	False	Not Cataloged	3374950
ora-01_02-06-2024_18_00_06_0582_0		1	Data	02/06/2024 6:00:26 PM	Unverified	False	Not Cataloged	3374903
ora-01_02-06-2024_17_59_01_1158_1		1	Log	02/06/2024 5:59:18 PM	Not Applicable	False	Not Cataloged	3374762

Summary Card:

- 3 Backups
- 1 Data Backup
- 2 Log Backups
- 0 Clones
- 0 Snapshots Locked

Database recovery

Database recovery via SnapCenter restores a snapshot copy of the database volume image point-in-time. The database is then rolled forward to a desired point by SCN/timestamp or a point as allowed by available archive logs in the backup set. The following section demonstrates the workflow of database recovery with SnapCenter UI.

1. From **Resources** tab, open the database **Primary Backup(s)** page. Choose the snapshot of database data volume, then click on **Restore** button to launch database recovery workflow. Note the SCN number or timestamp in the backup sets if you like to run the recovery by Oracle SCN or timestamp.

NTAP1 Topology

Backup to Object Store Protect Refresh

Manage Copies

3 Backups
0 Clones
Local copies

Summary Card

3 Backups
1 Data Backup
2 Log Backups
0 Clones
0 Snapshots Locked

Primary Backup(s)

search

Catalog Rename Clone **Restore** Mount Unmount Delete

Backup Name	Snapshot Lock Expiration	Count	Type	End Date	Verified	Mounted	RMAN Cataloged	SCN
ora-01_02-06-2024_18_00_06_0582_1		1	Log	02/06/2024 6:00:41 PM	Not Applicable	False	Not Cataloged	3374950
ora-01_02-06-2024_18_00_06_0582_0		1	Data	02/06/2024 6:00:26 PM	Unverified	False	Not Cataloged	3374903
ora-01_02-06-2024_17_59_01_1158_1		1	Log	02/06/2024 5:59:18 PM	Not Applicable	False	Not Cataloged	3374762

2. Select **Restore Scope**. For a container database, SnapCenter is flexible to perform a full container database (All Datafiles), pluggable databases, or tablespaces level restore.

Restore NTAP1 ×

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Restore Scope ⓘ

All Datafiles

Pluggable databases (PDBs)

Pluggable database (PDB) tablespaces

Control files

Database State

Change database state if needed for restore and recovery


Restore Mode ⓘ

Force in place restore

If this check box is not selected and if any of the in place restore criteria is not met, restore will be performed using the connect and copy method. The connect and copy restore method might take time based on the files being restored.

Previous Next

3. Select **Recovery Scope**. All logs means to apply all available archive logs in the backup set. Point-in-time recovery by SCN or timestamp are also available.



4. The `PreOps` allows execution of scripts against database before restore/recovery operation.

Restore NTAP1



1 Restore Scope

Specify optional scripts to run before performing a restore job ⓘ

2 Recovery Scope

Prescript full path Enter Prescript path

3 PreOps

Arguments

4 PostOps

Script timeout secs

5 Notification

6 Summary

Previous

Next

5. The `PostOps` allows execution of scripts against database after restore/recovery operation.

Restore NTAP1



1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Specify optional scripts to run after performing a restore job

Postscript full path

Arguments

Open the database or container database in READ-WRITE mode after recovery

Previous

Next

6. Notification via email if desired.

Restore NTAP1



1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Provide email settings


Email preference:

From:

To:

Subject:

Attach job report

 If you want to send notifications for Restore jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server.

Previous

Next

7. Restore job summary

Restore NTAP1 ✕

- 1 Restore Scope
- 2 Recovery Scope
- 3 PreOps
- 4 PostOps
- 5 Notification
- 6 Summary**

Summary

Backup name	ora-01_02-06-2024_18_00_06_0582_0
Backup date	02/06/2024 6:00:26 PM
Restore scope	All DataFiles
Recovery scope	All Logs
Options	Change database state if necessary , Open the database or container database in READ-WRITE mode after recovery
Prescript full path	None
Prescript arguments	
Postscript full path	None
Postscript arguments	
Send email	No

8. Click on running job to open Job Details window. The job status can also be opened and viewed from the Monitor tab.

Job Details



Restore 'ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net\NTAP1'

✓ ▾ Restore 'ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net\NTAP1'

✓ ▾ ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net

- ✓ ▶ Prescripts
- ✓ ▶ Mount log backups
- ✓ ▶ Pre Restore
- ✓ ▶ Restore
- ✓ ▶ Post Restore
- ✓ ▶ Unmount log backups
- ✓ ▶ Postscripts
- ✓ ▶ Post Restore Cleanup
- ✓ ▶ Data Collection

i Task Name: ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net Start Time: 02/06/2024 4:04:55 PM End Time: 02/06/2024 4:08:42 PM

View Logs

Cancel Job

Close

Database clone

Database clone via SnapCenter is accomplished by creating a new volume from a snapshot of a volume. The system uses the snapshot information to clone a new volume using the data on the volume when the snapshot was taken. More importantly, it is quick (a few minutes) and efficient compared with other methods to make a cloned copy of the production database to support development or testing. Thus, dramatically improve your database application lifecycle management. The following section demonstrates the workflow of database clone with SnapCenter UI.

1. From Resources tab, open the database Primary Backup(s) page. Choose the snapshot of database data volume, then click on clone button to launch database clone workflow.

The screenshot shows the SnapCenter interface for 'NTAP1 Topology'. It includes a 'Manage Copies' section with '3 Backups' and '0 Clones'. A 'Summary Card' on the right shows '3 Backups', '1 Data Backup', '2 Log Backups', '0 Clones', and '0 Snapshots Locked'. The 'Primary Backup(s)' section features a search bar and a toolbar with buttons for Catalog, Rename, Clone (highlighted), Restore, Mount, Unmount, and Delete. Below the toolbar is a table of backup details.

Backup Name	Snapshot Lock Expiration	Count	Type	End Date	Verified	Mounted	RMAN Cataloged	SCN
ora-01_02-06-2024_18_00_06_0582_1		1	Log	02/06/2024 6:00:41 PM	Not Applicable	False	Not Cataloged	3374950
ora-01_02-06-2024_18_00_06_0582_0		1	Data	02/06/2024 6:00:26 PM	Unverified	False	Not Cataloged	3374903
ora-01_02-06-2024_17_59_01_1158_1		1	Log	02/06/2024 5:59:18 PM	Not Applicable	False	Not Cataloged	3374762

2. Name the clone database SID. Optionally, for a container database, clone can be done at PDB level as well.

Clone from NTAP1



- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps
- 6 Notification
- 7 Summary

Capacity Pool Max. Throughput (MiB/s)



Complete Database Clone

Clone SID

Exclude PDBs:

PDB Clone

Previous

Next

3. Select the DB server where you want to place your cloned database copy. Keep the default file locations unless you want to name them differently.

✕
Clone from NTAP1

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Select the host to create a clone

Clone host:

Datafile locations ?

Reset

Control files ?

✕

✕ Reset

Redo logs ?

Group	Size	Unit	Number of files
▶ RedoGroup 1	✕ 200	MB	1
▶ RedoGroup 2	✕ 200	MB	1
▶ RedoGroup 3	✕ 200	MB	1

Previous
Next

4. Identical Oracle software stack as in source database should have been installed and configured on clone DB host. Keep the default credential but change Oracle Home Settings to match with settings on clone DB host.

Clone from NTAP1



1 Name

Database Credentials for the clone

2 Locations

Credential name for sys user

None



3 Credentials

Database port

1521

4 PreOps

Oracle Home Settings

5 PostOps

Oracle Home

/u01/app/oracle/product/19.0.0/NTAP2

6 Notification

Oracle OS User

oracle

7 Summary

Oracle OS Group

oinstall

Previous

Next

5. The `PreOps` allows execution of scripts before clone operation. Database parameters can be adjusted to meet a clone DB needs as versus a production database, such as reduced SGA target.

Clone from NTAP1



- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps**
- 5 PostOps
- 6 Notification
- 7 Summary

Specify scripts to run before clone operation ?

Prescript full path

Arguments

Script timeout

Database Parameter settings

processes	320	<input type="text" value="X"/>	<input type="text" value="▲"/>
remote_login_passwordfile	EXCLUSIVE	<input type="text" value="X"/>	<input type="text" value="+"/>
sga_target	<input type="text" value="3G"/>	<input type="text" value="X"/>	<input type="text" value="Reset"/>
undo_tablespace	UNDOTBS1	<input type="text" value="X"/>	<input type="text" value="▼"/>

Previous

Next

- The `PostOps` allows execution of scripts against database after clone operation. Clone database recovery can be SCN, timestamp based, or Until cancel (rolling forward database to last archived log in the backup set).

Clone from NTAP1



- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps**
- 6 Notification
- 7 Summary

Recover Database

Until Cancel ?

Date and Time



Date-time format: MM/DD/YYYY hh:mm:ss

Until SCN (System Change Number)



Specify external archive log locations ?

Create new DBID ?

Create tempfile for temporary tablespace ?

Enter SQL queries to apply when clone is created

Enter scripts to run after clone operation ?

Previous

Next

7. Notification via email if desired.

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps
- 6 Notification**
- 7 Summary

Provide email settings ⓘ

Email preference:

From:

To:

Subject:

Attach job report

⚠ If you want to send notifications for Clone jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server.

8. Clone job summary.

Clone from NTAP1



1 Name	Summary
2 Locations	Clone from backup ora-01_02-06-2024_18_00_06_0582_0
3 Credentials	Clone SID ntap1 dev
4 PreOps	Capacity Pool Max. Throughput (MiB/s) none
5 PostOps	Clone server ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net
6 Notification	Exclude PDBs none
7 Summary	Oracle home /u01/app/oracle/product/19.0.0/NTAP2
	Oracle OS user oracle
	Oracle OS group oinstall
	Datafile mountpaths /u02_ntap1 dev
	Control files /u02_ntap1 dev/ntap1 dev/control/control01.ctl /u02_ntap1 dev/ntap1 dev/control/control02.ctl
	Redo groups RedoGroup =1 TotalSize =200 Path =/u02_ntap1 dev/ntap1 dev/redolog/redo01_01.log RedoGroup =2 TotalSize =200 Path =/u02_ntap1 dev/ntap1 dev/redolog/redo02_01.log RedoGroup =3 TotalSize =200 Path =/u02_ntap1 dev/ntap1 dev/redolog/redo03_01.log
	Recovery scope Until Cancel
	Prescript full path none
	Prescript arguments
	Postscript full path none
	Postscript arguments
	Send email No

Previous Finish

- Click on running job to open Job Details window. The job status can also be opened and viewed from the Monitor tab.

Job Details

Clone from backup 'ora-01_02-06-2024_18_00_06_0582_0'

- ✓ ▾ Clone from backup 'ora-01_02-06-2024_18_00_06_0582_0'
 - ✓ ▾ ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net
 - ✓ ▶ Prescripts
 - ✓ ▶ Query Host Information
 - ✓ ▶ Prepare for Cloning
 - ✓ ▶ Cloning Resources
 - ✓ ▶ FileSystem Clone
 - ✓ ▶ Application Clone
 - ✓ ▶ Postscripts
 - ✓ ▶ Register Clone
 - ✓ ▶ Unmount Clone
 - ✓ ▶ Data Collection

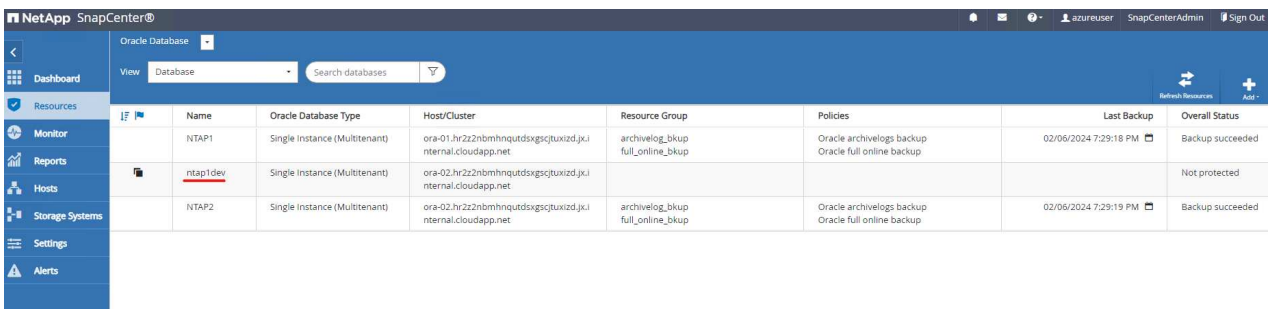
Task Name: ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net Start Time: 02/06/2024 6:21:59 PM End Time: 02/06/2024 6:28:10 PM

View Logs

Cancel Job

Close

10. Cloned database registers with SnapCenter immediately.



Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Overall Status
NTAP1	Single Instance (Multitenant)	ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net	archivelog_bkup full_online_bkup	Oracle archivelogs backup Oracle full online backup	02/06/2024 7:29:18 PM	Backup succeeded
ntap1dev	Single Instance (Multitenant)	ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net				Not protected
NTAP2	Single Instance (Multitenant)	ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net	archivelog_bkup full_online_bkup	Oracle archivelogs backup Oracle full online backup	02/06/2024 7:29:19 PM	Backup succeeded

11. Validate clone database on DB server host. For a cloned development database, database archive mode should be turned off.

```

[azureuser@ora-02 ~]$ sudo su
[root@ora-02 azureuser]# su - oracle
Last login: Tue Feb  6 16:26:28 UTC 2024 on pts/0

[oracle@ora-02 ~]$ uname -a
Linux ora-02 4.18.0-372.9.1.el8.x86_64 #1 SMP Fri Apr 15 22:12:19
EDT 2022 x86_64 x86_64 x86_64 GNU/Linux
[oracle@ora-02 ~]$ 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	49M	7.7G
1% /run			
tmpfs	7.8G	0	7.8G
0% /sys/fs/cgroup			
/dev/mapper/rootvg-rootlv	22G	17G	5.6G
75% /			
/dev/mapper/rootvg-usrlv	10G	2.0G	8.1G
20% /usr			
/dev/mapper/rootvg-homelv	1014M	40M	975M
4% /home			
/dev/sda1	496M	106M	390M
22% /boot			
/dev/mapper/rootvg-varlv	8.0G	958M	7.1G
12% /var			
/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-02-u03	250G	2.1G	248G
1% /u03			
172.30.136.68:/ora-02-u01	100G	10G	91G
10% /u01			
172.30.136.68:/ora-02-u02	250G	7.5G	243G
3% /u02			
tmpfs	1.6G	0	1.6G
0% /run/user/1000			
tmpfs	1.6G	0	1.6G
0% /run/user/0			
172.30.136.68:/ora-01-u02-Clone-020624161543077	250G	8.2G	242G

```
4% /u02_ntapldev
```

```
[oracle@ora-02 ~]$ 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.
```

```
#
```

```
#
```

```
NTAP2:/u01/app/oracle/product/19.0.0/NTAP2:Y
```

```
# SnapCenter Plug-in for Oracle Database generated entry (DO NOT  
REMOVE THIS LINE)
```

```
ntapldev:/u01/app/oracle/product/19.0.0/NTAP2:N
```

```
[oracle@ora-02 ~]$ export ORACLE_SID=ntapldev
```

```
[oracle@ora-02 ~]$ sqlplus / as sysdba
```

```
SQL*Plus: Release 19.0.0.0.0 - Production on Tue Feb 6 16:29:02 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
NTAP1DEV	READ WRITE	ARCHIVELOG

```
SQL> shutdown immediate;
```

```
Database closed.
```

```
Database dismounted.
```

```
ORACLE instance shut down.
```

```
SQL> startup mount;
```

```
ORACLE instance started.
```

```
Total System Global Area 3221223168 bytes
```

```
Fixed Size 9168640 bytes
```

```
Variable Size 654311424 bytes
```

```
Database Buffers 2550136832 bytes
```

```
Redo Buffers 7606272 bytes
```

```
Database mounted.
```

```
SQL> alter database noarchivelog;
```

```
Database altered.
```

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

```
Database altered.
```

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

NAME	OPEN_MODE	LOG_MODE
NTAP1DEV	READ WRITE	NOARCHIVELOG

```
SQL> show pdba
```

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

```
SQL> alter pluggable database all open;
```

Where to find additional information

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

- Azure NetApp Files

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

- SnapCenter Software documentation

<https://docs.netapp.com/us-en/snapcenter/index.html>

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

[Deployment Procedure](#)

TR-4977: Oracle Database backup, restore and clone with SnapCenter Services - Azure

Allen Cao, Niyaz Mohamed, NetApp

This solution provides overview and details for Oracle database backup, restore, clone using NetApp SnapCenter SaaS using BlueXP console.

Purpose

SnapCenter Services is the SaaS version of the classic SnapCenter database management UI tool that is available through the NetApp BlueXP cloud management console. It is an integral part of the NetApp cloud-backup, data-protection offering for databases such as Oracle and HANA running on Azure NetApp Files. This SaaS-based service simplifies traditional SnapCenter standalone server deployment that generally requires a Windows server operating in a Windows domain environment.

In this documentation, we demonstrate how you can set up SnapCenter Services to backup, restore, and clone Oracle databases deployed on Azure NetApp Files volumes and Azure compute instances. It is very easy to setup data protection for Oracle database deployed on Azure NetApp Files with web based BlueXP user interface.

This solution addresses the following use cases:

- Database backup with snapshots for Oracle databases hosted in Azure NetApp Files and Azure VMs
- Oracle database recovery in the case of a failure
- Fast cloning of primary databases for dev, test environments or other use cases

Audience

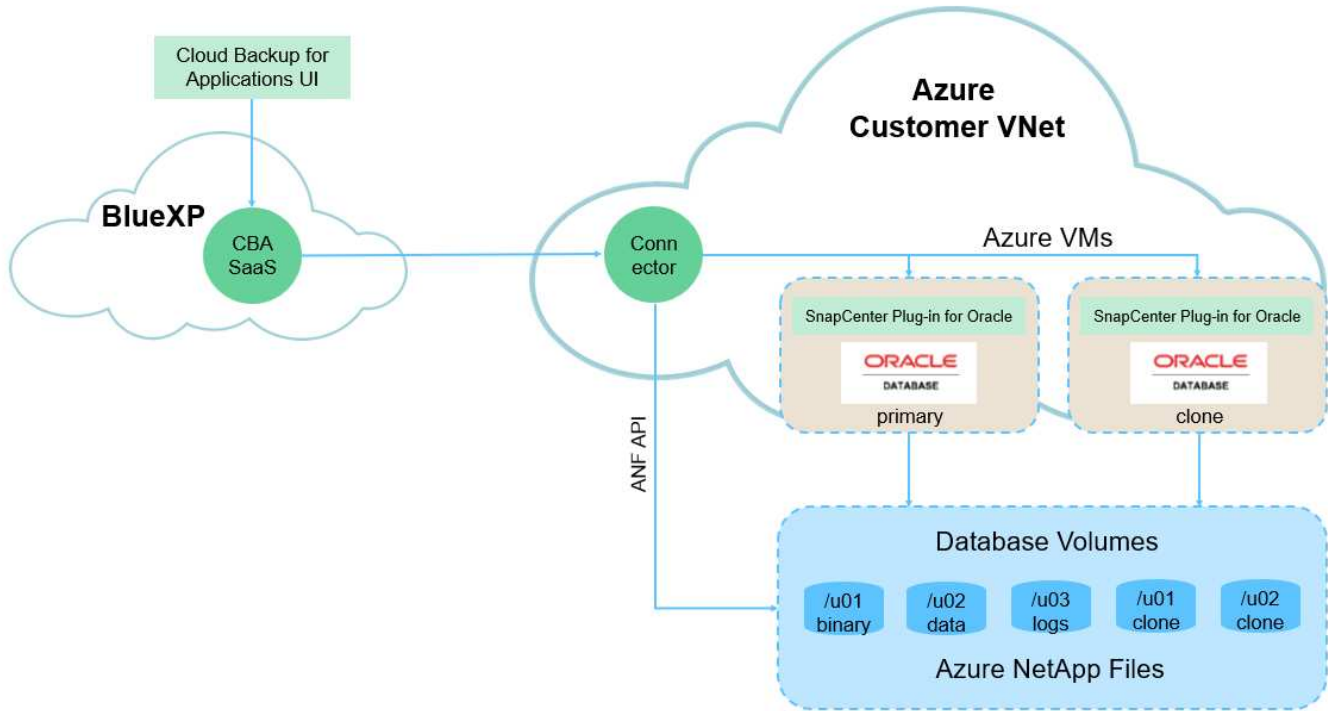
This solution is intended for the following audiences:

- The DBA who manages Oracle databases running on Azure NetApp Files storage
- The solution architect who is interested in testing Oracle database backup, restore, and clone in Azure
- The storage administrator who supports and manages the Azure NetApp Files storage
- The application owner who owns applications that are deployed to Azure NetApp Files storage and Azure

Solution test and validation environment

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

Architecture



This image provides a detailed picture of BlueXP backup and recovery for applications within the BlueXP console, including the UI, the connector, and the resources it manages.

Hardware and software components

Hardware

Azure NetApp Files storage	Premium Service level	Auto QoS type, and 4TB in storage capacity in testing
Azure instance for compute	Standard B4ms (4 vcpus, 16 GiB memory)	Two instances deployed, one as primary DB server and the other as clone DB server

Software

RedHat Linux	Red Hat Enterprise Linux 8.7 (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
SnapCenter Service	Version v2.5.0-2822	Agent Version v2.5.0-2822

Key factors for deployment consideration

- **Connector to be deployed in the same virtual network / subnet as databases and Azure NetApp Files.** When possible, the connector should be deployed in the same Azure virtual networks and resource groups, which enables connectivity to the Azure NetApp Files storage and the Azure compute instances.
- **An Azure user account or Active Directory service principle created at Azure portal for SnapCenter connector.** Deploying a BlueXP Connector requires specific permissions to create and configure a virtual machine and other compute resources, to configure networking, and to get access to the Azure subscription. It also requires permissions to later create roles and permissions for the Connector to operate. Create a custom role in Azure with permissions and assign to the user account or service principle. Review the following link for details: [Set up Azure permissions](#).
- **A ssh key pair created in the Azure resource group.** The ssh key pair is assigned to the Azure VM user for logging into the connector host and also the database VM host for deploying and executing a plug-in. BlueXP console UI uses the ssh key to deploy SnapCenter service plugin to database host for one-step plugin installation and application host database discovery.
- **A credential added to the BlueXP console setting.** To add Azure NetApp Files storage to the BlueXP working environment, a credential that grants permissions to access Azure NetApp Files from the BlueXP console needs to be set up in the BlueXP console setting.
- **java-11-openjdk installed on the Azure VM database instance host.** SnapCenter service installation requires java version 11. It needs to be installed on application host before plugin deployment attempt.

Solution deployment

There is extensive NetApp documentation with a broader scope to help you protect your cloud-native application data. The goal of this documentation is to provide step-by-step procedures that cover SnapCenter Service deployment with the BlueXP console to protect your Oracle database deployed on an Azure NetApp Files storage and an Azure compute instance.

To get started, complete the following steps:

- Read the general instructions [Protect your cloud native applications data](#) and the sections related to Oracle and Azure NetApp Files.
- Watch the following video walkthrough

[Video of deployment of Oracle and ANF](#)

Prerequisites for SnapCenter service deployment

Deployment requires the following prerequisites.

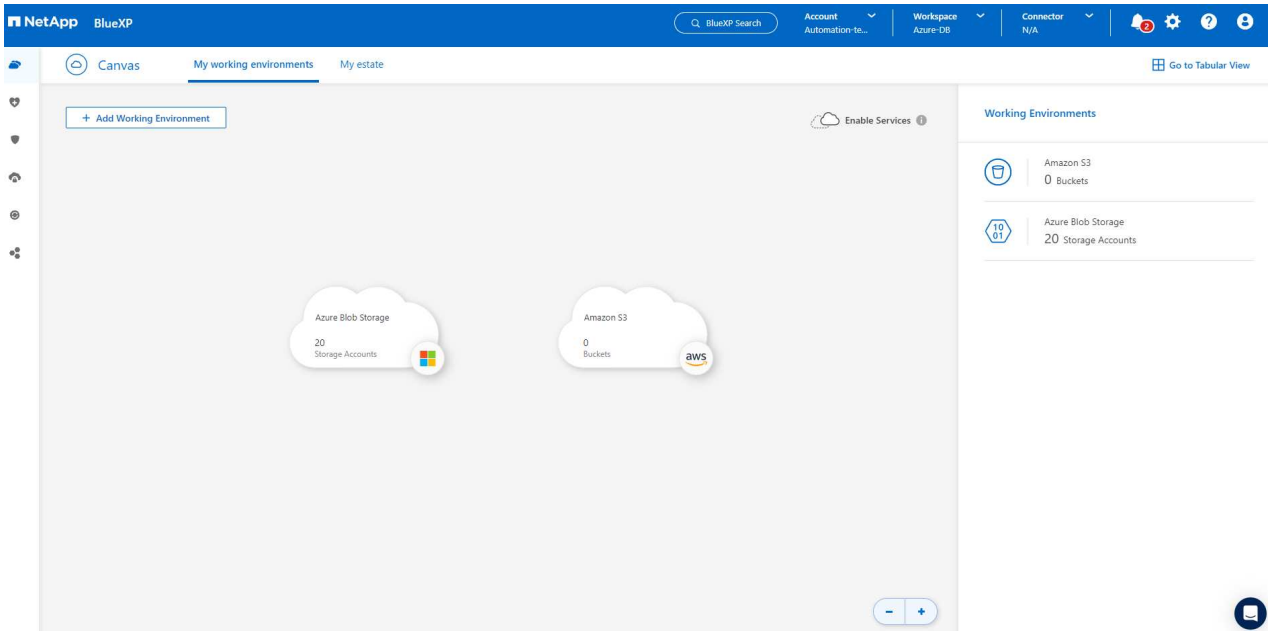
1. A primary Oracle database server on an Azure VM instance with an Oracle database fully deployed and running.
2. An Azure NetApp Files storage service capacity pool deployed in Azure that has capacity to meet the database storage needs listed in hardware component section.
3. A secondary database server on an Azure VM instance that can be used for testing the cloning of an Oracle database to an alternate host for the purpose of supporting a dev/test workload or any use cases that requires a full data set of production Oracle database.
4. For additional information for Oracle database deployment on Azure NetApp Files and Azure compute instance, see [Oracle Database Deployment and Protection on Azure NetApp Files](#).

Onboarding to BlueXP preparation


1. Use the link [NetApp BlueXP](#) to sign up for BlueXP console access.
2. Create an Azure user account or an Active Directory service principle and grant permissions with role in Azure portal for Azure connector deployment.
3. To set up BlueXP to manage Azure resources, add a BlueXP credential with details of an Active Directory service principal that BlueXP can use to authenticate with Azure Active Directory (App client ID), a client secret for the service principal application (Client Secret), and the Active Directory ID for your organization (Tenant ID).
4. You also need the Azure virtual network, resources group, security group, an SSH key for VM access, etc. ready for connector provisioning and database plugin installation.

Deploy a connector for SnapCenter services

1. Login to the BlueXP console.



2. Click on **Connector** drop down arrow and **Add Connector** to launch the connector provisioning workflow.



3. Choose your cloud provider (in this case, **Microsoft Azure**).

Provider

Choose the cloud provider where you want to run the BlueXP Connector:



[Deploy the Connector on your premises](#)

Continue



4. Skip the **Permission**, **Authentication**, and **Networking** steps if you already have them set up in your Azure account. If not, you must configure these before proceeding. From here, you could also retrieve the permissions for the Azure policy that is referenced in the previous section "[Onboarding to BlueXP preparation](#)."

Deploying a BlueXP Connector

The BlueXP Connector is a crucial component for the day-to-day use of BlueXP.

It's used to connect BlueXP's services to your hybrid-cloud environments.

The BlueXP Connector can then manage the resources and processes within your public cloud environment.

Before you begin the deployment process, ensure that you have completed the required preparations. This guide will enable you to focus on the minimum requirements for BlueXP Connector installation.

Permissions

Ensure that the Azure user or service principal you've provided has sufficient permissions

Authentication

Choose between two methods: an [Azure user account](#) or an [Active Directory service principal](#)

Networking

Ensure that you have details on the VNet and subnet in which the BlueXP Connector will reside

[Skip to Deployment](#)

Previous

Continue



5. Click on **Skip to Deployment** to configure your connector **Virtual Machine Authentication**. Add the SSH key pair you have created in Azure resource group during onboarding to BlueXP preparation for connector OS authentication.

1 VM Authentication 2 Details 3 Network 4 Security Group 5 Review

Virtual Machine Authentication

You are logged in with Azure user: [acao@netapp.com](#) | Tenant: Hybrid Cloud TME

Subscription

Hybrid Cloud TME Onprem

Location

South Central US

Resource Group

Create New Use Existing

Resource Group

ANFAVSRG

Authentication Method

Password Public Key

User Name

azureuser

Enter SSH Public Key

-----BEGIN RSA PRIVATE KEY----- MIIGSAIBAAKCA...

Previous

Next


6. Provide a name for the connector instance, select **Create** and accept default **Role Name** under **Details**, and choose the subscription for the Azure account.

1 VM Authentication 2 Details 3 Network 4 Security Group 5 Review

Details

Connector Instance Name

AzureConnector

 Add Tags to Connector Instance

Connector Role

Create Attach existing Manual

Role Name

BlueXP Operator-5519248

Subscriptions to apply with the role

Hybrid Cloud TME Onprem

Previous

Next

7. Configure networking with the proper **VNet**, **Subnet**, and disable **Public IP** but ensure that the connector has the internet access in your Azure environment.

The screenshot shows the 'Add BlueXP Connector - Azure' configuration wizard. The title bar includes 'Add BlueXP Connector - Azure' and a 'More Information' link with a close icon. Below the title bar is a progress indicator with five steps: 1. VM Authentication, 2. Details, 3. Network (current step), 4. Security Group, and 5. Review. The main content area is titled 'Network' and is divided into two columns: 'Connectivity' and 'Proxy Configuration (Optional)'. Under 'Connectivity', there are three dropdown menus: 'VNet' (set to 'ANFAVSVaI'), 'Subnet' (set to 'VM_Sub'), and 'Public IP' (set to 'Disable'). Under 'Proxy Configuration (Optional)', there is an 'HTTP Proxy' text input field with the example 'http://172.16.254.1:8080', a 'Define Credentials for this Proxy' dropdown, and an 'Upload a root certificate' dropdown. A notice at the bottom of the 'Connectivity' section states: 'Notice: Ensure that the subnet has internet connectivity through a NAT device or proxy server so that the Connector can communicate with Azure services.' At the bottom of the wizard, there are 'Previous' and 'Next' buttons, and a help icon in the bottom right corner.

8. Configure the **Security Group** for the connector that allows HTTP, HTTPS, and SSH access.

VM Authentication Details Network 4 Security Group 5 Review

Security Group

The security group must allow inbound HTTP, HTTPS and SSH access.

Assign a security group: Create a new security group Select an existing security group

HTTP (Port 80)	HTTPS (Port 443)	SSH (Port 22)
Source Type <input type="text" value="Anywhere"/>	Source Type <input type="text" value="Anywhere"/>	Source Type <input type="text" value="Anywhere"/>
Source (CIDR) <input type="text" value="0.0.0.0/0"/>	Source (CIDR) <input type="text" value="0.0.0.0/0"/>	Source (CIDR) <input type="text" value="0.0.0.0/0"/>

Previous

Next



- Review the summary page and click **Add** to start connector creation. It generally takes about 10 mins to complete deployment. Once completed, the connector instance VM appears in the Azure portal.

✓ VM Authentication ✓ Details ✓ Network ✓ Security Group **5** Review

Review

[Code for Terraform Automation](#)

BlueXP Connector Name	AzureConnector
Subscription	Hybrid Cloud TME Onprem
Location	South Central US
Resource Group	Existing - ANFAVSRG
Role	New - BlueXP Operator-5519248
Authentication Method	Password (user: azureuser)
VNet	ANFAVVal
Subnet	VM_Sub
Public IP	Enable
Proxy	None
Security Group	HTTP: 0.0.0.0/0, HTTPS: 0.0.0.0/0, SSH: 0.0.0.0/0

[Previous](#)

[Add](#)

10. After the connector is deployed, the newly created connector appears under **Connector** drop-down.

NetApp BlueXP

Account Automation-to... Workspace Azure-DB Connector AzureConnector

Canvas My working environments My estate Go to Tabular View

+ Add Working Environment Enable Services

Azure Blob Storage
20 Storage Accounts

Amazon S3
0 Buckets


Working Environments

Amazon S3
0 Buckets

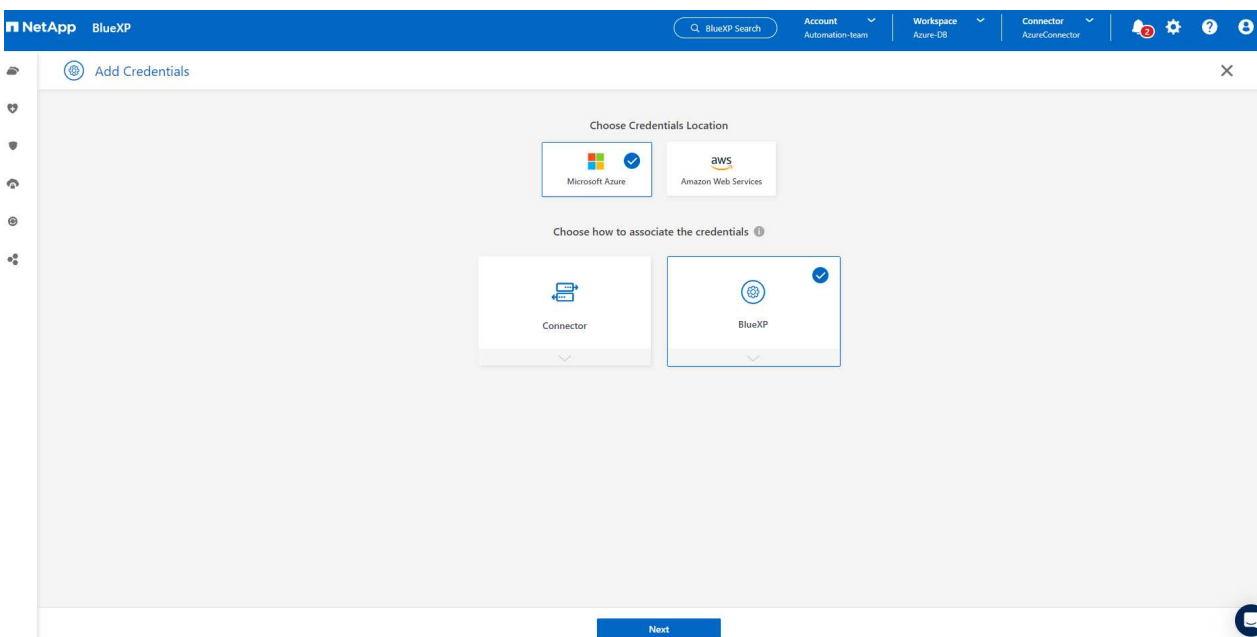
Azure Blob Storage
20 Storage Accounts

Define a credential in BlueXP for Azure resources access

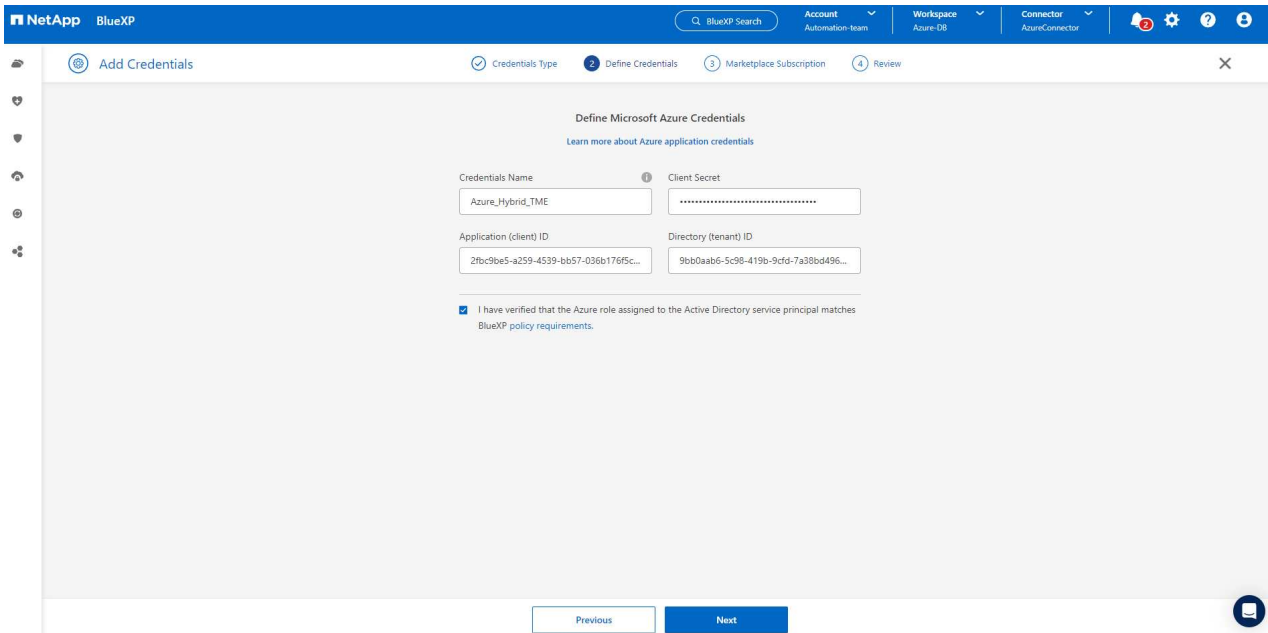
1. Click on setting icon on top right corner of BlueXP console to open **Account credentials** page, click **Add credentials** to start credential configuration workflow.



2. Choose credential location as - **Microsoft Azure - BlueXP**.



3. Define Azure credentials with proper **Client Secret**, **Client ID**, and **Tenant ID**, which should have been gathered during previous BlueXP onboarding process.



4. Review and **Add**.

image::snapctr_svcs_credential_04-azure.png["Screenshot showing this step in the GUI."]


5. You may also need to associate a **Marketplace Subscription** with the credential.

image::snapctr_svcs_credential_05-azure.png["Screenshot showing this step in the GUI."]

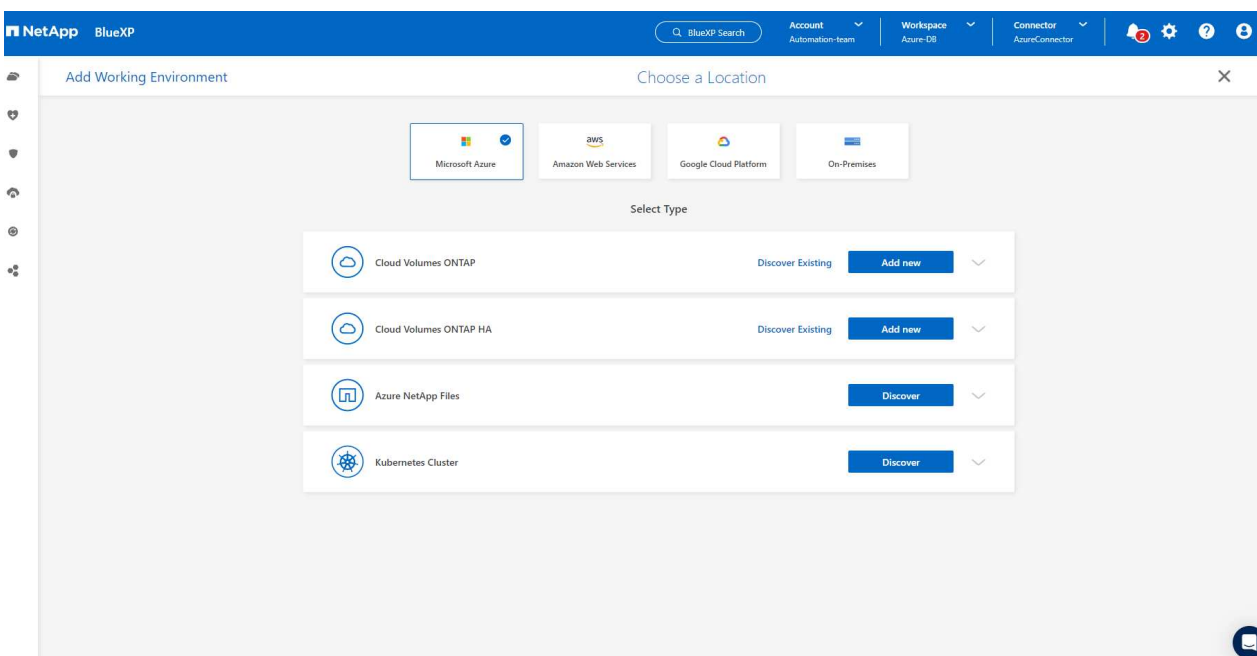
SnapCenter services setup

With the Azure credential configured, SnapCenter services can now be set up with the following procedures:


1. Back to Canvas page, from **My Working Environment** click **Add working Environment** to discover Azure NetApp Files deployed in Azure.




2. Choose **Microsoft Azure** as the location and click on **Discover**.




3. Name **Working Environment** and choose **Credential Name** created in previous section, and click **Continue**.



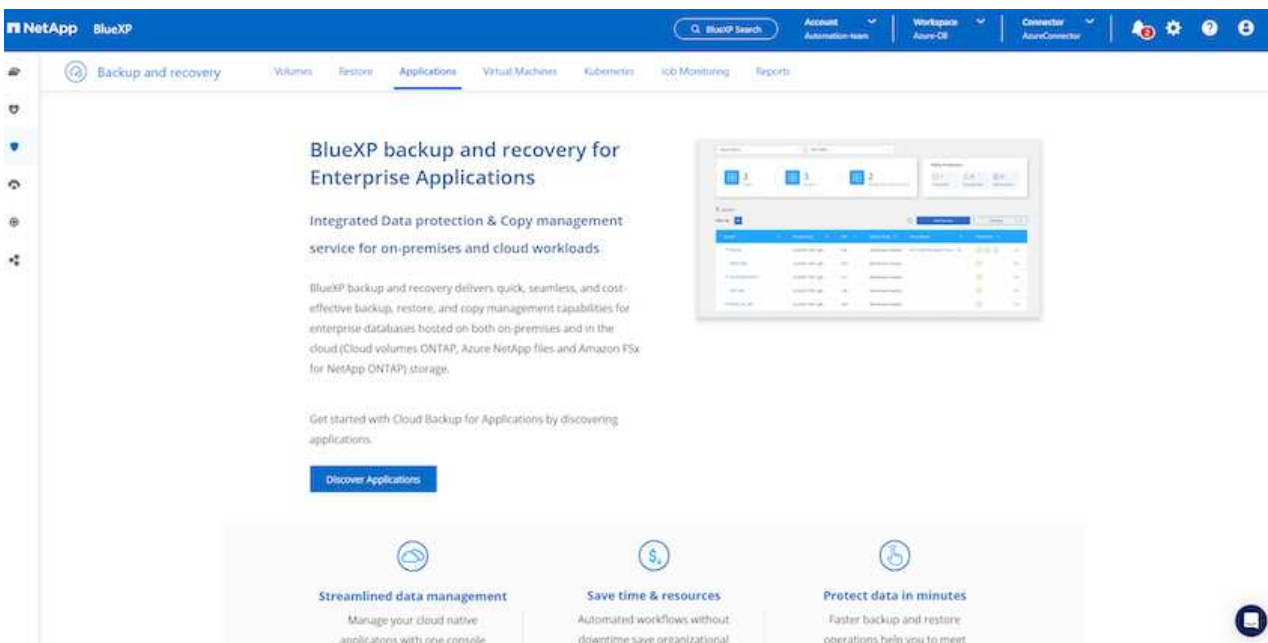
4. BlueXP console returns to **My working environments** and discovered Azure NetApp Files from Azure now appears on **Canvas**.




5. Click on **Azure NetApp Files** icon, then **Enter Working Environment** to view Oracle database volumes deployed in Azure NetApp Files storage.




6. From the left-hand sidebar of the console, hover your mouse over the protection icon, and then click **Protection > Applications** to open the Applications launch page. Click **Discover Applications**.



7. Select **Cloud Native** as the application source type.



8. Choose **Oracle** for the application type, click on **Next** to open host details page.



9. Select **Using SSH** and provide the Oracle Azure VM details such as **IP address**, **Connector**, Azure VM management **Username** such as azureuser. Click on **Add SSH Private Key** to paste in the SSH key pair that you used to deploy the Oracle Azure VM. You will also be prompted to confirm the fingerprint.

NetApp BlueXP

Discover Applications

1 Host Details 2 Configuration 3 Review

Select host type

Provide the following details to add host and discover applications

Host Installation Type Manual Using SSH

Host FQDN or IP: 172.30.137.142

Connector: AzureConnector

Username: azureuser

SSH Port: 22

Plug-in Port: 8145

Previous Next

Discover Applications

1 Host Details 2 Configuration 3 Review

Select host type

Provide the following details to add host and discover applications

Host Installation Type Manual Using SSH

Validate fingerprint

Algorithm: ssh-rsa


Fingerprint: AAAAE2VjZHNhLXNoYTItbmlzdHh0YTYAAAIbmlzdHh0YTYAAAB...

By proceeding further, I confirm that the above fingerprint for host is valid.

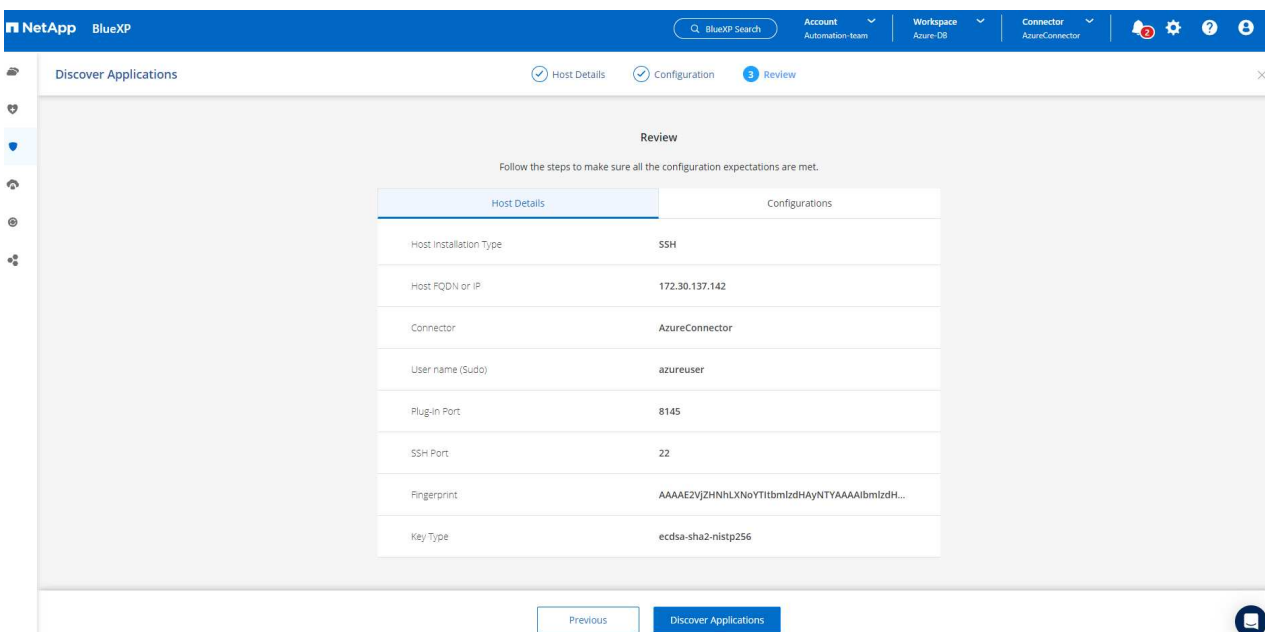
Proceed Cancel

Previous Next

10. Move on to next **Configuration** page to setup sudoer access on Oracle Azure VM.



11. Review and click on **Discover Applications** to install a plugin on the Oracle Azure VM and discover Oracle database on the VM in one step.



12. Discovered Oracle databases on Azure VM are added to **Applications**, and the **Applications** page lists the number of hosts and Oracle databases within the environment. The database **Protection Status** initially shows as **Unprotected**.

The screenshot displays the NetApp BlueXP interface for Oracle applications. At the top, there are navigation tabs for Backup and recovery, Volumes, Restore, Applications (selected), Virtual Machines, Kubernetes, Job Monitoring, and Reports. A search bar and account/workspace information are also visible. Below the navigation, there are filters for 'Cloud Native' and 'Oracle'. Three summary cards show: 3 Hosts, 3 ORACLE clones, and 0 Clones. An 'Application Protection' summary shows 0 Protected and 3 Unprotected items. A table titled '3 Databases' lists the following data:

Name	Host Name	Policy Name	Protection Status
NTAP	172.30.137.142		Unprotected
db1	172.30.15.99		Unprotected
db1st	172.30.15.124		Unprotected

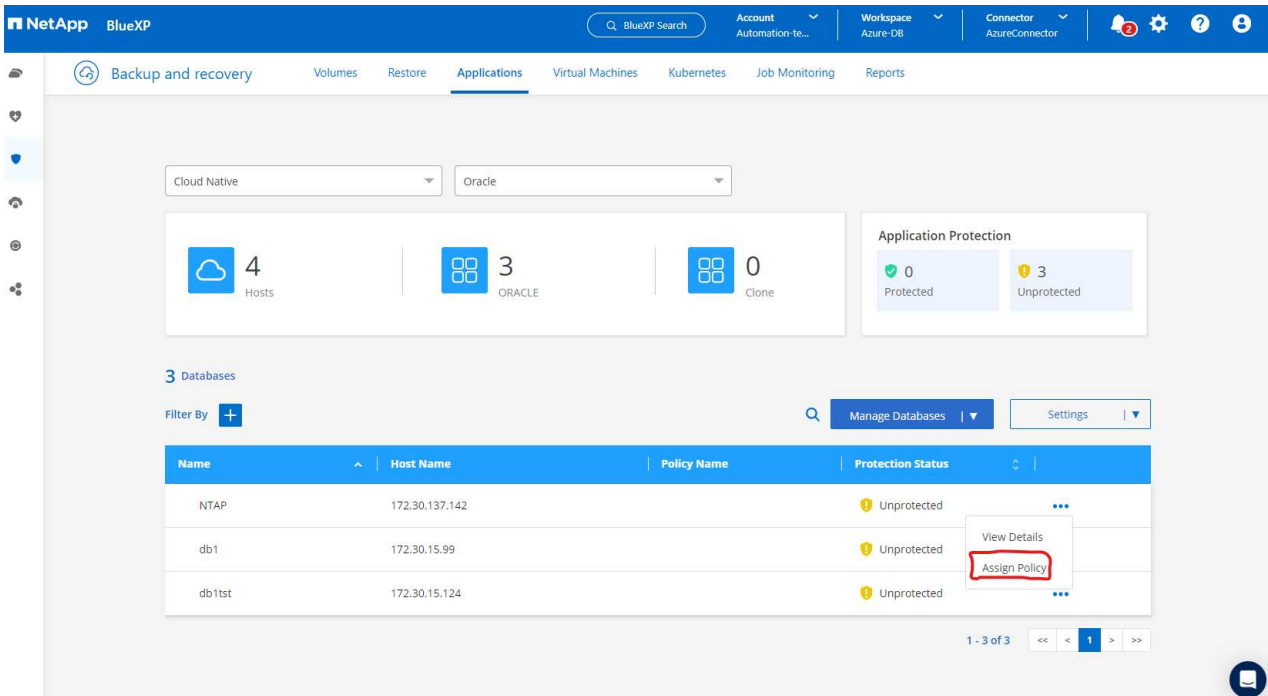
This completes the initial setup of SnapCenter services for Oracle. The next three sections of this document describe Oracle database backup, restore, and clone operations.

Oracle database backup

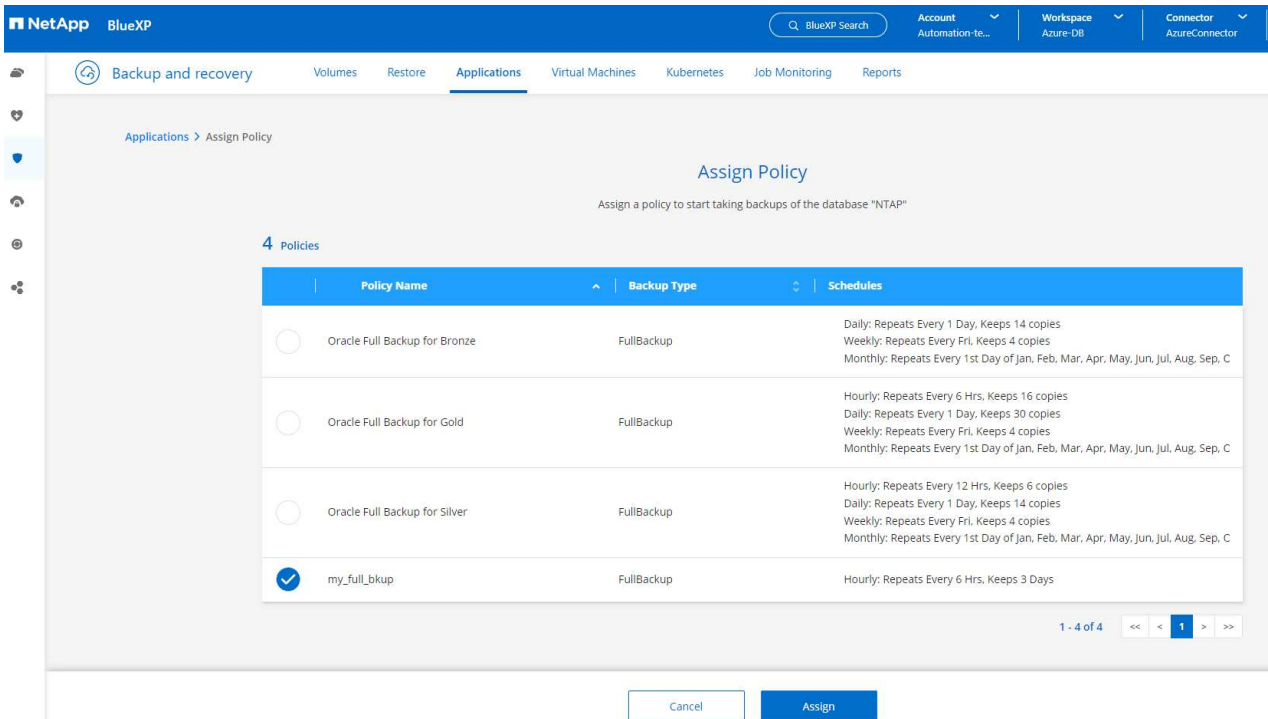
1. Our test Oracle database in Azure VM is configured with three volumes with an aggregate total storage about 1.6 TiB. This gives context about the timing for the snapshot backup, restore, and clone of a database of this size.

```
[oracle@acao-ora01 ~]$ df -h
Filesystem                Size      Used Avail Use% Mounted on
devtmpfs                  7.9G         0  7.9G   0% /dev
tmpfs                     7.9G         0  7.9G   0% /dev/shm
tmpfs                     7.9G      17M  7.9G   1% /run
tmpfs                     7.9G         0  7.9G   0% /sys/fs/cgroup
/dev/mapper/rootvg-rootlv 40G       23G   15G  62% /
/dev/mapper/rootvg-usrlv  9.8G     1.6G   7.7G  18% /usr
/dev/sda2                 496M     115M  381M  24% /boot
/dev/mapper/rootvg-varlv  7.9G     787M   6.7G  11% /var
/dev/mapper/rootvg-homelv 976M     323M   586M  36% /home
/dev/mapper/rootvg-optlv  2.0G     9.6M   1.8G   1% /opt
/dev/mapper/rootvg-tmplv  2.0G     22M   1.8G   2% /tmp
/dev/sda1                 500M     6.8M  493M   2% /boot/efi
172.30.136.68:/ora01-u01 100G     23G   78G  23% /u01
172.30.136.68:/ora01-u03 500G    117G  384G  24% /u03
172.30.136.68:/ora01-u02 1000G   804G  197G  81% /u02
tmpfs                     1.6G         0  1.6G   0% /run/user/1000
[oracle@acao-ora01 ~]$
```


1. To protect database, click the three dots next to the database **Protection Status**, and then click **Assign Policy** to view the default preloaded or user defined database protection policies that can be applied to your Oracle databases. Under **Settings - Policies**, you have option to create your own policy with a customized backup frequency and backup data-retention window.



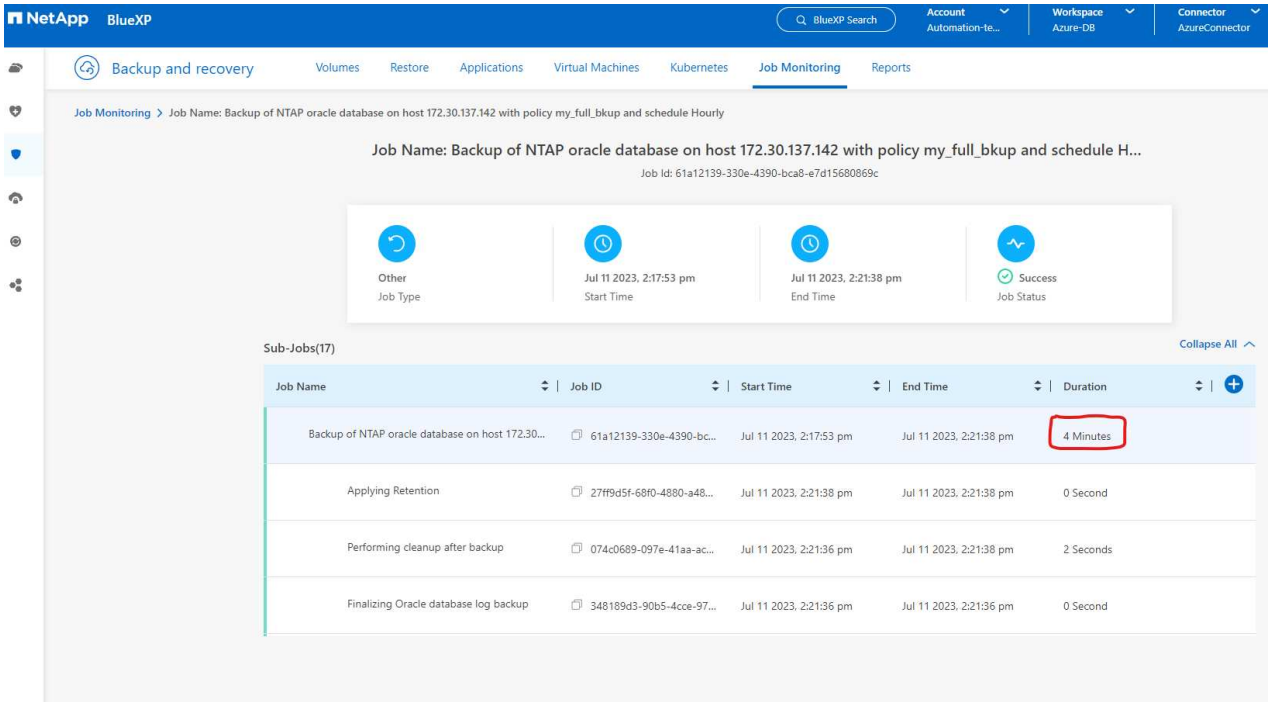
2. When you are happy with the policy configuration, you can then **Assign** your policy of choice to protect the database.




3. After the policy is applied, the database protection status changed to **Protected** with a green check mark. BlueXP executes the snapshot backup according to the schedule defined. In addition, **ON-Demand Backup** is available from the three-dot drop down menu as shown below.




4. From **Job Monitoring** tab, backup job details can be viewed. Our test results showed that it took about 4 minutes to backup an Oracle database about 1.6 TiB.



5. From three-dot drop down menu **View Details**, you can view the backup sets created from snapshot backup.

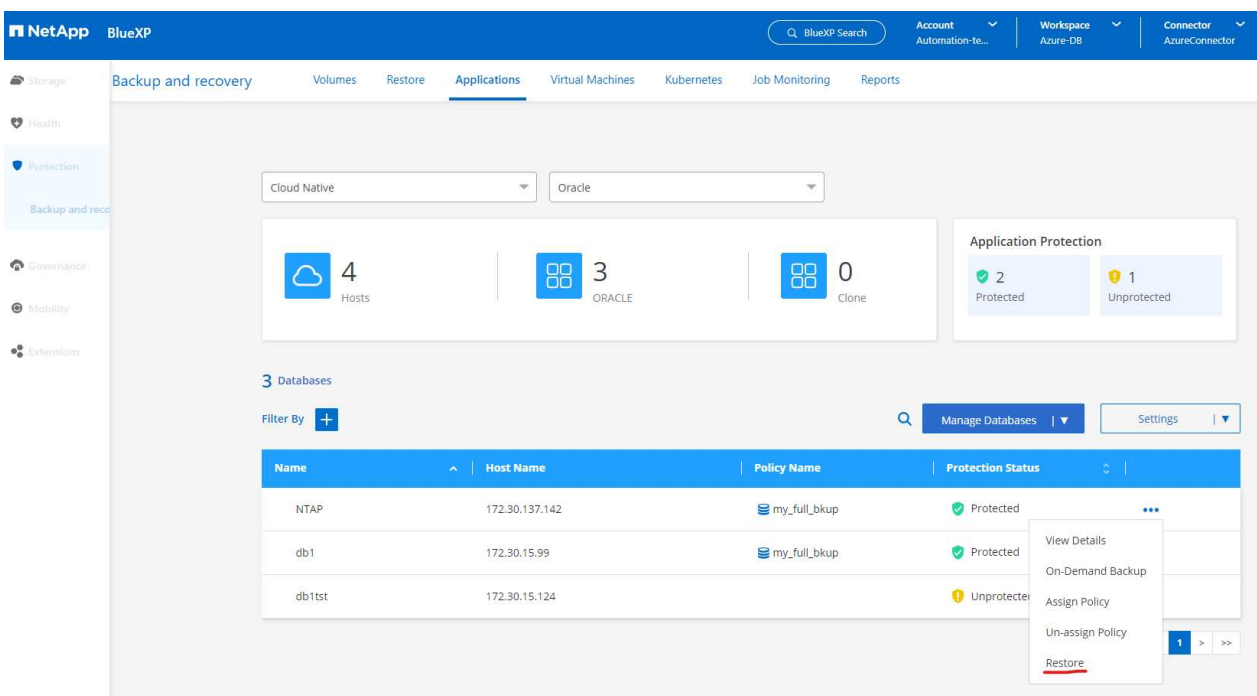


6. Database backup details include the **Backup Name**, **Backup Type**, **SCN**, **RMAN Catalog**, and **Backup Time**. A backup set contains application-consistent snapshots for data volume and log volume respectively. A log volume snapshot takes place right after a database data volume snapshot. You could apply a filter if you are looking for a particular backup in the backup list.




Oracle database restore and recovery


1. For a database restore, click the three-dot drop down menu for the particular database to be restored in **Applications**, then click **Restore** to initiate database restore and recovery workflow.



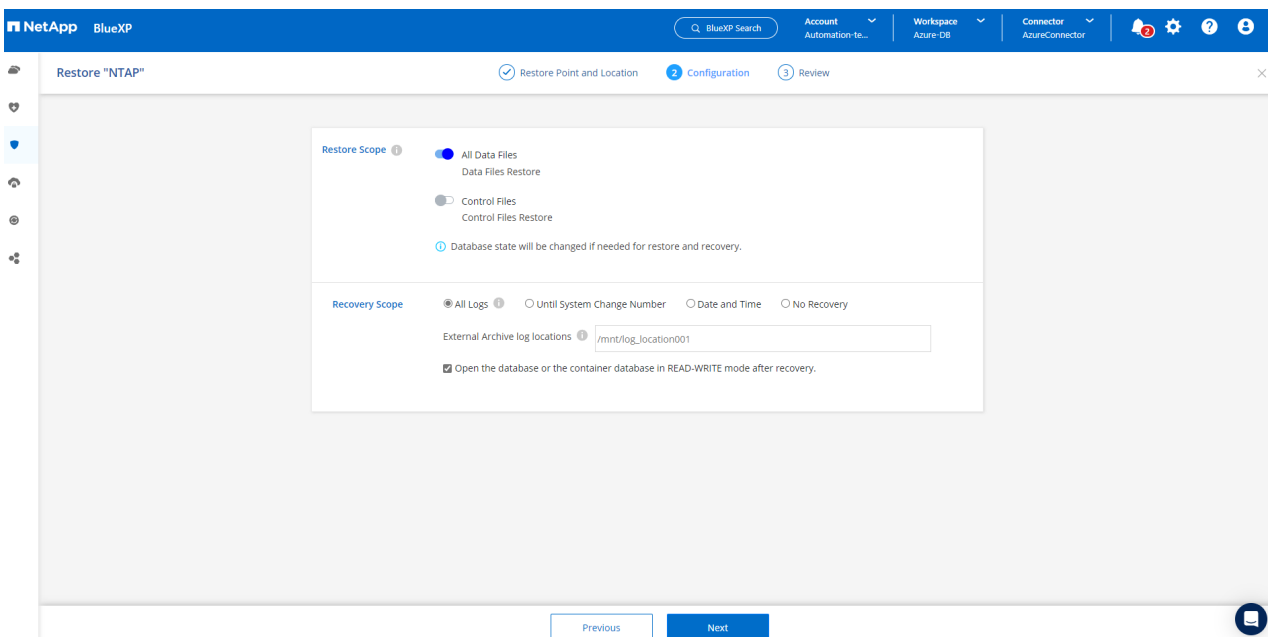
2. Choose your **Restore Point** by time stamp. Each time stamp in the list represents an available database backup set.




3. Choose your **Restore Location** to **original location** for an Oracle database in place restore and recovery.




4. Define your **Restore Scope**, and **Recovery Scope**. All Logs mean a full recovery up to date including current logs.



5. Review and **Restore** to start database restore and recovery.



6. From the **Job Monitoring** tab, we observed that it took 2 minutes to run a full database restore and recovery up to date.



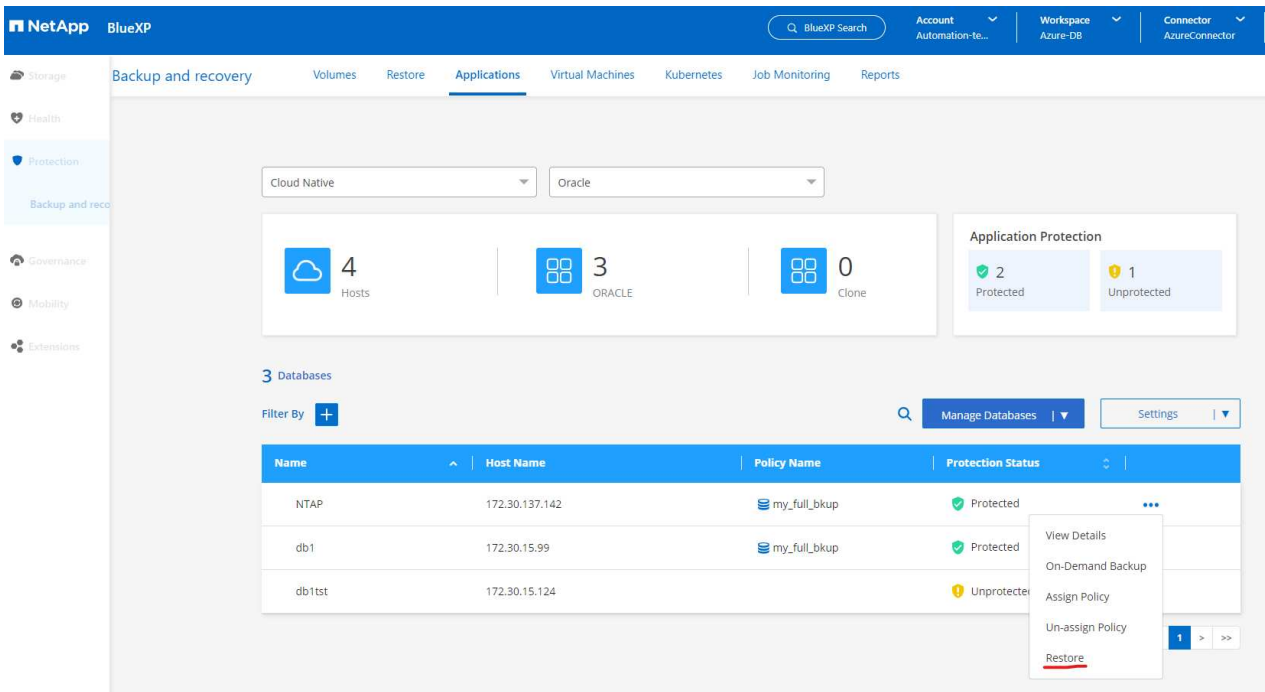
Oracle database clone

Database clone procedures are similar to restore but to an alternate Azure VM with identical Oracle software stack pre-installed and configured.

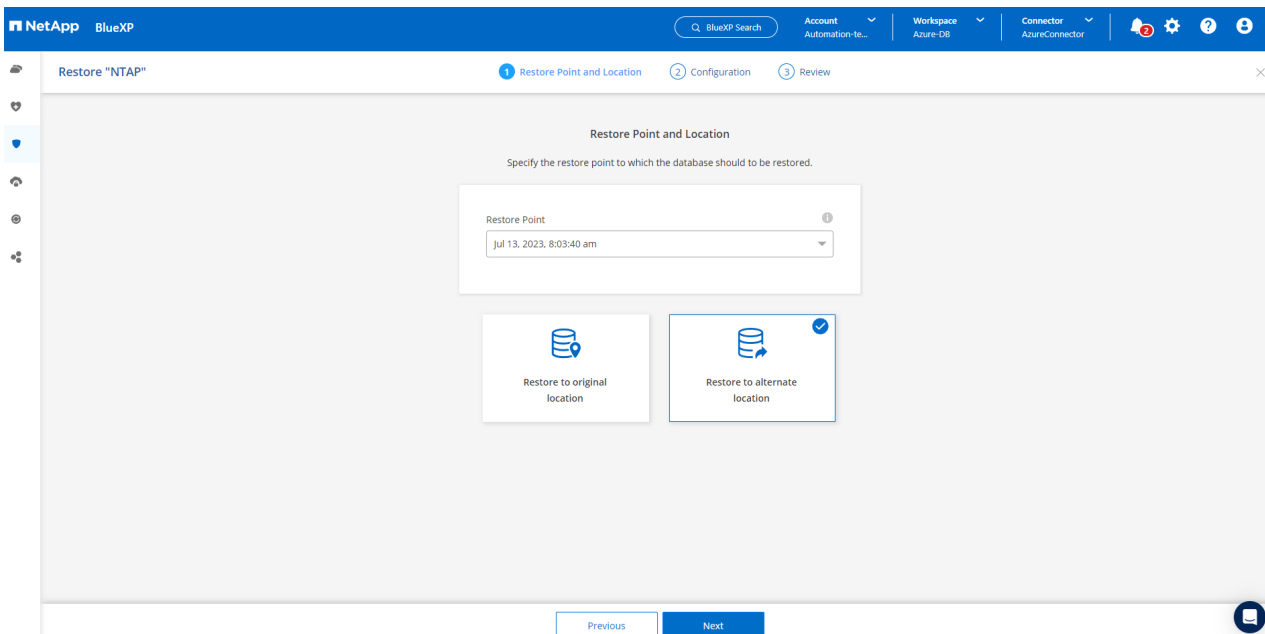


Ensure that your Azure NetApp File storage has sufficient capacity for a cloned database the same size as the primary database to be cloned. The alternate Azure VM has been added to **Applications**.

1. Click the three-dot drop down menu for the particular database to be cloned in **Applications**, then click **Restore** to initiate clone workflow.




2. Select the **Restore Point** and check the **Restore to alternate location**.

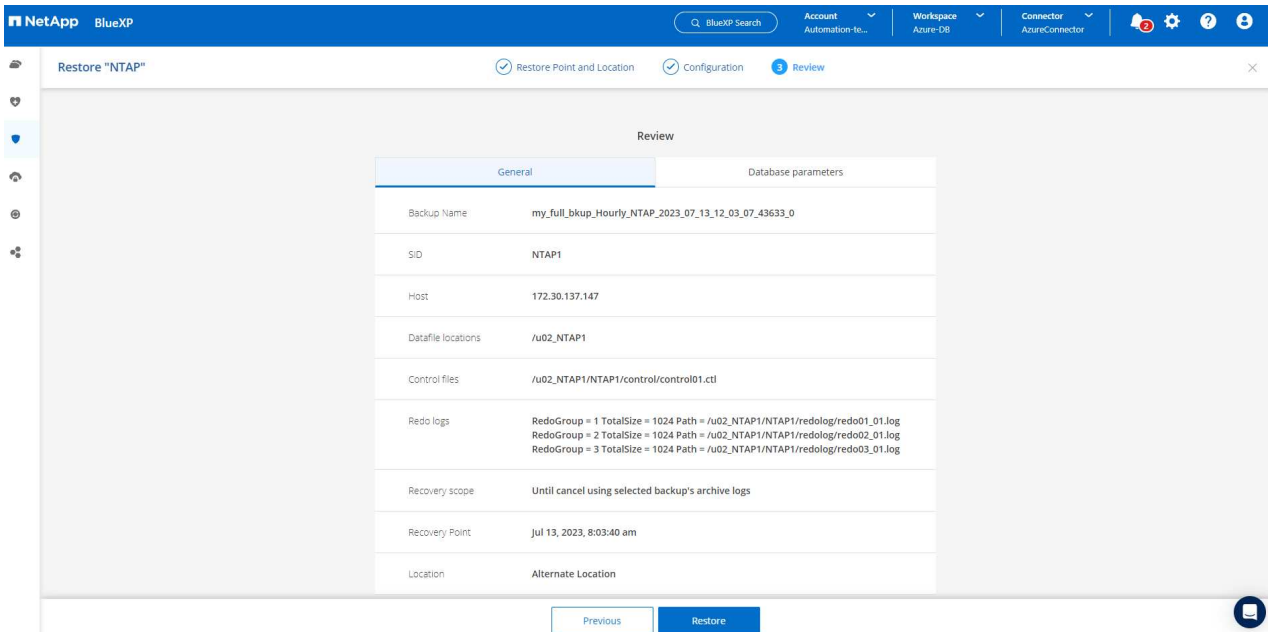


3. In the next **Configuration** page, set alternate **Host**, new database **SID**, and **Oracle Home** as

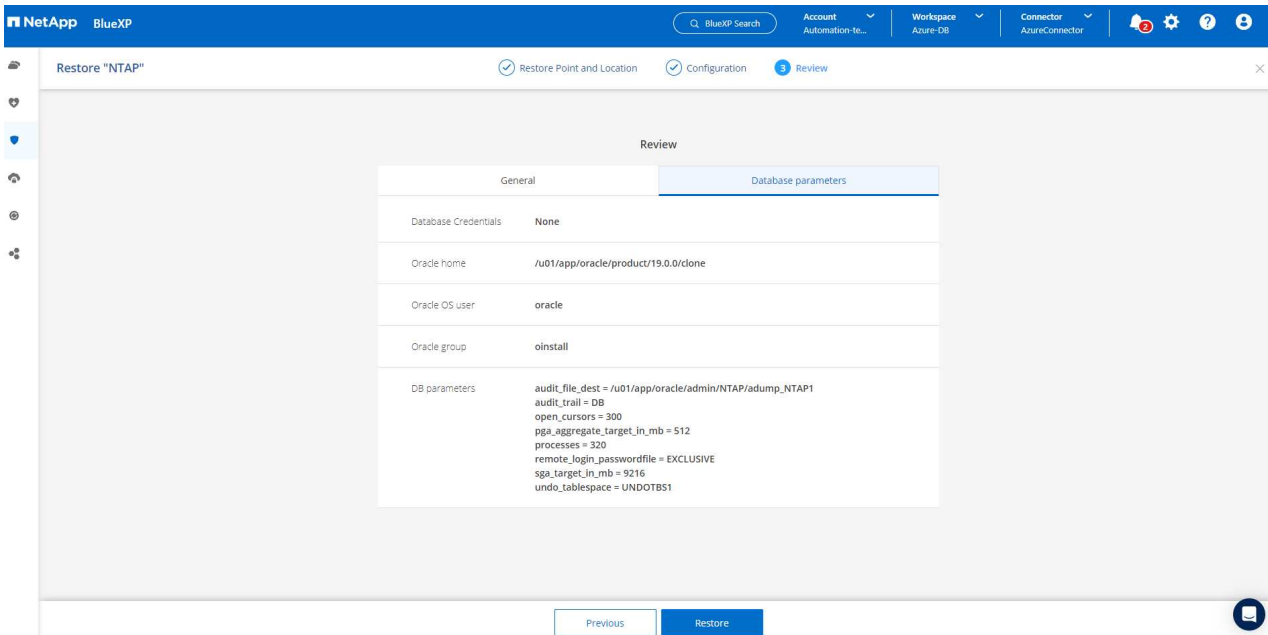
configured at alternate Azure VM.




4. Review **General** page shows the details of cloned database such as SID, alternate host, data file locations, recovery scope etc.



5. Review **Database parameters** page shows the details of cloned database configuration as well as some database parameters setting.



6. Monitor the cloning job status from the **Job Monitoring** tab, we observed that it took 8 minutes to clone a 1.6 TiB Oracle database.



7. Validate the cloned database in BlueXP **Applications** page that showed the cloned database was immediately registered with BlueXP.

NetApp BlueXP

Account Automation-te... Workspace Azure-DB Connector AzureConnector

Backup and recovery Volumes Restore Applications Virtual Machines Kubernetes Job Monitoring Reports

Cloud Native Oracle

4 Hosts 4 ORACLE 0 Clone

Application Protection 2 Protected 2 Unprotected

4 Databases

Filter By + Manage Databases Settings

Name	Host Name	Policy Name	Protection Status
NTAP	172.30.137.142	my_full_bkup	Protected
NTAP1	172.30.137.147	my_full_bkup	Unprotected
db1	172.30.15.99	my_full_bkup	Protected
db1tst	172.30.15.124	my_full_bkup	Unprotected

1 - 4 of 4

8. Validate the cloned database on the Oracle Azure VM that showed the cloned database was running as expected.

```

[oracle@acao-ora02 admin]$ 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 $ORACLE_SID are not allowed.
#
#
# SnapCenter Plug-in for Oracle Database generated entry (DO NOT REMOVE THIS LINE)
NTAP1:/u01/app/oracle/product/19.0.0/clone:N
[oracle@acao-ora02 admin]$ export ORACLE_SID=NTAP1
[oracle@acao-ora02 admin]$ export ORACLE_HOME=/u01/app/oracle/product/19.0.0/clone
[oracle@acao-ora02 admin]$ export PATH=$PATH:$ORACLE_HOME/bin
[oracle@acao-ora02 admin]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Thu Jul 13 17:16:31 2023
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$databases;

NAME          OPEN_MODE          LOG_MODE
-----
NTAP1         READ WRITE         NOARCHIVELOG

```

This completes the demonstration of an Oracle database backup, restore, and clone in Azure with NetApp BlueXP console using SnapCenter Service.

Additional information

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

- Set up and administer BlueXP

<https://docs.netapp.com/us-en/cloud-manager-setup-admin/index.html>

- BlueXP backup and recovery documentation

<https://docs.netapp.com/us-en/cloud-manager-backup-restore/index.html>

- Azure NetApp Files

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

- Get started with Azure

<https://azure.microsoft.com/en-us/get-started/>

TR-4964: Oracle Database backup, restore and clone with SnapCenter Services - AWS

This solution provides overview and details for Oracle database backup, restore, clone using NetApp SnapCenter SaaS using BlueXP console in Azure cloud.

Allen Cao, Niyaz Mohamed, NetApp

Purpose

SnapCenter Services is the SaaS version of the classic SnapCenter database management UI tool that is available through the NetApp BlueXP cloud management console. It is an integral part of the NetApp cloud-backup, data-protection offering for databases such as Oracle and HANA running on NetApp cloud storage. This SaaS-based service simplifies traditional SnapCenter standalone server deployment that generally requires a Windows server operating in a Windows domain environment.

In this documentation, we demonstrate how you can set up SnapCenter Services to backup, restore, and clone Oracle databases deployed to Amazon FSx for ONTAP storage and EC2 compute instances. Although it is much easier to set up and use, SnapCenter Services deliver key functionalities that are available in the legacy SnapCenter UI tool.

This solution addresses the following use cases:

- Database backup with snapshots for Oracle databases hosted in Amazon FSx for ONTAP
- Oracle database recovery in the case of a failure
- Fast and storage-efficient cloning of primary databases for a dev/test environment or other use cases

Audience

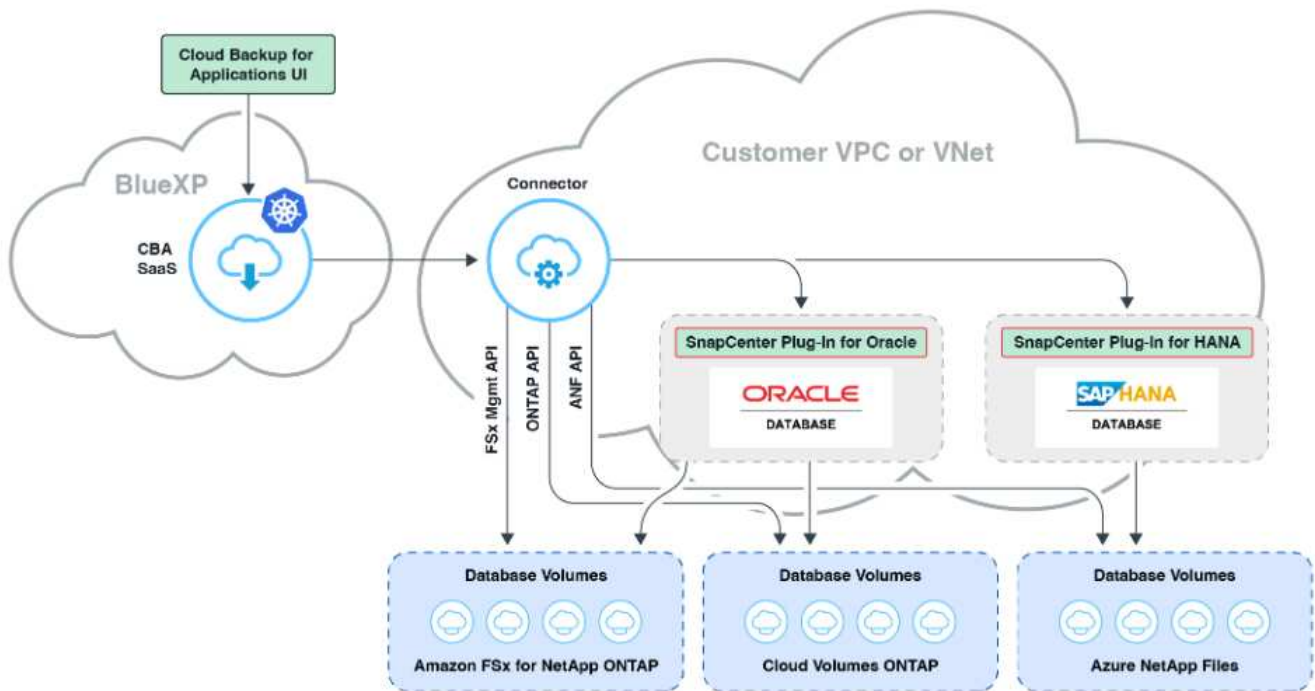
This solution is intended for the following audiences:

- The DBA who manages Oracle databases running on Amazon FSx for ONTAP storage
- The solution architect who is interested in testing Oracle database backup, restore, and clone in the public AWS cloud
- The storage administrator who supports and manages the Amazon FSx for ONTAP storage
- The application owner who owns applications that are deployed to Amazon FSx for ONTAP storage

Solution test and validation environment

The testing and validation of this solution was performed in an AWS FSx and EC2 environment that might not match the final deployment environment. For more information, see the section [Key factors for deployment consideration](#).

Architecture



This image provides a detailed picture of BlueXP backup and recovery for applications within the BlueXP console, including the UI, the connector, and the resources it manages.

Hardware and software components

Hardware

FSx ONTAP storage	Current version offered by AWS	One FSx HA cluster in the same VPC and availability zone
EC2 instance for compute	t2.xlarge/4vCPU/16G	Two EC2 T2 xlarge EC2 instances, one as primary DB server and the other as clone DB server

Software

RedHat Linux	RHEL-8.6.0_HVM-20220503-x86_64-2-Hourly2-GP2	Deployed RedHat subscription for testing
Oracle Grid Infrastructure	Version 19.18	Applied RU patch p34762026_190000_Linux-x86-64.zip
Oracle Database	Version 19.18	Applied RU patch p34765931_190000_Linux-x86-64.zip
Oracle OPatch	Version 12.2.0.1.36	Latest patch p6880880_190000_Linux-x86-64.zip

Key factors for deployment consideration

- **Connector to be deployed in the same VPC as database and FSx.** When possible, the connector should be deployed in the same AWS VPC, which enables connectivity to the FSx storage and the EC2 compute instance.
- **An AWS IAM policy created for SnapCenter connector.** The policy in JSON format is available in the detailed SnapCenter service documentation. When you launch connector deployment with the BlueXP console, you are also prompted to set up the prerequisites with details of required permission in JSON format. The policy should be assigned to the AWS user account that owns the connector.
- **The AWS account access key and the SSH key pair created in the AWS account.** The SSH key pair is assigned to the ec2-user for logging into the connector host and then deploying a database plug-in to the EC2 DB server host. The access key grants permission for provisioning the required connector with IAM policy above.
- **A credential added to the BlueXP console setting.** To add Amazon FSx for ONTAP to the BlueXP working environment, a credential that grants BlueXP permissions to access Amazon FSx for ONTAP is set up in the BlueXP console setting.
- **java-11-openjdk installed on the EC2 database instance host.** SnapCenter service installation requires java version 11. It needs to be installed on application host before plugin deployment attempt.

Solution deployment

There is extensive NetApp documentation with a broader scope to help you protect your cloud-native application data. The goal of this documentation is to provide step-by-step procedures that cover SnapCenter Service deployment with the BlueXP console to protect your Oracle database deployed to Amazon FSx for ONTAP and an EC2 compute instance. This document fills in certain details that might be missing from more general instructions.

To get started, complete the following steps:

- Read the general instructions [Protect your cloud native applications data](#) and the sections related to Oracle and Amazon FSx for ONTAP.
- Watch the following video walkthrough.

[Solution Deployment](#)

Prerequisites for SnapCenter service deployment

Deployment requires the following prerequisites.

1. A primary Oracle database server on an EC2 instance with an Oracle database fully deployed and running.
2. An Amazon FSx for ONTAP cluster deployed in AWS that is hosting the database volumes above.
3. An optional database server on an EC2 instance that can be used for testing the cloning of an Oracle database to an alternate host for the purpose of supporting a dev/test workload or any use cases that requires a full data set of a production Oracle database.
4. If you need help to meet the above prerequisites for Oracle database deployment on Amazon FSx for ONTAP and EC2 compute instance, see [Oracle Database Deployment and Protection in AWS FSx/EC2 with iSCSI/ASM](#) or white paper [Oracle Database Deployment on EC2 and FSx Best Practices](#)

Onboarding to BlueXP preparation

1. Use the link [NetApp BlueXP](#) to sign up for BlueXP console access.
2. Login to your AWS account to create an IAM policy with proper permissions and assign the policy to the AWS account that will be used for BlueXP connector deployment.

The screenshot shows the AWS IAM console interface. On the left is the 'Identity and Access Management (IAM)' navigation menu. The main content area is titled 'Policies > snapcenter' and shows the 'Summary' page for a policy. The policy details include:

- Policy ARN:** arn:aws:iam::541696183547:policy/snapcenter
- Description:** Policy to grant snapcenter service permission to create connector in AWS.

Below the summary, there are tabs for 'Permissions', 'Policy usage', 'Tags', 'Policy versions', and 'Access Advisor'. The 'Permissions' tab is active, showing a 'Policy summary' section with a JSON string for the policy's permissions:

```

1 {
2   "Version": "2012-10-17",
3   "Statement": [
4     {
5       "Effect": "Allow",
6       "Action": [
7         "iam:CreateRole",
8         "iam>DeleteRole",
9         "iam:PutRolePolicy",
10        "iam:CreateInstanceProfile",
11        "iam>DeleteRolePolicy",
12        "iam:AddRoleToInstanceProfile",
13        "iam:RemoveRoleFromInstanceProfile",
14        "iam>DeleteInstanceProfile",
15        "iam:PassRole",
16        "iam:ListRoles",
17        "ec2:DescribeInstanceStatus",
18        "ec2:RunInstances",
19        "ec2:ModifyInstanceAttribute",
20        "ec2:CreateSecurityGroup",
21        "ec2>DeleteSecurityGroup",
22        "ec2:DescribeSecurityGroups",
23        "ec2:RevokeSecurityGroupEgress",
24        "ec2:AuthorizeSecurityGroupEgress",
25        "ec2:AuthorizeSecurityGroupIngress",
26        "ec2:RevokeSecurityGroupIngress",
27        "ec2:CreateNetworkInterface",
28        "ec2:DescribeNetworkInterfaces"


```

The policy should be configured with a JSON string that is available in NetApp documentation. The JSON string can also be retrieved from the page when connector provisioning is launched and you are prompted for the prerequisites permissions assignment.

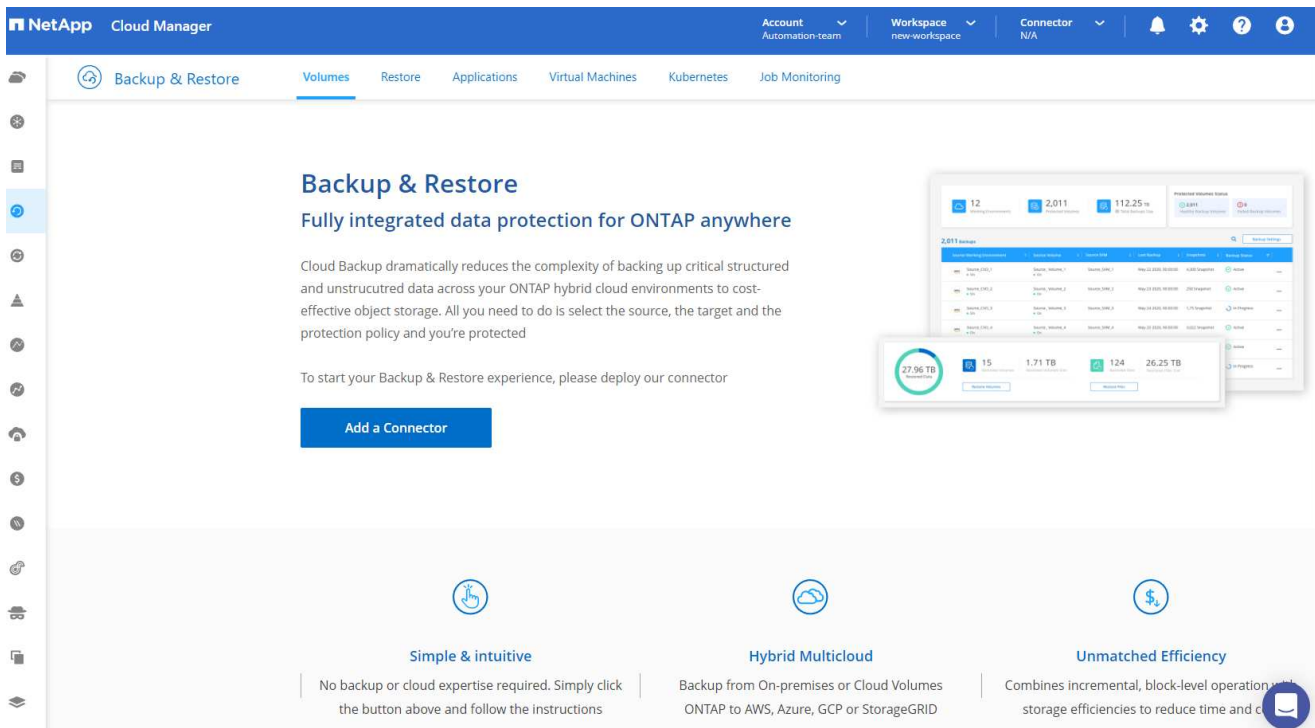
3. You also need the AWS VPC, subnet, security group, an AWS user account access key and secrets, an SSH key for ec2-user, and so on ready for connector provisioning.

Deploy a connector for SnapCenter services


1. Login to the BlueXP console. For a shared account, it is a best practice to create an individual workspace by clicking **Account > Manage Account > Workspace** to add a new workspace.



2. Click **Add a Connector** to launch the connector provisioning workflow.



1. Choose your cloud provider (in this case, **Amazon Web Services**).



1. Skip the **Permission**, **Authentication**, and **Networking** steps if you already have them set up in your AWS account. If not, you must configure these before proceeding. From here, you could also retrieve the permissions for the AWS policy that is referenced in the previous section "[Onboarding to BlueXP preparation.](#)"

Deploying a Connector

The Connector is a crucial component for the day-to-day use of Cloud Manager. It's used to connect Cloud Manager's services to your hybrid-cloud environments. The Connector can then manage the resources and processes within your public cloud environment.

Before you begin the deployment process, ensure that you have completed the required preparations. This guide will enable you to focus on the minimum requirements for Connector installation.

Permissions Set up an IAM role with the required permissions	Authentication Choose between two AWS authentication methods: AWS keys or assuming an IAM role	Networking Obtain details about the VPC and subnet in which the Connector will reside
--	--	---

[Skip to Deployment](#)



1. Enter your AWS account authentication with **Access Key** and **Secret Key**.

- 1 AWS Credentials
- 2 Details
- 3 Network
- 4 Security Group
- 5 Review

AWS Authentication

Region
us-east-1 | US East (N. Virginia) ▼

Select the Authentication Method: Assume Role AWS Keys

AWS Access Key
AKIA6JRXA6ZVGVFSHMO3

AWS Secret Key
.....

Want to launch an instance without AWS Credentials? ▼



2. Name the connector instance and select **Create Role** under **Details**.

Add Connector - AWS More Information ✕

1 AWS Credentials 2 **Details** 3 Network 4 Security Group 5 Review

Details

Connector Instance Name i

+ Add Tags to Connector Instance

Connector Role i
 Create Role Select an existing Role

Role Name

AWS Managed Encryption i
Master Key: aws/ebs (default) [Change Key](#)

Previous Next 🗨️

1. Configure networking with the proper **VPC**, **Subnet**, and SSH **Key Pair** for connector access.

- AWS Credentials
- Details
- 3 Network**
- 4 Security Group
- 5 Review

Network

Connectivity

VPC

vpc-0b522d5e982a50ceb - 172.30.15.0/25

Subnet

172.30.15.0/25 | priv-subnet-01

Key Pair

sufi_new

Public IP

Use subnet settings (Disable)

Notice: Ensure that the subnet has internet connectivity through a NAT device or proxy server so that the Connector can communicate with AWS services.

Proxy Configuration (Optional)

HTTP Proxy

Example: http://172.16.254.1:8080

Define Credentials for this Proxy

Upload a root certificate

Previous

Next



2. Set the Security Group for the connector.

- AWS Credentials
- Details
- Network
- 4 Security Group**
- 5 Review

Security Group

The security group must allow inbound HTTP, HTTPS and SSH access.

Assign a security group: Create a new security group Select an existing security group

1 Security Group

Security Group Name	Description
<input checked="" type="radio"/> default	default VPC security group

Previous

Next




3. Review the summary page and click **Add** to start connector creation. It generally takes about 10 mins to complete deployment. Once completed, the connector instance appears in the AWS EC2 dashboard.

Add BlueXP Connector - AWS More Information ×

✓ AWS Credentials ✓ Details ✓ Network ✓ Security Group **5** Review

Review [Code for Terraform Automation](#)

BlueXP Connector Name	aws-snapctr-us-east
AWS Access Key	AKIAX4H43ZT56IWWR3TI
Region	us-east-1
VPC	vpc-0b522d5e982a50ceb - 172.30.15.0/25
Subnet	172.30.15.0/25 priv-subnet-01
Key Pair	sufi_new
Public IP	Use subnet settings (Disable)
Proxy	None
Security Group	default

Previous Add 

Define a credential in BlueXP for AWS resources access

1. First, from AWS EC2 console, create a role in **Identity and Access Management (IAM)** menu **Roles, Create role** to start role creation workflow.

The screenshot shows the AWS IAM console 'Roles' page. The left sidebar contains the 'Identity and Access Management (IAM)' menu with 'Roles' selected. The main content area displays a list of roles with columns for 'Role name', 'Trusted entities', and 'Last activity'. The 'Create role' button is visible in the top right corner.

Role name	Trusted entities	Last activity
AmazonEC2RoleforAWSWizard	AWS Service: ec2	-
AmazonS3RoleforInstancesQuickSetup	AWS Service: ec2	156 days ago
aws-controltower-AdministratorExecutionRole	Account: 982617961887	-
aws-controltower-ConfigRecorderRole	AWS Service: config	-
aws-controltower-ForwarderStatisticsRole	AWS Service: lambda	2 days ago
aws-controltower-ReadOnlyExecutionRole	Account: 982617961887	-
AWS-QuickSetup-StackSet-Local-AdministrationRole	AWS Service: cloudformation	354 days ago
AWS-QuickSetup-StackSet-Local-ExecutionRole	Account: 541696183547	354 days ago
AWSControlTowerExecution	Account: 292306980405	232 days ago
AWSReservedSSO_AWSAdministratorAccess_30e65a909921a	Identity Provider: am:aws:iam:541696183547:saml-provider/AWSSSO_91b222f38b29f441_DO_NOT_DELETE	26 days ago
AWSReservedSSO_AWSOrganizationsAccess_6c3be705667e2b3	Identity Provider: am:aws:iam:541696183547:saml-provider/AWSSSO_91b222f38b29f441_DO_NOT_DELETE	-
AWSReservedSSO_AWSPowerUserAccess_50950a9f640ed1	Identity Provider: am:aws:iam:541696183547:saml-provider/AWSSSO_91b222f38b29f441_DO_NOT_DELETE	-
AWSReservedSSO_AWSReadOnlyAccess_234340b7e71bb118	Identity Provider: am:aws:iam:541696183547:saml-provider/AWSSSO_91b222f38b29f441_DO_NOT_DELETE	-
AWSReservedSSO_SAA-Dev-ReadOnly_85e81e993e815e7	Identity Provider: am:aws:iam:541696183547:saml-provider/AWSSSO_91b222f38b29f441_DO_NOT_DELETE	-

2. In **Select trusted entity** page, choose **AWS account, Another AWS account**, and paste in the BlueXP account ID, which can be retrieved from BlueXP console.

The screenshot shows the AWS IAM console 'Create role' page, Step 1: Select trusted entity. The 'AWS account' option is selected. Under 'An AWS account', 'Another AWS account' is selected, and the account ID '952013314444' is entered in the text field.

Select trusted entity

Trusted entity type

- AWS service
- AWS account
- Web identity
- SAML 2.0 federation
- Custom trust policy

An AWS account

This account (541696183547)

Another AWS account

Account ID
 Identifier of the account that can use this role

Account ID is a 12-digit number.

Options

- Require external ID (Best practice when a third party will assume this role)
- Require MFA

Cancel

3. Filter permission policies by fsx and add **Permissions policies** to the role.

Step 1
Select trusted entityStep 2
Add permissionsStep 3
Name, review, and createAdd permissions InfoPermissions policies (Selected 1/889) Info

Choose one or more policies to attach to your new role.

 4 matches

<input type="checkbox"/>	Policy name <small>Info</small>	Type	Description
<input type="checkbox"/>	AmazonFSxReadOnlyAccess	AWS ma...	Provides read only access to Amazon FSx.
<input checked="" type="checkbox"/>	AmazonFSxFullAccess	AWS ma...	Provides full access to Amazon FSx and access to related AWS services.
<input type="checkbox"/>	AmazonFSxConsoleReadOnlyAccess	AWS ma...	Provides read only access to Amazon FSx and access to related AWS services via the AWS Management Console.
<input type="checkbox"/>	AmazonFSxConsoleFullAccess	AWS ma...	Provides full access to Amazon FSx and access to related AWS services via the AWS Management Console.

Set permissions boundary - optional Info

Set a permissions boundary to control the maximum permissions this role can have. This is not a common setting, but you can use it to delegate permission management to others.

4. In **Role details** page, name the role, add a description, then click **Create role**.

Step 1
Select trusted entityStep 2
Add permissionsStep 3
Name, review, and create

Name, review, and create

Role details

Role name

Enter a meaningful name to identify this role.

Maximum 64 characters. Use alphanumeric and '+, -, @, _' characters.

Description

Add a short explanation for this role.

Maximum 1000 characters. Use alphanumeric and '+, -, @, _' characters.

Step 1: Select trusted entities

```

1- {
2-   "Version": "2012-10-17",
3-   "Statement": [
4-     {
5-       "Effect": "Allow",
6-       "Action": "sts:AssumeRole",
7-       "Principal": {
8-         "AWS": "992013314444"
9-       },
10-      "Condition": {}
11-     }
12-   ]
13- }

```

5. Back to BlueXP console, click on setting icon on top right corner of the console to open **Account credentials** page, click **Add credentials** to start credential configuration workflow.

NetApp BlueXP

BlueXP Search Account Automation-16... Workspace Database-2 Connector acio-aws-conn...

Credentials Account credentials User credentials


BlueXP and the Connector use account-level credentials to deploy and manage resources in your cloud environment.

5 Credentials

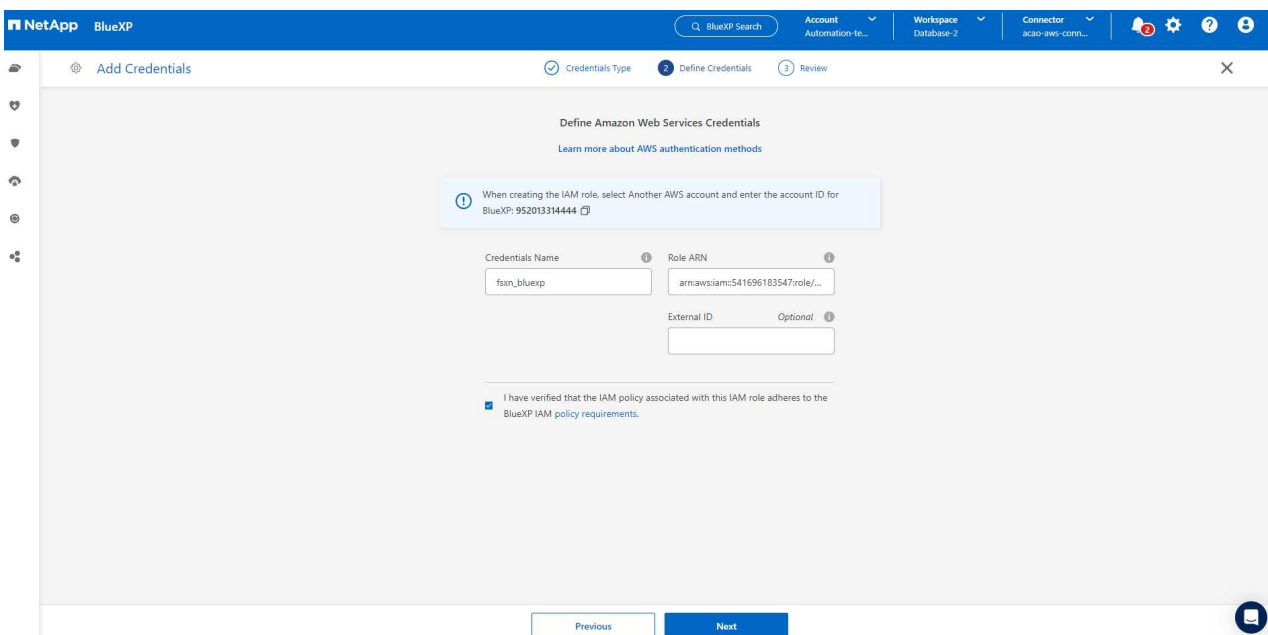
shantanucreds
Type: Assume Role | BlueXP

210811600188 AWS Account ID | nkarthik_kafka_nfs_role_FSxN Assume Role

6. Choose credential location as - **Amazon Web Services - BlueXP**.



7. Define AWS credentials with proper **Role ARN**, which can be retrieved from AWS IAM role created in step one above. BlueXP **account ID**, which is used for creating AWS IAM role in step one.



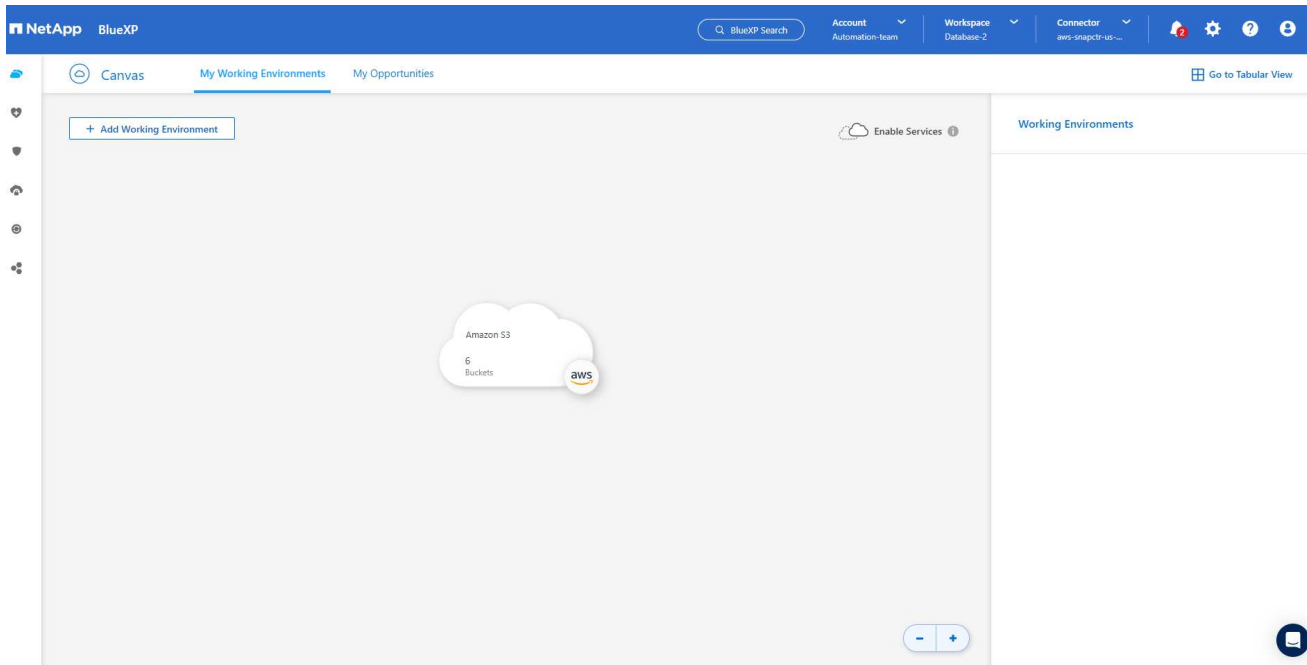
8. Review and **Add**.

image::snapctr_svcs_credential_08-aws.png["Screenshot showing this step in the GUI."]

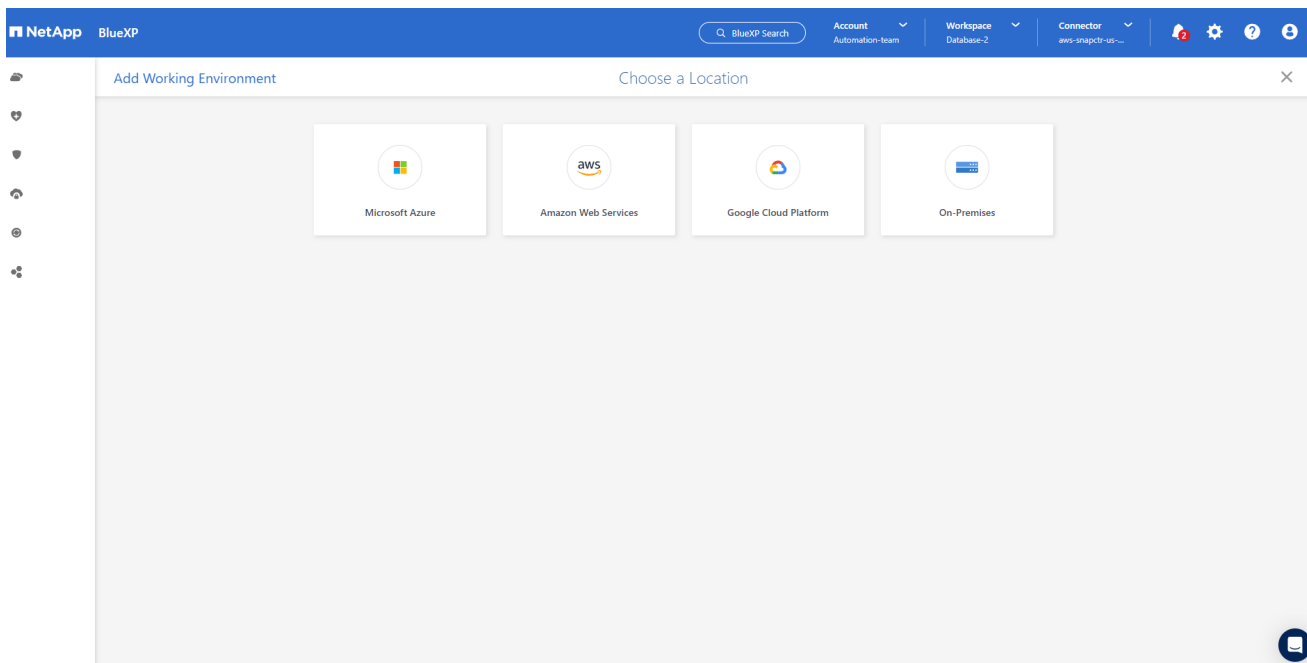
SnapCenter services setup

With the connector deployed and the credential added, SnapCenter services can now be set up with the following procedure:


1. From **My Working Environment** click **Add working Environment** to discover FSx deployed in AWS.




1. Choose **Amazon Web Services** as the location.



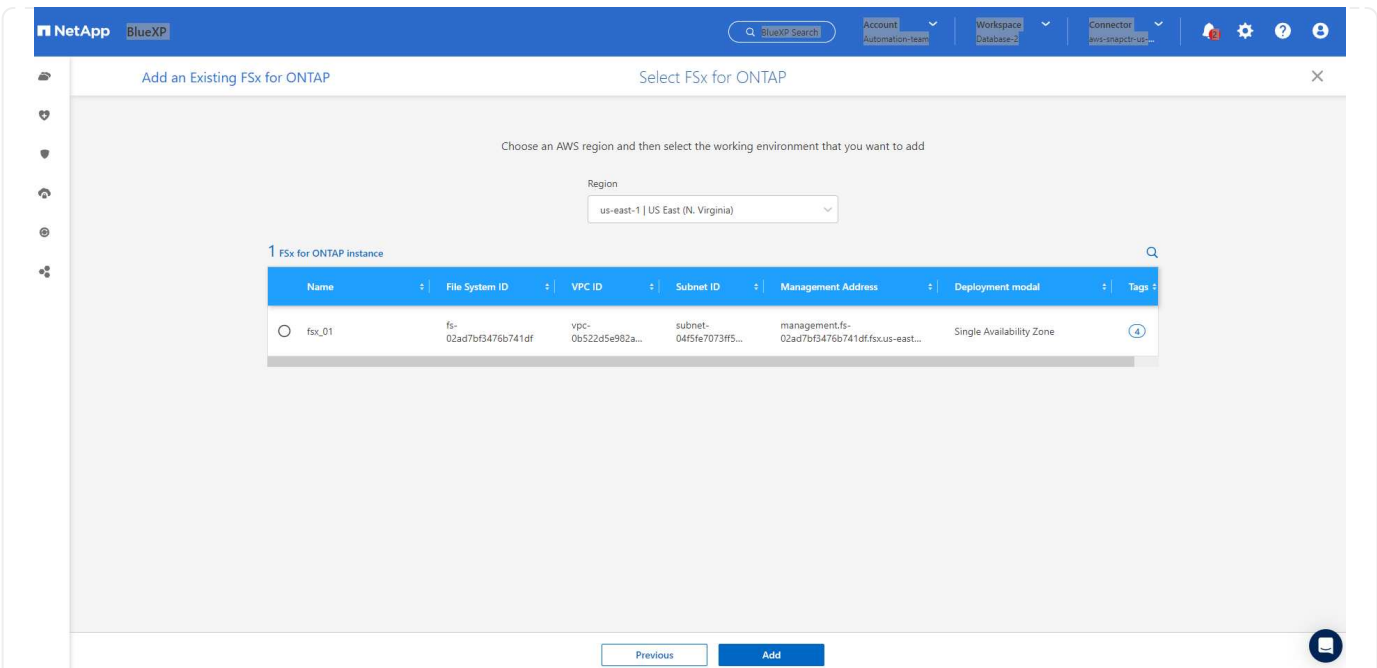
1. Click **Discover Existing** next to **Amazon FSx for ONTAP**.



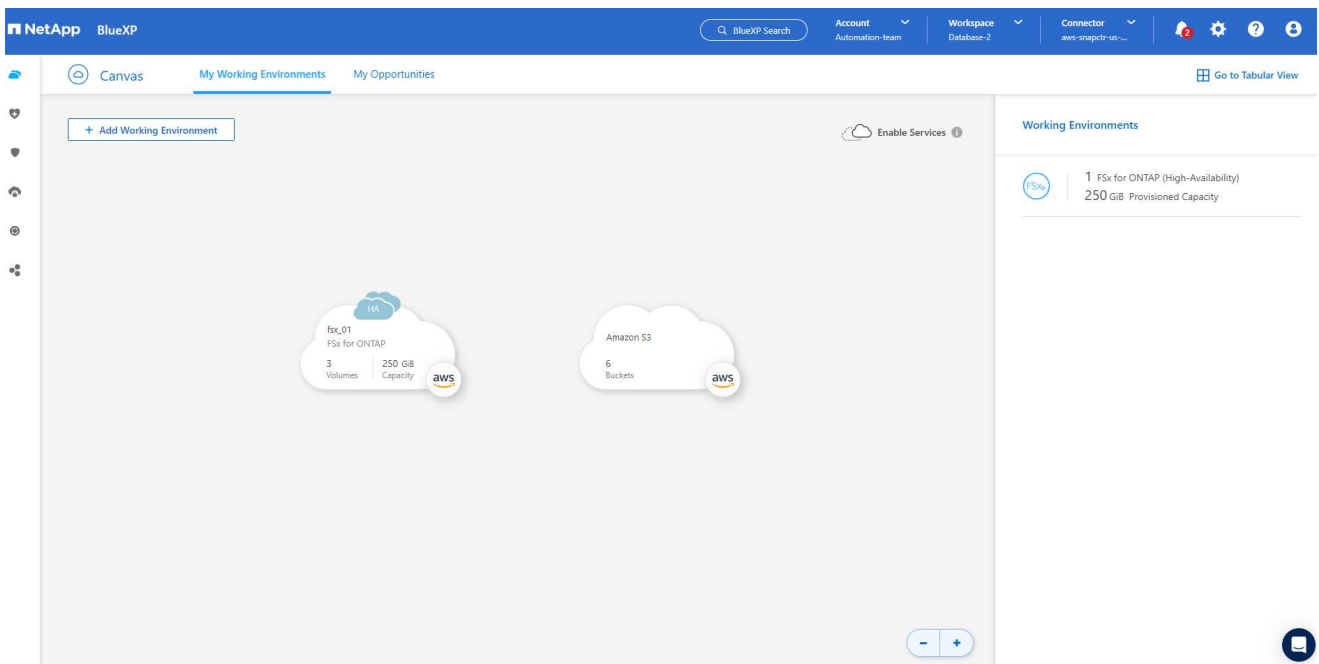
1. Select the **Credentials Name** that you have created in previous section to grant BlueXP with the permissions that it needs to manage FSx for ONTAP. If you have not added credentials, you can add it from the **Settings** menu at the top right corner of the BlueXP console.



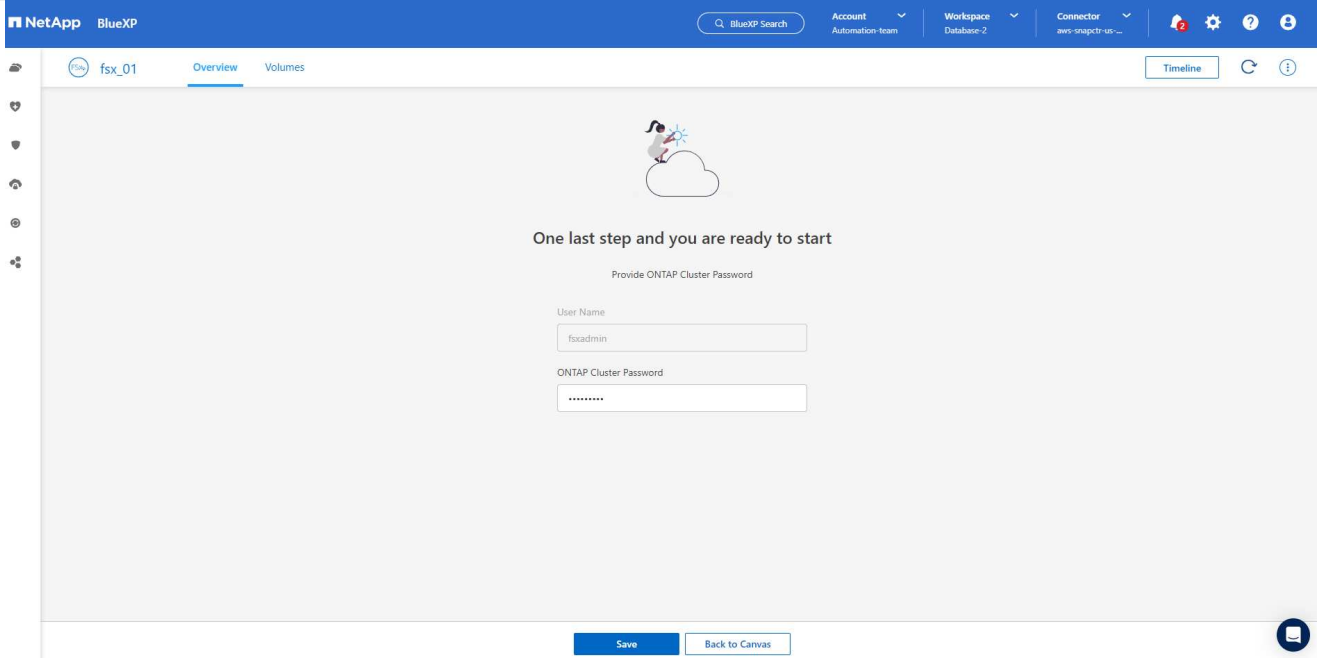
2. Choose the AWS region where Amazon FSx for ONTAP is deployed, select the FSx cluster that is hosting the Oracle database and click Add.



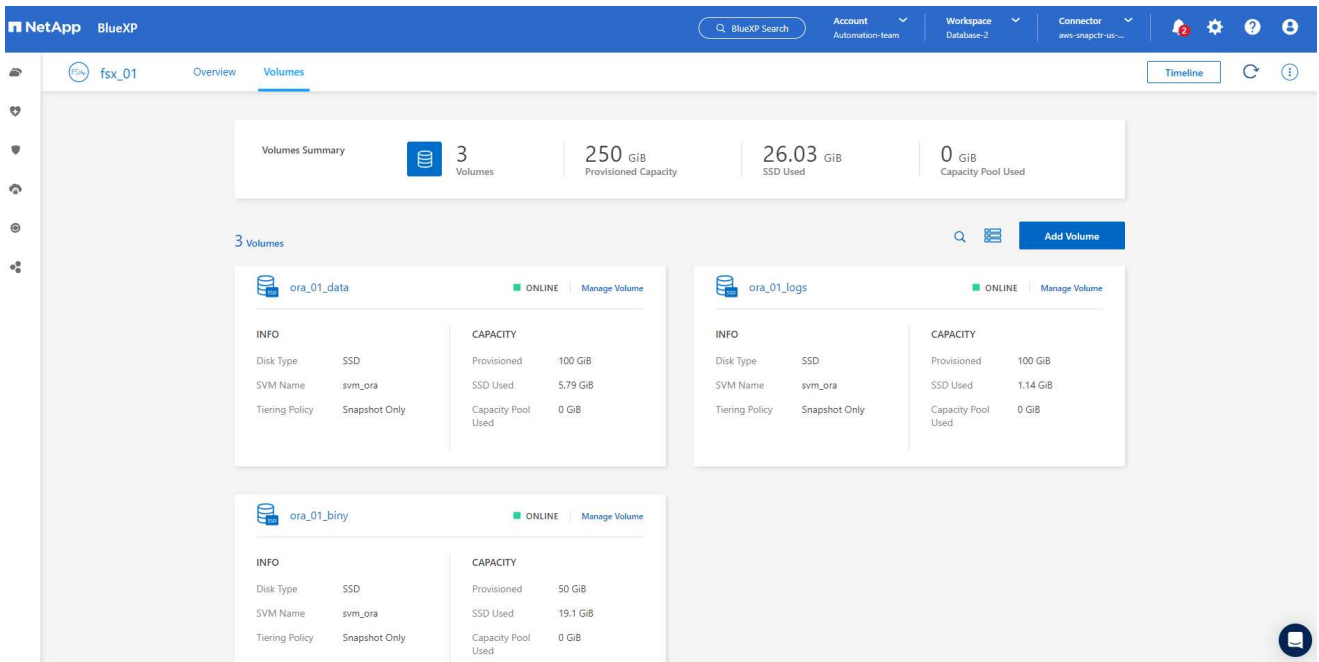
1. The discovered Amazon FSx for ONTAP instance now appears in the working environment.




1. You can log into the FSx cluster with your fsxadmin account credentials.



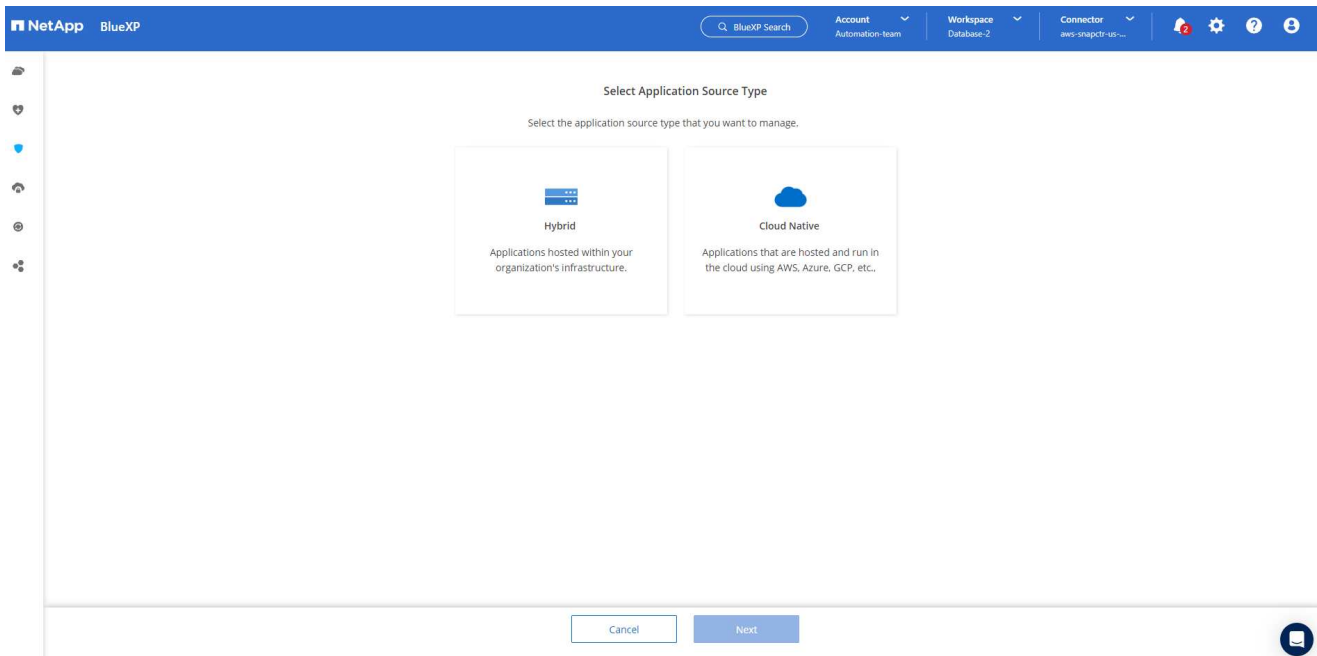
1. After you log into Amazon FSx for ONTAP, review your database storage information (such as database volumes).




1. From the left-hand sidebar of the console, hover your mouse over the protection icon, and then click **Protection > Applications** to open the Applications launch page. Click **Discover Applications**.



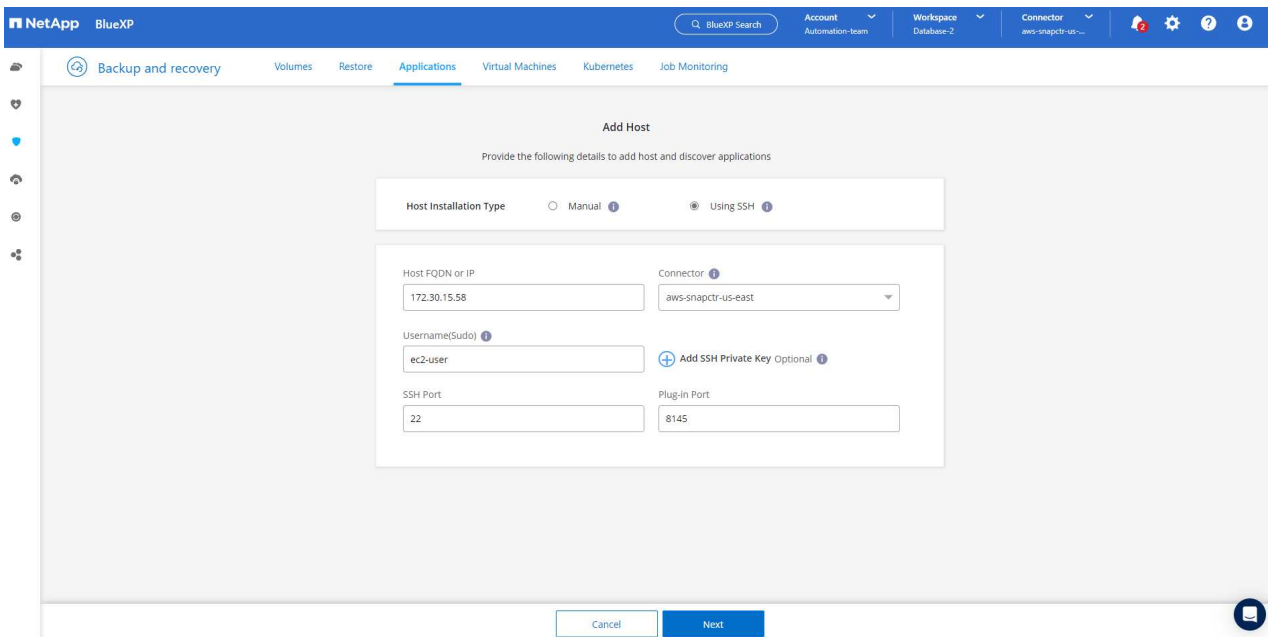
1. Select **Cloud Native** as the application source type.




1. Choose **Oracle** for the application type.



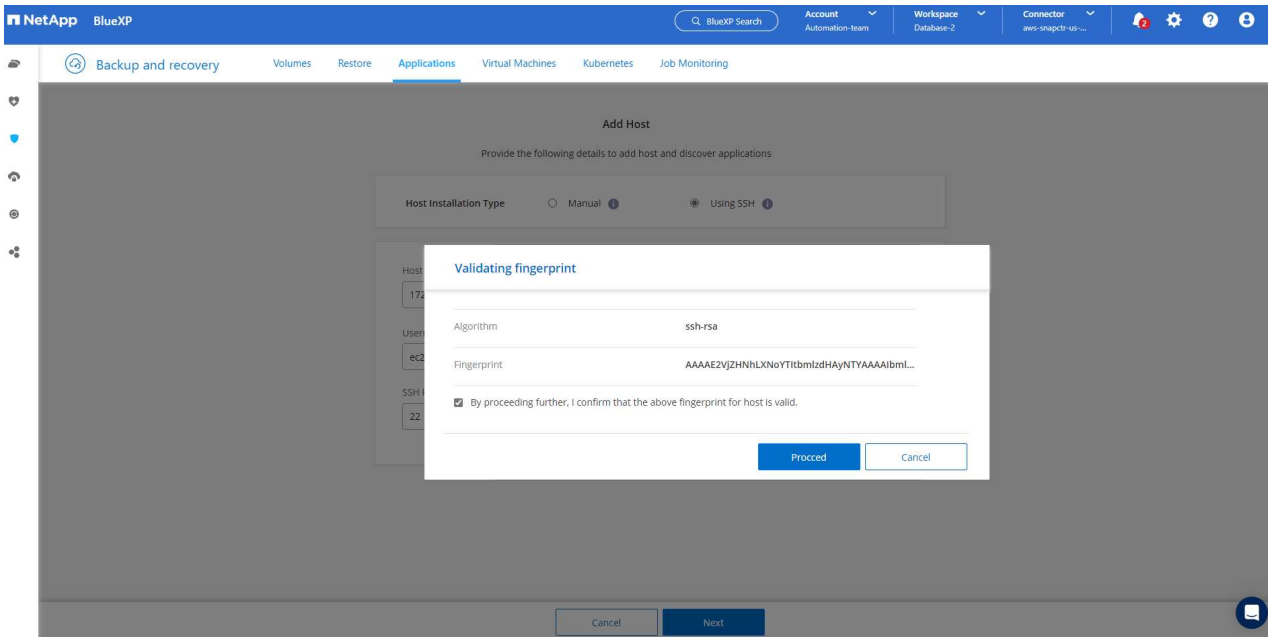
1. Fill in the AWS EC2 Oracle application host details. Choose **Using SSH** as **Host Installation Type** for one step plugin installation and database discovery. Then, click on **Add SSH Private Key**.




2. Paste in your ec2-user SSH key for the database EC2 host and click on **Validate** to proceed.



3. You will be prompted for **Validating fingerprint** to proceed.




4. Click on **Next** to install an Oracle database plugin and discover the Oracle databases on the EC2 host. Discovered databases are added to **Applications**. The database **Protection Status** shows as **Unprotected** when initially discovered.




This completes the initial setup of SnapCenter services for Oracle. The next three sections of this document describe Oracle database backup, restore, and clone operations.

Oracle database backup


1. Click the three dots next to the database **Protection Status**, and then click **Policies** to view the default preloaded database protection policies that can be applied to protect your Oracle databases.



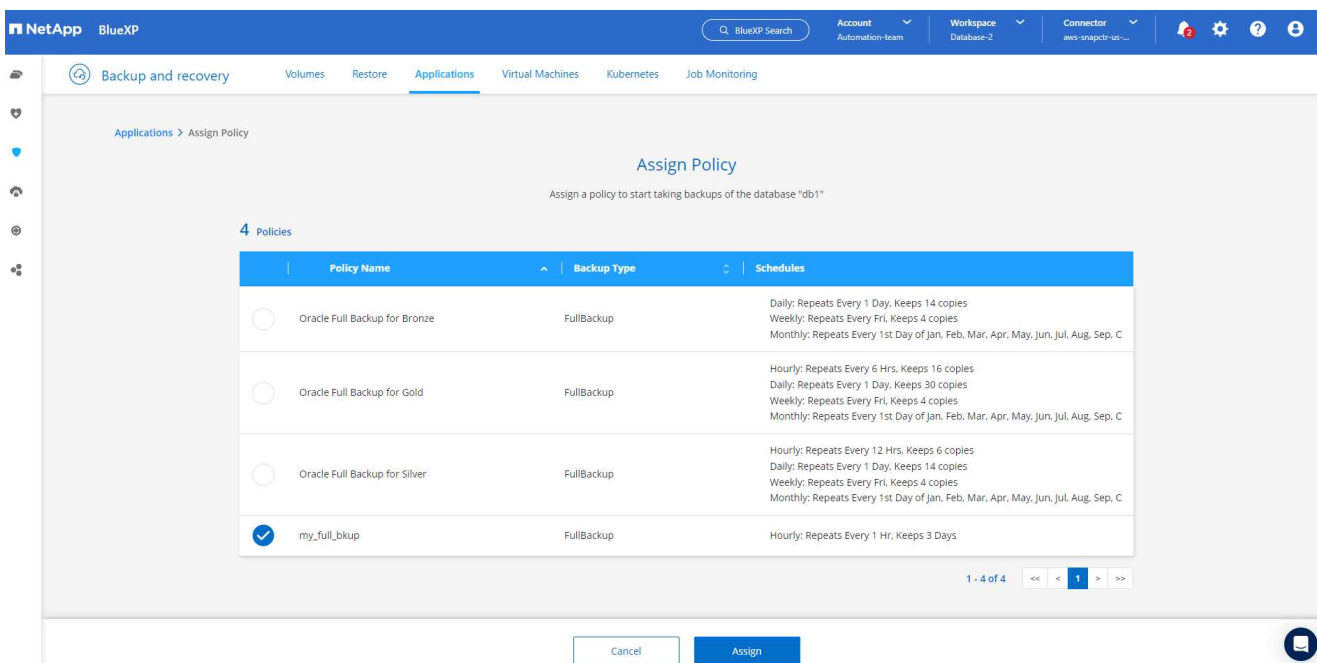
1. You can also create your own policy with a customized backup frequency and backup data-retention window.



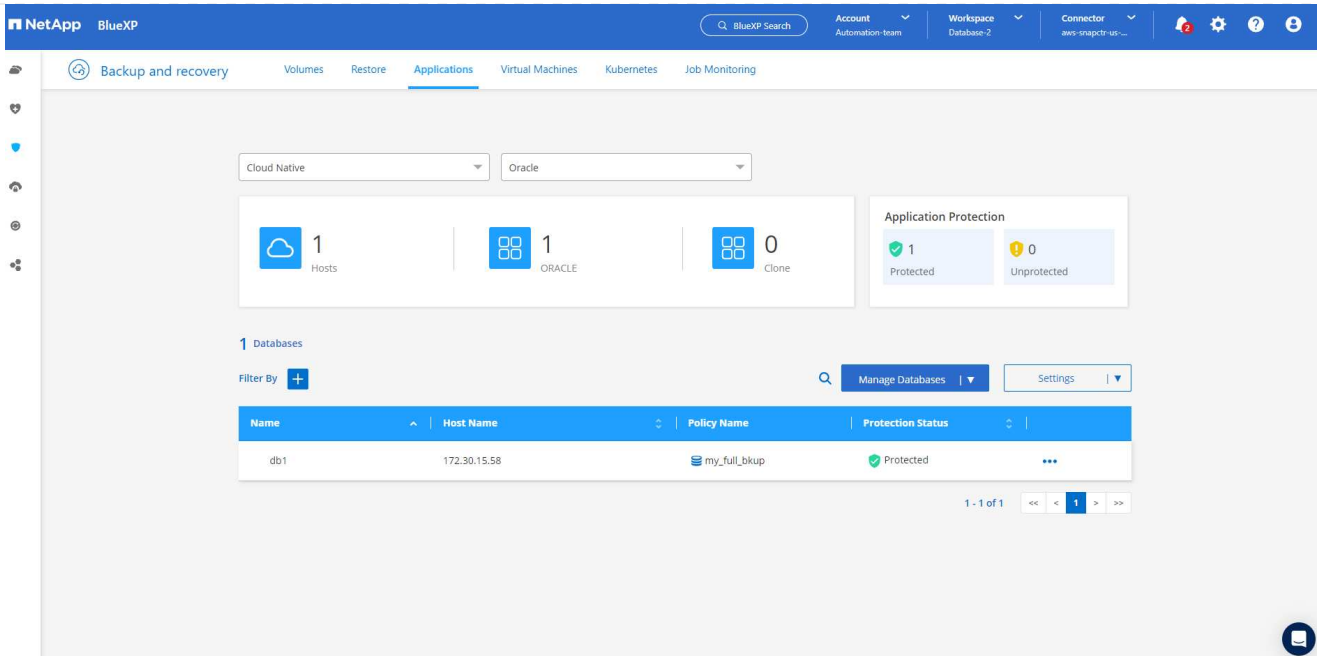
1. When you are happy with the policy configuration, you can then assign your policy of choice to protect the database.



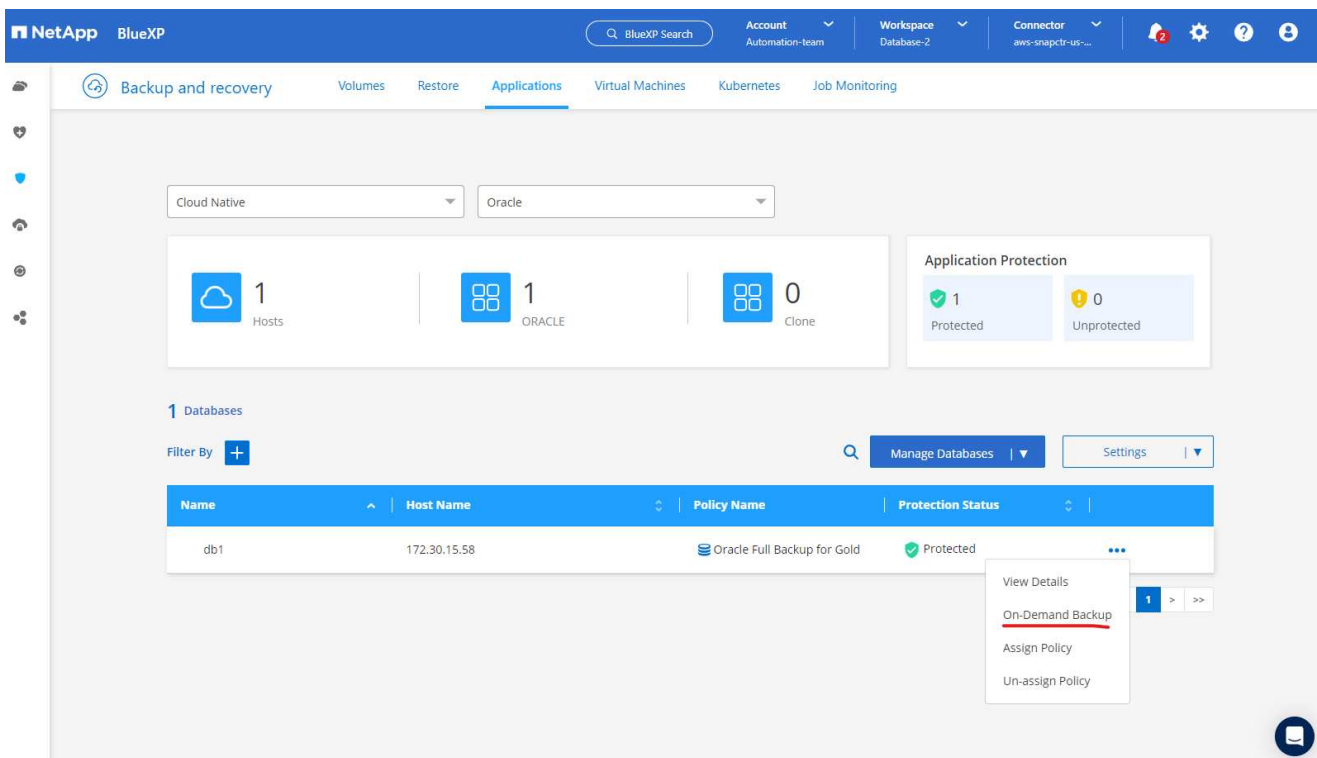
1. Choose the policy to assign to the database.



1. After the policy is applied, the database protection status changed to **Protected** with a green check mark.



1. The database backup runs on a predefined schedule. You can also run a one-off on-demand backup as shown below.



1. The database backups details can be viewed by clicking **View Details** from the menu list. This includes the backup name, backup type, SCN, and backup date. A backup set covers a snapshot for both data volume and log volume. A log volume snapshot takes place right after a database volume snapshot. You can apply a filter if you are looking for a particular backup in a long list.

NetApp BlueXP

Account Automation-team | Workspace Database-2 | Connector aws-snapctr-us...

Backup and recovery | Volumes | Restore | Applications | Virtual Machines | Kubernetes | Job Monitoring

Applications > Database Details

Database Details

db1 Database Name	Protected Protection	Oracle Full Backup for Gold Policy Names	Database Type
172.30.15.58 Host Name	FSx Host Storage	Unreachable Database Version	bKed8yv2T19Bj0V5QyqvA... Agent Id
- Clones	- Parent Database		

8 Backups

Filter By +

Select Timeframe

Backup Name	Backup Type	SCN	Backup Date	
Oracle_Full_Backup_for_Gold_Weekly_db1_2023_03_24_19_12_18_60900_1	Log	2589354	Mar 24, 2023, 3:12:34 pm	Delete
Oracle_Full_Backup_for_Gold_Weekly_db1_2023_03_24_19_11_51_51476_0	Data	2589306	Mar 24, 2023, 3:12:18 pm	...
Oracle_Full_Backup_for_Gold_Hourly_db1_2023_03_24_18_10_31_71953_1	Log	2586621	Mar 24, 2023, 2:10:45 pm	Delete
Oracle_Full_Backup_for_Gold_Hourly_db1_2023_03_24_18_10_03_70535_0	Data	2586557	Mar 24, 2023, 2:10:31 pm	...

Oracle database restore and recovery

1. For a database restore, choose the right backup, either by the SCN or backup time. Click the three dots from the database data backup, and then click **Restore** to initiate database restore and recovery.

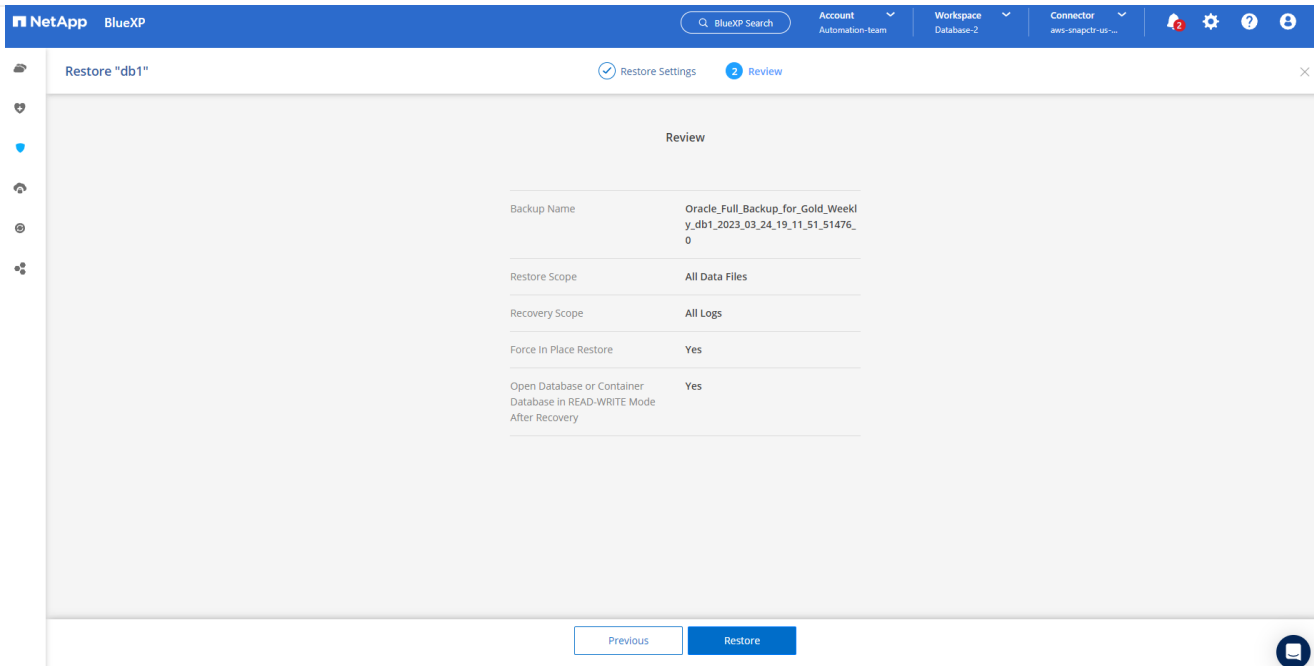
The screenshot shows the 'Database Details' page in NetApp BlueXP. It displays information for a database named 'db1', including its host name, storage, and backup details. Below this, there is a 'Backups' section with a table listing several backup entries. The table has columns for Backup Name, Backup Type, SCN, Backup Date, and actions. The 'Restore' button is highlighted in red on the third row of the table.

Backup Name	Backup Type	SCN	Backup Date	Actions
Oracle_Full_Backup_for_Gold_Hourly_db1_2023_03_24_18_10_31_71953_1	Log	2586621	Mar 24, 2023, 2:10:45 pm	Delete
Oracle_Full_Backup_for_Gold_Hourly_db1_2023_03_24_18_10_03_70535_0	Data	2586557	Mar 24, 2023, 2:10:31 pm	...
Oracle_Full_Backup_for_Gold_Hourly_db1_2023_03_24_15_37_04_98851_1	Log	2580577	Mar 24, 2023, 11:37:11	Restore, Delete
Oracle_Full_Backup_for_Gold_Hourly_db1_2023_03_24_15_36_33_27205_0	Data	2580524	Mar 24, 2023, 11:37:00	Clone

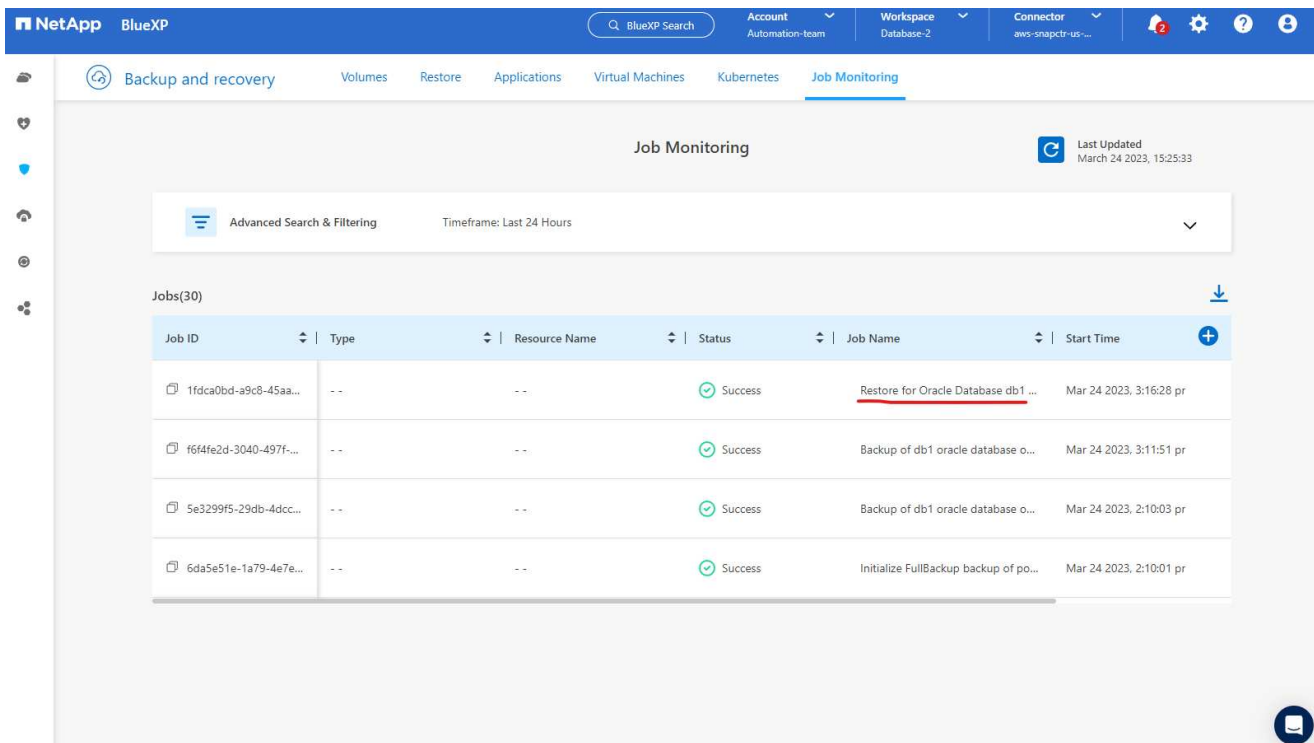
1. Choose your restore setting. If you are sure that nothing has changed in the physical database structure after the backup (such as the addition of a data file or a disk group), you can use the **Force in place restore** option, which is generally faster. Otherwise, do not check this box.

The screenshot shows the 'Restore Settings' page in NetApp BlueXP. It displays various configuration options for restoring the database 'db1'. The 'Restore Scope' section includes options for 'All Data Files', 'Control Files', and 'Force in place restore' (which is checked). The 'Recovery Scope' section includes options for 'All Logs', 'Until System Change Number', 'Date and Time', and 'No Recovery'. There is also a field for 'Archive Log Files Locations' and a checkbox for 'Open the database or the container database in READ-WRITE mode after recovery'.

1. Review and start database restore and recovery.



1. From the **Job Monitoring** tab, you can view the status of the restore job as well as any details while it is running.



NetApp BlueXP Account Automation-team Workspace Database-2 Connector aws-snapctr-us-...

Backup and recovery Volumes Restore Applications Virtual Machines Kubernetes Job Monitoring

Job Monitoring > Job Id: 1fdca0bd-a9c8-45aa-9d7a-05a07cb291f4

Job Details

Job Id: 1fdca0bd-a9c8-45aa-9d7a-05a07cb291f4 Expand All


Sub-Jobs(6)

Job Name	Job ID	Start Time	End Time	Duration
Restore for Oracle Database db1 using backup ...	1fdca0bd-a9c8-45aa-9d...	Mar 24 2023, 3:16:28 pm	Mar 24 2023, 3:23:33 pm	7 Minutes
Post Restore Cleanup	2096a8e4-889d-4b2a-9...	Mar 24 2023, 3:23:18 pm	Mar 24 2023, 3:23:32 pm	14 Seconds
Post Restore	fb7b1171-966f-4228-9e...	Mar 24 2023, 3:20:06 pm	Mar 24 2023, 3:23:19 pm	3 Minutes
Restore	0f4580d0-6598-458b-a7...	Mar 24 2023, 3:17:49 pm	Mar 24 2023, 3:20:07 pm	2 Minutes


Oracle database clone

To clone a database, launch the clone workflow from the same database backup details page.

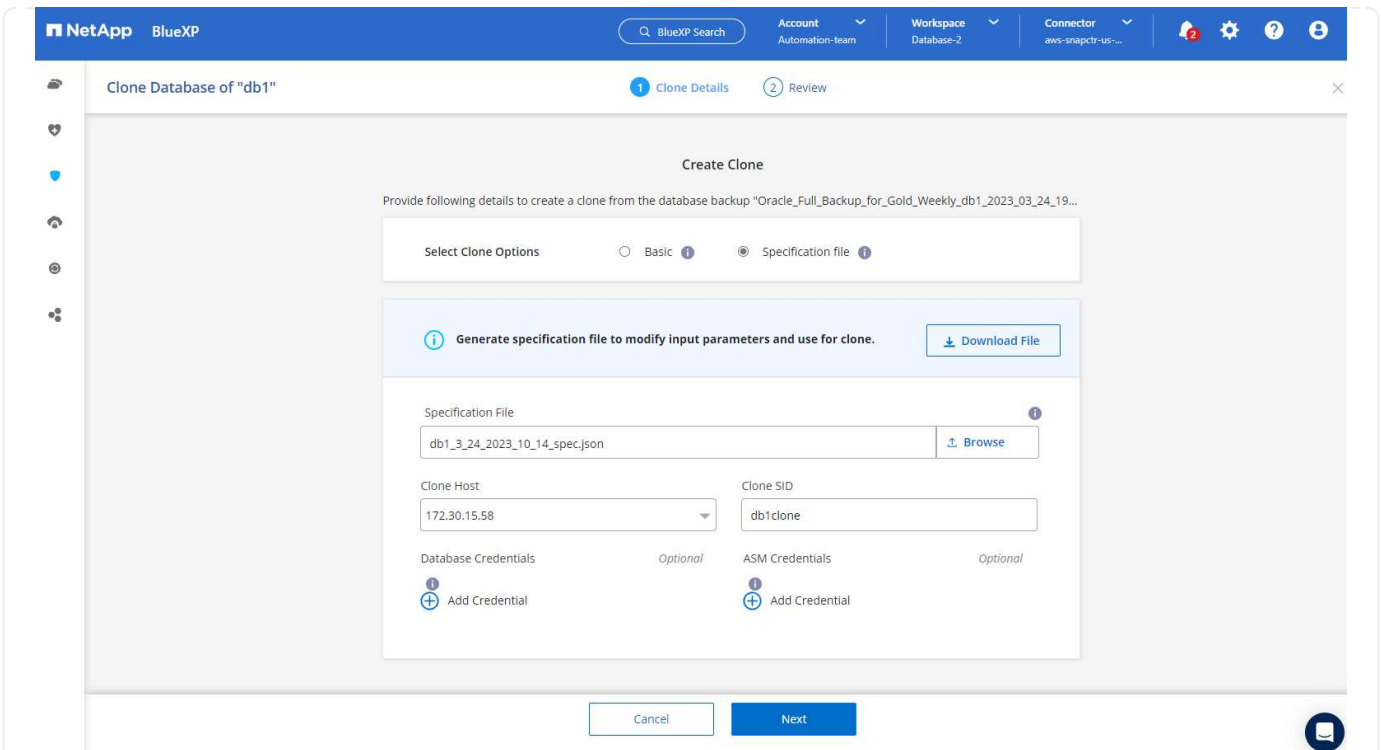
1. Select the right database backup copy, click the three dots to view the menu, and choose the **Clone** option.



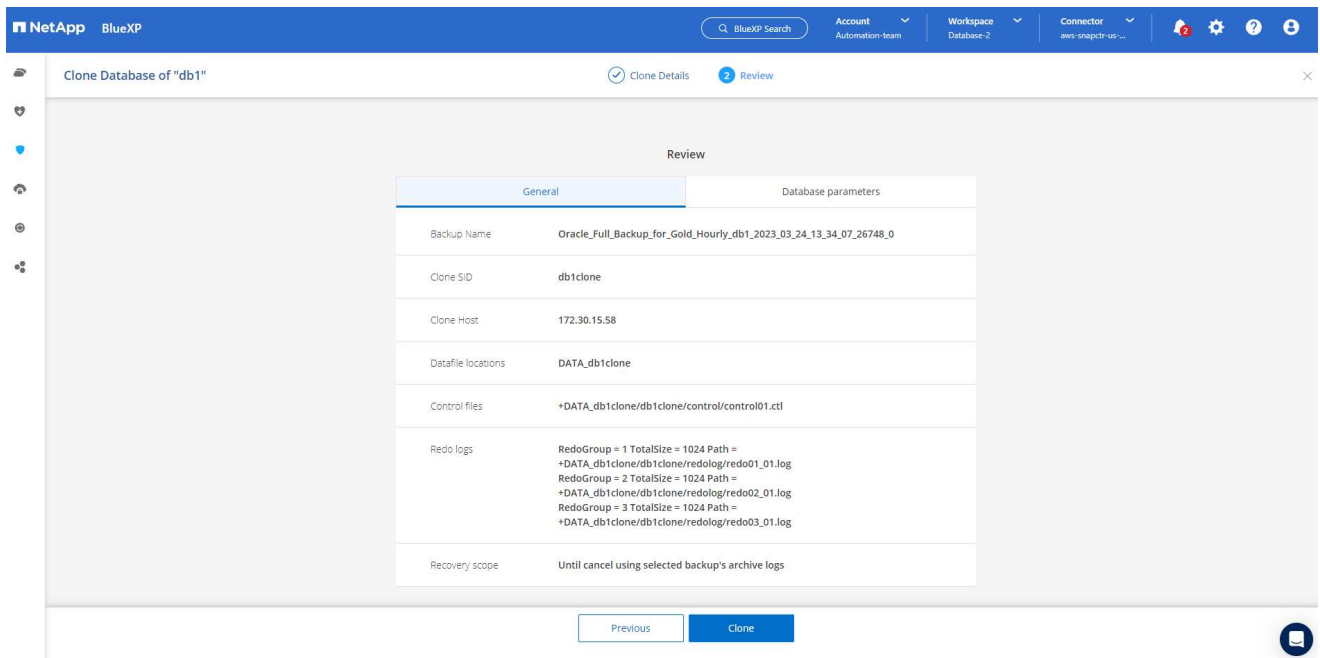
1. Select the **Basic** option if you don't need to change any cloned database parameters.



1. Alternatively, select **Specification file**, which gives you the option of downloading the current init file, making changes, and then uploading it back to the job.



1. Review and launch the job.



1. Monitor the cloning job status from the **Job Monitoring** tab.

The screenshot displays the NetApp BlueXP interface. The top navigation bar includes the NetApp logo, 'BlueXP', a search bar, and user/role information: 'Account Automation-team', 'Workspace Database-2', and 'Connector aws-snapc1r-1b...'. The main menu shows 'Backup and recovery' as the active section, with sub-sections for 'Volumes', 'Restore', 'Applications', 'Virtual Machines', 'Kubernetes', and 'Job Monitoring'. The 'Job Monitoring' page shows a specific job with ID 'cd30abaf-fbe2-4052-a6db-4bf965a8d29b'. The 'Job Details' section shows 'Sub-Jobs(2)' and an 'Expand All' button. A table lists the sub-jobs:

Job Name	Job ID	Start Time	End Time	Duration
Cloning Oracle Database db1 as db1clone on h...	cd30abaf-fbe2-4052-a6...	Mar 24 2023, 1:30:36 pm		--
Running pre scripts	511f52c1-853a-4e6-a4f...	Mar 24 2023, 1:30:41 pm	Mar 24 2023, 1:30:41 pm	0 Second
Validating clone request	f93a6c44-2eb2-4c5e-9f...	Mar 24 2023, 1:30:35 pm	Mar 24 2023, 1:30:42 pm	7 Seconds

1. Validate the cloned database on the EC2 instance host.

```

#
# Multiple entries with the same $ORACLE_SID are not allowed.
#
#
+ASM:/u01/app/oracle/product/19.0.0/grid:N
db1:/u01/app/oracle/product/19.0.0/db1:N
# SnapCenter Plug-in for Oracle Database generated entry (DO NOT REMOVE THIS LINE)
db1clone:/u01/app/oracle/product/19.0.0/db1:N
[oracle@ip-172-30-15-58 ~]$ crsctl stat res -t
-----
Name                Target  State        Server                    State details
-----
Local Resources
-----
ora.DATA.dg
      ONLINE  ONLINE      ip-172-30-15-58          STABLE
ora.DATA_DB1CLONE.dg
      ONLINE  ONLINE      ip-172-30-15-58          STABLE
ora.LISTENER.lsnr
      ONLINE  ONLINE      ip-172-30-15-58          STABLE
ora.LOGS.dg
      ONLINE  ONLINE      ip-172-30-15-58          STABLE
ora.LOGS_SCO_2748138658.dg
      ONLINE  ONLINE      ip-172-30-15-58          STABLE
ora.asm
      ONLINE  ONLINE      ip-172-30-15-58          Started,STABLE
ora.ons
      OFFLINE OFFLINE      ip-172-30-15-58          STABLE
-----
Cluster Resources
-----
ora.cssd
      1        ONLINE  ONLINE      ip-172-30-15-58          STABLE
ora.db1.db
      1        ONLINE  ONLINE      ip-172-30-15-58          Open,HOME=/u01/app/oracle/product/19.0.0/db1,STABLE
ora.db1clone.db
      1        ONLINE  ONLINE      ip-172-30-15-58          Open,HOME=/u01/app/oracle/product/19.0.0/db1,STABLE
ora.diskmon
      1        OFFLINE OFFLINE
      STABLE
ora.driver.afd
      1        ONLINE  ONLINE      ip-172-30-15-58          STABLE
ora.evmd
      1        ONLINE  ONLINE      ip-172-30-15-58          STABLE
-----
[oracle@ip-172-30-15-58 ~]$ █

```

```

[oracle@ip-172-30-15-58 ~]$ export ORACLE_HOME=/u01/app/oracle/product/19.0.0/db1
[oracle@ip-172-30-15-58 ~]$ export ORACLE_SID=db1clone
[oracle@ip-172-30-15-58 ~]$ export PATH=$ORACLE_HOME/bin:$PATH
[oracle@ip-172-30-15-58 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Fri Mar 24 18:32:21 2023
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 from v$databases;

NAME          OPEN_MODE
-----
DB1CLONE     READ WRITE

SQL> █

```

Additional information

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

- Set up and administer BlueXP

<https://docs.netapp.com/us-en/cloud-manager-setup-admin/index.html>

- BlueXP backup and recovery documentation

<https://docs.netapp.com/us-en/cloud-manager-backup-restore/index.html>

- Amazon FSx for NetApp ONTAP

<https://aws.amazon.com/fsx/netapp-ontap/>

- Amazon EC2

https://aws.amazon.com/pm/ec2/?trk=36c6da98-7b20-48fa-8225-4784bced9843&sc_channel=ps&s_kwcid=AL!4422!3!467723097970!e!!g!!aws%20ec2&ef_id=Cj0KCQiA54KfBhCKARIsAJzSrdqwQrghn6I71jiWzSeaT9Uh1-vY-VfhJixF-xnv5rWwn2S7RqZOTQ0aAh7eEALw_wcB:G:s&s_kwcid=AL!4422!3!467723097970!e!!g!!aws%20ec2

Hybrid Cloud Database Solutions with SnapCenter

TR-4908: Hybrid Cloud Database Solutions with SnapCenter Overview

Alan Cao, Felix Melligan, NetApp

This solution provides NetApp field and customers with instructions and guidance for configuring, operating, and migrating databases to a hybrid cloud environment using the NetApp SnapCenter GUI-based tool and the NetApp storage service CVO in public clouds for the following use cases:


- Database dev/test operations in the hybrid cloud
- Database disaster recovery in the hybrid cloud

Today, many enterprise databases still reside in private corporate data centers for performance, security, and/or other reasons. This hybrid cloud database solution enables enterprises to operate their primary databases on site while using a public cloud for dev/test database operations as well as for disaster recovery to reduce licensing and operational costs.

Many enterprise databases, such as Oracle, SQL Server, SAP HANA, and so on, carry high licensing and operational costs. Many customers pay a one-time license fee as well as annual support costs based on the number of compute cores in their database environment, whether the cores are used for development, testing, production, or disaster recovery. Many of those environments might not be fully utilized throughout the application lifecycle.

The solutions provide an option for customers to potentially reduce their licensable cores count by moving their database environments devoted to development, testing, or disaster recovery to the cloud. By using public-cloud scale, redundancy, high availability, and a consumption-based billing model, the cost saving for licensing and operation can be substantial, while not sacrificing any application usability or availability.

Beyond potential database license-cost savings, the NetApp capacity-based CVO license model allows customers to save storage costs on a per-GB basis while empowering them with high level of database manageability that is not available from competing storage services. The following chart shows a storage cost comparison of popular storage services available in the public cloud.



This solution demonstrates that, by using the SnapCenter GUI-based software tool and NetApp SnapMirror technology, hybrid cloud database operations can be easily setup, implemented, and operated.

The following videos demonstrate SnapCenter in action:


- [Backup of an Oracle database across a Hybrid Cloud using SnapCenter](#)
- [SnapCenter- Clone DEV/TEST to AWS Cloud for an Oracle database](#)

Notably, although the illustrations throughout this document show CVO as a target storage instance in the public cloud, the solution is also fully validated for the new release of the FSx ONTAP storage engine for AWS.

To test drive the solution and use cases for yourself, a NetApp Lab-on-Demand SL10680 can be requested at following xref:./databases/ [TL_AWS_004](#) HCoD: [AWS - NW,SnapCenter\(OnPrem\)](#).

Solution Architecture

The following architecture diagram illustrates a typical implementation of enterprise database operation in a hybrid cloud for dev/test and disaster recovery operations.



In normal business operations, synchronized database volumes in the cloud can be cloned and mounted to dev/test database instances for applications development or testing. In the event of a failure, the synchronized database volumes in the cloud can then be activated for disaster recovery.

SnapCenter Requirements

This solution is designed in a hybrid cloud setting to support on-premises production databases that can burst to all of the popular public clouds for dev/test and disaster recovery operations.

This solution supports all databases that are currently supported by SnapCenter, although only Oracle and SQL Server databases are demonstrated here. This solution is validated with virtualized database workloads, although bare-metal workloads are also supported.

We assume that production database servers are hosted on-premises with DB volumes presented to DB hosts from a ONTAP storage cluster. SnapCenter software is installed on-premises for database backup and data replication to the cloud. An Ansible controller is recommended but not required for database deployment automation or OS kernel and DB configuration syncing with a standby DR instance or dev/test instances in the public cloud.

Requirements

Environment	Requirements
On-premises	Any databases and versions supported by SnapCenter
	SnapCenter v4.4 or higher
	Ansible v2.09 or higher
	ONTAP cluster 9.x
	Intercluster LIFs configured
	Connectivity from on-premises to a cloud VPC (VPN, interconnect, and so on)
	Networking ports open - ssh 22 - tcp 8145, 8146, 10000, 11104, 11105
Cloud - AWS	Cloud Manager Connector
	Cloud Volumes ONTAP
	Matching DB OS EC2 instances to On-prem
Cloud - Azure	Cloud Manager Connector
	Cloud Volumes ONTAP
	Matching DB OS Azure Virtual Machines to On-prem
Cloud - GCP	Cloud Manager Connector
	Cloud Volumes ONTAP
	Matching DB OS Google Compute Engine instances to on-premises

Prerequisites configuration

Certain prerequisites must be configured both on-premises and in the cloud before the execution of hybrid cloud database workloads. The following section provides a high-level summary of this process, and the following links provide further information about necessary system configuration.

On premises

- SnapCenter installation and configuration
- On-premises database server storage configuration
- Licensing requirements
- Networking and security
- Automation

Public cloud

- A NetApp Cloud Central login
- Network access from a web browser to several endpoints
- A network location for a connector

- Cloud provider permissions
- Networking for individual services

Important considerations:

1. Where to deploy the Cloud Manager Connector?
2. Cloud Volume ONTAP sizing and architecture
3. Single node or high availability?

The following links provide further details:

[On Premises](#)

[Public Cloud](#)

Prerequisites on-premises

The following tasks must be completed on-premises to prepare the SnapCenter hybrid-cloud database workload environment.

SnapCenter installation and configuration

The NetApp SnapCenter tool is a Windows-based application that typically runs in a Windows domain environment, although workgroup deployment is also possible. It is based on a multitiered architecture that includes a centralized management server (the SnapCenter server) and a SnapCenter plug-in on the database server hosts for database workloads. Here are a few key considerations for hybrid-cloud deployment.


- **Single instance or HA deployment.** HA deployment provides redundancy in the case of a single SnapCenter instance server failure.
- **Name resolution.** DNS must be configured on the SnapCenter server to resolve all database hosts as well as on the storage SVM for forward and reverse lookup. DNS must also be configured on database servers to resolve the SnapCenter server and the storage SVM for both forward and reverse lookup.
- **Role-based access control (RBAC) configuration.** For mixed database workloads, you might want to use RBAC to segregate management responsibility for different DB platform such as an admin for Oracle database or an admin for SQL Server. Necessary permissions must be granted for the DB admin user.
- **Enable policy-based backup strategy.** To enforce backup consistency and reliability.
- **Open necessary network ports on the firewall.** For the on-premises SnapCenter server to communicate with agents installed in the cloud DB host.
- **Ports must be open to allow SnapMirror traffic between on-prem and public cloud.** The SnapCenter server relies on ONTAP SnapMirror to replicate onsite Snapshot backups to cloud CVO storage SVMs.

After careful pre-installation planning and consideration, click this [SnapCenter installation workflow](#) for details of SnapCenter installation and configuration.


On-premises database server storage configuration

Storage performance plays an important role in the overall performance of databases and applications. A well-designed storage layout can not only improve DB performance but also make it easy to manage database backup and recovery. Several factors should be considered when defining your storage layout, including the size of the database, the rate of expected data change for the database, and the frequency with which you perform backups.

Directly attaching storage LUNs to the guest VM by either NFS or iSCSI for virtualized database workloads generally provides better performance than storage allocated via VMDK. NetApp recommends the storage layout for a large SQL Server database on LUNs depicted in the following figure.



The following figure shows the NetApp recommended storage layout for small or medium SQL Server database on LUNs.



The Log directory is dedicated to SnapCenter to perform transaction log rollup for database recovery. For an extra large database, multiple LUNs can be allocated to a volume for better performance.

For Oracle database workloads, SnapCenter supports database environments backed by ONTAP storage that are mounted to the host as either physical or virtual devices. You can host the entire database on a single or multiple storage devices based on the criticality of the environment. Typically, customers isolate data files on dedicated storage from all other files such as control files, redo files, and archive log files. This helps administrators to quickly restore (ONTAP single-file SnapRestore) or clone a large critical database (petabyte scale) using Snapshot technology within few seconds to minutes.



For mission critical workloads that are sensitive to latency, a dedicated storage volume should be deployed to different types of Oracle files to achieve the best latency possible. For a large database, multiple LUNs (NetApp recommends up to eight) per volume should be allocated to data files.



For smaller Oracle databases, SnapCenter supports shared storage layouts in which you can host multiple databases or part of a database on the same storage volume or LUN. As an example of this layout, you can host data files for all the databases on a +DATA ASM disk group or a volume group. The remainder of the files (redo, archive log, and control files) can be hosted on another dedicated disk group or volume group (LVM). Such a deployment scenario is illustrated below.



To facilitate the relocation of Oracle databases, the Oracle binary should be installed on a separate LUN that is included in the regular backup policy. This ensures that in the case of database relocation to a new server host, the Oracle stack can be started for recovery without any potential issues due to an out-of-sync Oracle binary.

Licensing requirements

SnapCenter is licensed software from NetApp. It is generally included in an on-premises ONTAP license. However, for hybrid cloud deployment, a cloud license for SnapCenter is also required to add CVO to SnapCenter as a target data replication destination. Please review following links for SnapCenter standard capacity-based license for details:

[SnapCenter standard capacity-based licenses](#)

Networking and security

In a hybrid database operation that requires an on-premises production database that is burstable to cloud for dev/test and disaster recovery, networking and security is important factor to consider when setting up the environment and connecting to the public cloud from an on-premises data center.

Public clouds typically use a virtual private cloud (VPC) to isolate different users within a public-cloud platform. Within an individual VPC, security is controlled using measures such as security groups that are configurable based on user needs for the lockdown of a VPC.

The connectivity from the on-premises data center to the VPC can be secured through a VPN tunnel. On the VPN gateway, security can be hardened using NAT and firewall rules that block attempts to establish network

connections from hosts on the internet to hosts inside the corporate data center.

For networking and security considerations, review the relevant inbound and outbound CVO rules for your public cloud of choice:

- [Security group rules for CVO - AWS](#)
- [Security group rules for CVO - Azure](#)
- [Firewall rules for CVO - GCP](#)

Using Ansible automation to sync DB instances between on-premises and the cloud - optional

To simplify management of a hybrid-cloud database environment, NetApp highly recommends but does not require that you deploy an Ansible controller to automate some management tasks, such as keeping compute instances on-premises and in the cloud in sync. This is particularly important because an out-of-sync compute instance in the cloud might render the recovered database in the cloud error prone because of missing kernel packages and other issues.

The automation capability of an Ansible controller can also be used to augment SnapCenter for certain tasks, such as breaking up the SnapMirror instance to activate the DR data copy for production.

Follow these instructions to set up your Ansible control node for RedHat or CentOS machines:
include::_include/automation_rhel_centos_setup.adoc[]

Follow these instructions to set up your Ansible control node for Ubuntu or Debian machines:
include::_include/automation_ubuntu_debian_setup.adoc[]

Prerequisites for the public cloud

Before we install the Cloud Manager connector and Cloud Volumes ONTAP and configure SnapMirror, we must perform some preparation for our cloud environment. This page describes the work that needs to be done as well as the considerations when deploying Cloud Volumes ONTAP.

Cloud Manager and Cloud Volumes ONTAP deployment prerequisites checklist

- A NetApp Cloud Central login
- Network access from a web browser to several endpoints
- A network location for a Connector
- Cloud provider permissions
- Networking for individual services

For more information about what you need to get started, visit our [cloud documentation](#).

Considerations

1. What is a Cloud Manager connector?

In most cases, a Cloud Central account admin must deploy a connector in your cloud or on-premises network. The connector enables Cloud Manager to manage resources and processes within your public cloud environment.

For more information about Connectors, visit our [cloud documentation](#).

2. Cloud Volumes ONTAP sizing and architecture

When deploying Cloud Volumes ONTAP, you are given the choice of either a predefined package or the creation of your own configuration. Although many of these values can be changed later on nondisruptively, there are some key decisions that need to be made before deployment based on the workloads to be deployed in the cloud.

Each cloud provider has different options for deployment and almost every workload has its own unique properties. NetApp has a [CVO sizing tool](#) that can help size deployments correctly based on capacity and performance, but it has been built around some basic concepts which are worth considering:

- Capacity required
- Network capability of the cloud virtual machine
- Performance characteristics of cloud storage

The key is to plan for a configuration that not only satisfies the current capacity and performance requirements, but also looks at future growth. This is generally known as capacity headroom and performance headroom.

If you would like further information, read the documentation about planning correctly for [AWS](#), [Azure](#), and [GCP](#).

3. Single node or high availability?

In all clouds, there is the option to deploy CVO in either a single node or in a clustered high availability pair with two nodes. Depending on the use case, you might wish to deploy a single node to save costs or an HA pair to provide further availability and redundancy.

For a DR use case or spinning up temporary storage for development and testing, single nodes are common since the impact of a sudden zonal or infrastructure outage is lower. However, for any production use case, when the data is in only a single location, or when the dataset must have more redundancy and availability, high availability is recommended.

For further information about the architecture of each cloud's version of high availability, visit the documentation for [AWS](#), [Azure](#) and [GCP](#).

Getting started overview

This section provides a summary of the tasks that must be completed to meet the prerequisite requirements as outlined in previous section. The following section provide a high level tasks list for both on-premises and public cloud operations. The detailed processes and procedures can be accessed by clicking on the relevant links.

On-premises

- Setup database admin user in SnapCenter
- SnapCenter plugin installation prerequisites
- SnapCenter host plugin installation
- DB resource discovery
- Setup storage cluster peering and DB volume replication

- Add CVO database storage SVM to SnapCenter
- Setup database backup policy in SnapCenter
- Implement backup policy to protect database
- Validate backup

AWS public cloud

- Pre-flight check
- Steps to deploy Cloud Manager and Cloud Volumes ONTAP in AWS
- Deploy EC2 compute instance for database workload

Click the following links for details:

[On Premises, Public Cloud - AWS](#)

Getting started on premises

The NetApp SnapCenter tool uses role based access control (RBAC) to manage user resources access and permission grants, and SnapCenter installation creates prepopulated roles. You can also create custom roles based on your needs or applications.

On Premises

1. Setup database admin user in SnapCenter

It makes sense to have a dedicated admin user ID for each database platform supported by SnapCenter for database backup, restoration, and/or disaster recovery. You can also use a single ID to manage all databases. In our test cases and demonstration, we created a dedicated admin user for both Oracle and SQL Server, respectively.

Certain SnapCenter resources can only be provisioned with the SnapCenterAdmin role. Resources can then be assigned to other user IDs for access.

In a pre-installed and configured on-premises SnapCenter environment, the following tasks might have already have been completed. If not, the following steps create a database admin user:

1. Add the admin user to Windows Active Directory.
2. Log into SnapCenter using an ID granted with the SnapCenterAdmin role.
3. Navigate to the Access tab under Settings and Users, and click Add to add a new user. The new user ID is linked to the admin user created in Windows Active Directory in step 1. . Assign the proper role to the user as needed. Assign resources to the admin user as applicable.

The screenshot shows the NetApp SnapCenter interface. The top navigation bar includes 'Global Settings', 'Policies', 'Users and Access', 'Roles', 'Credential', and 'Software'. The 'Users and Access' tab is active. Below the navigation bar is a search bar labeled 'Search by Name' and a dropdown menu for 'Type' set to 'All'. A table lists three users:

Name	Type	Roles	Domain
administrator	User	SnapCenterAdmin	demo
oradba	User	App Backup and Clone Admin	demo
sqlclba	User	App Backup and Clone Admin	demo

2. SnapCenter plugin installation prerequisites

SnapCenter performs backup, restore, clone, and other functions by using a plugin agent running on the DB hosts. It connects to the database host and database via credentials configured under the Setting and Credentials tab for plugin installation and other management functions. There are specific privilege requirements based on the target host type, such as Linux or Windows, as well as the type of database.

DB hosts credentials must be configured before SnapCenter plugin installation. Generally, you want to use an administrator user accounts on the DB host as your host connection credentials for plugin installation. You can also grant the same user ID for database access using OS-based authentication. On the other hand, you can also employ database authentication with different database user IDs for DB management access. If you decide to use OS-based authentication, the OS admin user ID must be granted DB access. For Windows domain-based SQL Server installation, a domain admin account can be used to manage all SQL Servers within the domain.

Windows host for SQL server:

1. If you are using Windows credentials for authentication, you must set up your credential before installing plugins.
2. If you are using a SQL Server instance for authentication, you must add the credentials after installing plugins.
3. If you have enabled SQL authentication while setting up the credentials, the discovered instance or database is shown with a red lock icon. If the lock icon appears, you must specify the instance or database credentials to successfully add the instance or database to a resource group.
4. You must assign the credential to a RBAC user without sysadmin access when the following conditions are met:
 - The credential is assigned to a SQL instance.
 - The SQL instance or host is assigned to an RBAC user.
 - The RBAC DB admin user must have both the resource group and backup privileges.

Unix host for Oracle:

1. You must have enabled the password-based SSH connection for the root or non-root user by editing `sshd.conf` and restarting the `sshd` service. Password-based SSH authentication on AWS instance is turned off by default.
2. Configure the sudo privileges for the non-root user to install and start the plugin process. After installing the plugin, the processes run as an effective root user.
3. Create credentials with the Linux authentication mode for the install user.

4. You must install Java 1.8.x (64-bit) on your Linux host.
5. Installation of the Oracle database plugin also installs the SnapCenter plugin for Unix.

3. SnapCenter host plugin installation

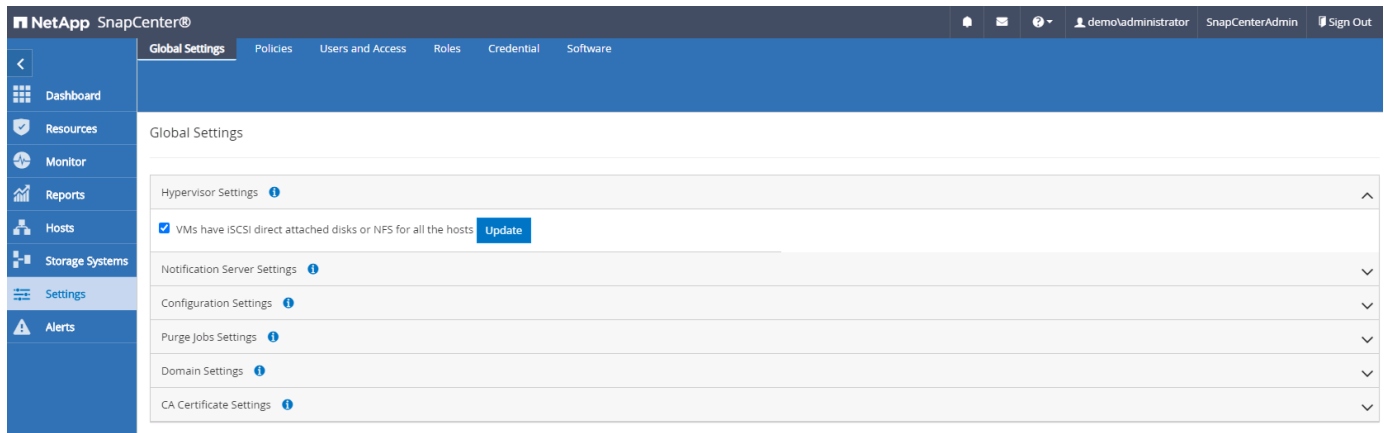


Before attempting to install SnapCenter plugins on cloud DB server instances, make sure that all configuration steps have been completed as listed in the relevant cloud section for compute instance deployment.

The following steps illustrate how a database host is added to SnapCenter while a SnapCenter plugin is installed on the host. The procedure applies to adding both on-premises hosts and cloud hosts. The following demonstration adds a Windows or a Linux host residing in AWS.


Configure SnapCenter VMware global settings

Navigate to Settings > Global Settings. Select "VMs have iSCSI direct attached disks or NFS for all the hosts" under Hypervisor Settings and click Update.




Add Windows host and installation of plugin on the host

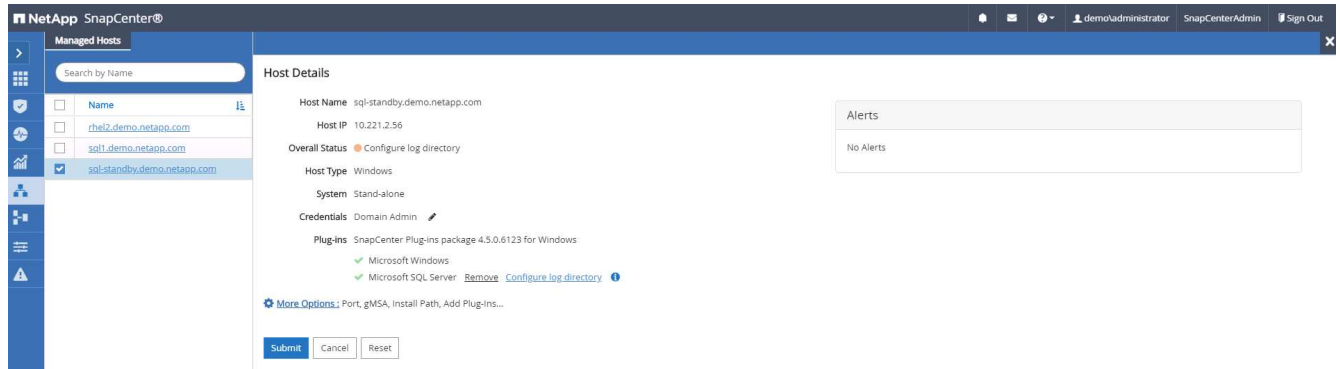
1. Log into SnapCenter with a user ID with SnapCenterAdmin privileges.
2. Click the Hosts tab from the left-hand menu, and then click Add to open the Add Host workflow.
3. Choose Windows for Host Type; the Host Name can be either a host name or an IP address. The host name must be resolved to the correct host IP address from the SnapCenter host. Choose the host credentials created in step 2. Choose Microsoft Windows and Microsoft SQL Server as the plugin packages to be installed.



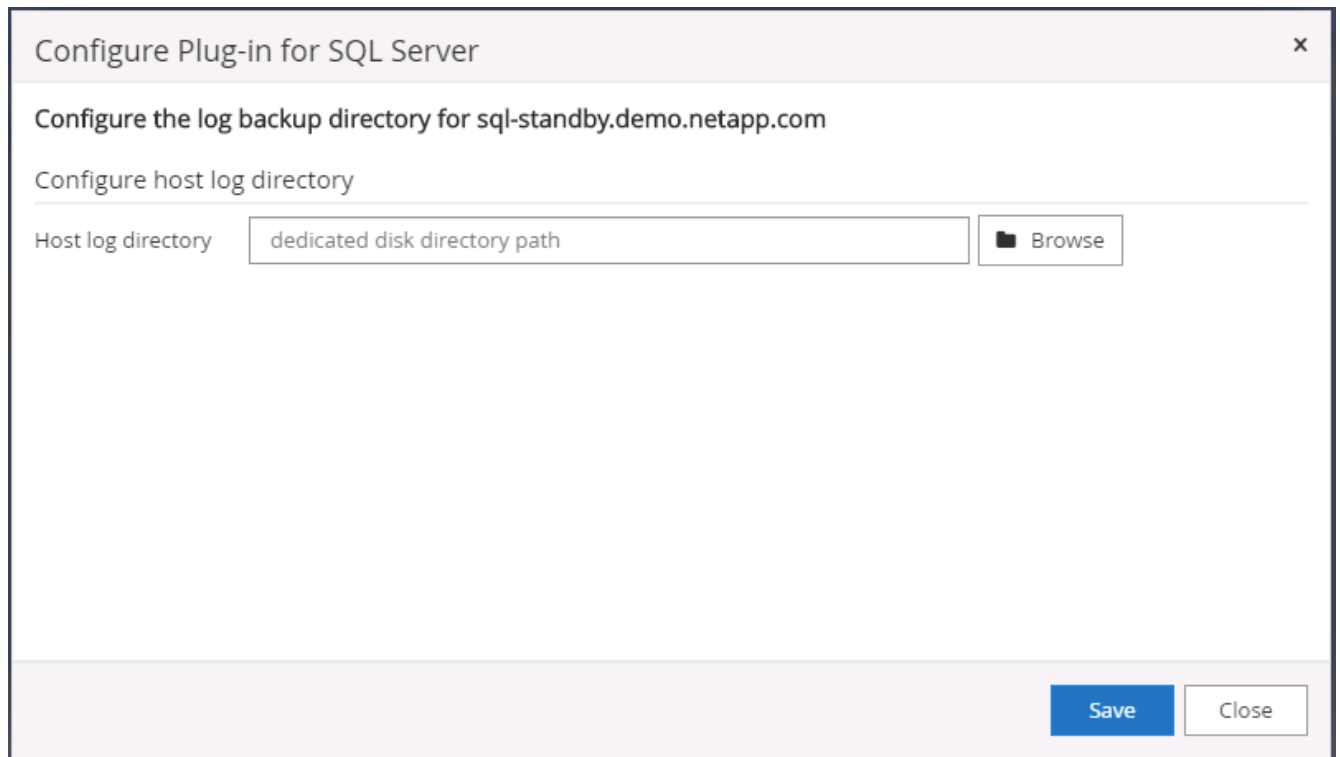
4. After the plugin is installed on a Windows host, its Overall Status is shown as "Configure log directory."




5. Click the Host Name to open the SQL Server log directory configuration.



6. Click "Configure log directory" to open "Configure Plug-in for SQL Server."



7. Click Browse to discover NetApp storage so that a log directory can be set; SnapCenter uses this log directory to roll up the SQL server transaction log files. Then click Save.





For NetApp storage provisioned to a DB host to be discovered, the storage (on-prem or CVO) must be added to SnapCenter, as illustrated in step 6 for CVO as an example.

8. After the log directory is configured, the Windows host plugin Overall Status is changed to Running.



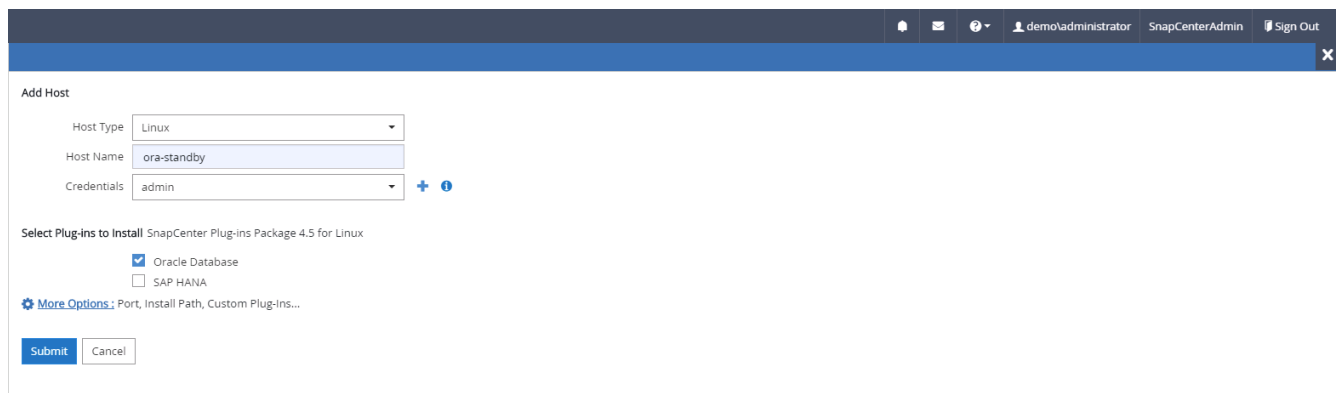
9. To assign the host to the database management user ID, navigate to the Access tab under Settings and Users, click the database management user ID (in our case the sqldba that the host needs to be assigned to), and click Save to complete host resource assignment.





Add Unix host and installation of plugin on the host

1. Log into SnapCenter with a user ID with SnapCenterAdmin privileges.
2. Click the Hosts tab from left-hand menu, and click Add to open the Add Host workflow.
3. Choose Linux as the Host Type. The Host Name can be either the host name or an IP address. However, the host name must be resolved to correct host IP address from SnapCenter host. Choose host credentials created in step 2. The host credentials require sudo privileges. Check Oracle Database as the plug-in to be installed, which installs both Oracle and Linux host plugins.



4. Click More Options and select "Skip preinstall checks." You are prompted to confirm the skipping of the preinstall check. Click Yes and then Save.

More Options

Port:

Installation Path:

Skip preinstall checks

Add all hosts in the oracle RAC

Custom Plug-ins

Choose a File

5. Click Submit to start the plugin installation. You are prompted to Confirm Fingerprint as shown below.

Confirm Fingerprint



Authenticity of the host cannot be determined

Host name	Fingerprint	Valid
ora-standby.demo.netapp.com	ssh-rsa 3072 5C:02:EF:6B:63:54:59:10:84:DF:4D:6B:AB:FB:61:67	

6. SnapCenter performs host validation and registration, and then the plugin is installed on the Linux host. The status is changed from Installing Plugin to Running.

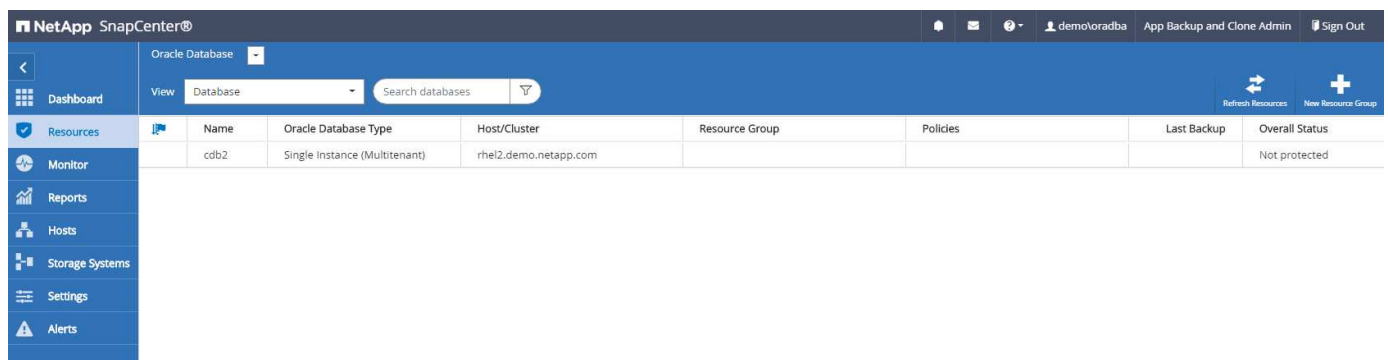
Name	Type	System	Plug-in	Version	Overall Status
ora-standby.demo.netapp.com	Linux	Stand-alone	UNIX, Oracle Database	4.5	Running
rhel2.demo.netapp.com	Linux	Stand-alone	UNIX, Oracle Database	4.5	Running
sql1.demo.netapp.com	Windows	Stand-alone	Microsoft Windows Server, Microsoft SQL Server	4.5	Running
sql-standby.demo.netapp.com	Windows	Stand-alone	Microsoft Windows Server, Microsoft SQL Server	4.5	Running

7. Assign the newly added host to the proper database management user ID (in our case, oradba).



4. Database resource discovery

With successful plugin installation, the database resources on the host can be immediately discovered. Click the Resources tab in the left-hand menu. Depending on the type of database platform, a number of views are available, such as the database, resources group, and so on. You might need to click the Refresh Resources tab if the resources on the host are not discovered and displayed.



When the database is initially discovered, the Overall Status is shown as "Not protected." The previous

screenshot shows an Oracle database not protected yet by a backup policy.

When a backup configuration or policy is set up and a backup has been executed, the Overall Status for the database shows the backup status as "Backup succeeded" and the timestamp of the last backup. The following screenshot shows the backup status of a SQL Server user database.

Name	Instance	Host	Last Backup	Overall Status	Type
master	sql1	sql1.demo.netapp.com		Not available for backup	System database
model	sql1	sql1.demo.netapp.com		Not available for backup	System database
msdb	sql1	sql1.demo.netapp.com		Not available for backup	System database
tempdb	sql1	sql1.demo.netapp.com		Not available for backup	System database
tpcc	sql1	sql1.demo.netapp.com	09/14/2021 2:35:07 PM	Backup succeeded	User database

If database access credentials are not properly set up, a red lock button indicates that the database is not accessible. For example, if Windows credentials do not have sysadmin access to a database instance, then database credentials must be reconfigured to unlock the red lock.

Name	Host	Resource Groups	Policies	State	Type
sql-standby	sql-standby.demo.netapp.com			Running	Standalone ()
sql1	sql1.demo.netapp.com			Running	Standalone (15.0.2000)

The screenshot shows the 'Instance - Credentials' configuration page for the 'sql-standby' instance. A red lock icon is present in the top left corner. A message at the top states: "The Microsoft SQL server or Windows credentials are necessary to unlock the selected instance. Click Refresh Resources to run a discovery with the associated Auth." Below this, the following fields are visible:

- Name: sql-standby
- Resource Group: None
- Policy: None
- Selectable: Not available for backup. DB is not on NetApp storage, auto-close is enabled or in recovery mode.

After the appropriate credentials are configured either at the Windows level or the database level, the red lock disappears and SQL Server Type information is gathered and reviewed.

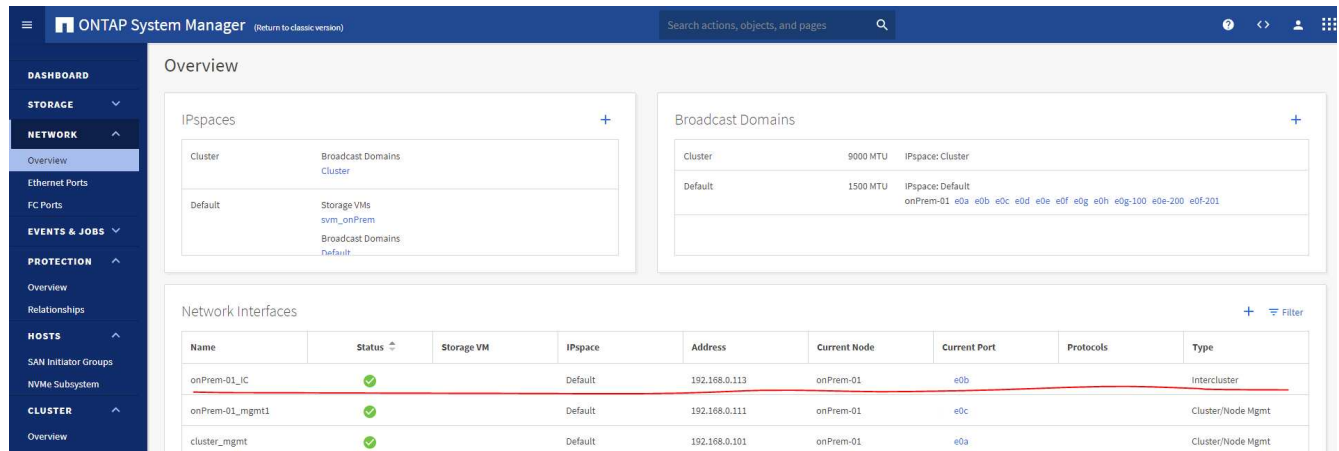
Name	Host	Resource Groups	Policies	State	Type
sql1	sql1.demo.netapp.com			Running	Standalone (15.0.2000)
sql-standby	sql-standby.demo.netapp.com			Running	Standalone (15.0.2000)

5. Setup storage cluster peering and DB volumes replication

To protect your on-premises database data using a public cloud as the target destination, on-premises ONTAP cluster database volumes are replicated to the cloud CVO using NetApp SnapMirror technology. The replicated target volumes can then be cloned for DEV/OPS or disaster recovery. The following high-level steps enable you to set up cluster peering and DB volumes replication.

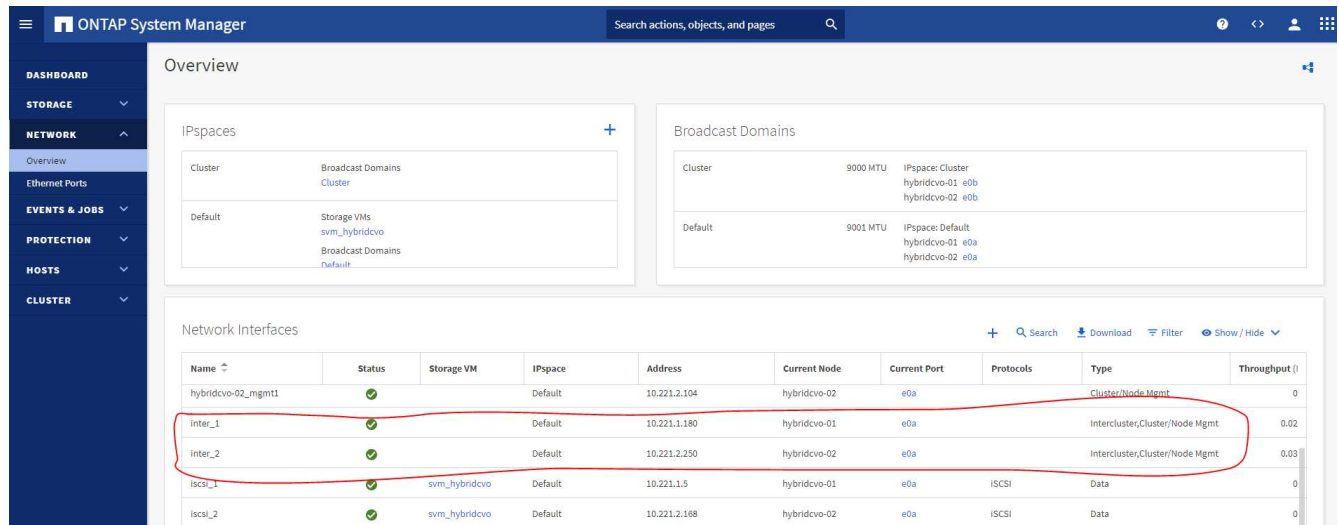
1. Configure intercluster LIFs for cluster peering on both the on-premises cluster and the CVO cluster instance. This step can be performed with ONTAP System Manager. A default CVO deployment has intercluster LIFs configured automatically.

On-premises cluster:



Name	Status	Storage VM	IPspace	Address	Current Node	Current Port	Protocols	Type
onPrem-01_IC	✓		Default	192.168.0.113	onPrem-01	e0b		Intercluster
onPrem-01_mgmt1	✓		Default	192.168.0.111	onPrem-01	e0c		Cluster/Node Mgmt
cluster_mgmt	✓		Default	192.168.0.101	onPrem-01	e0a		Cluster/Node Mgmt

Target CVO cluster:




Name	Status	Storage VM	IPspace	Address	Current Node	Current Port	Protocols	Type	Throughput (I)
hybridcvo-02_mgmtt	✓		Default	10.221.2.104	hybridcvo-02	e0a		Cluster/Node Mgmt	0
Inter_1	✓		Default	10.221.1.180	hybridcvo-01	e0a		Intercluster,Cluster/Node Mgmt	0.02
Inter_2	✓		Default	10.221.2.250	hybridcvo-02	e0a		Intercluster,Cluster/Node Mgmt	0.03
iscsi_1	✓	svm_hybridcvo	Default	10.221.1.5	hybridcvo-01	e0a	ISCSI	Data	0
iscsi_2	✓	svm_hybridcvo	Default	10.221.2.168	hybridcvo-02	e0a	ISCSI	Data	0


2. With the intercluster LIFs configured, cluster peering and volume replication can be set up by using drag-and-drop in NetApp Cloud Manager. See ["Getting Started - AWS Public Cloud"](#) for details.

Alternatively, cluster peering and DB volume replication can be performed by using ONTAP System Manager as follows:


3. Log into ONTAP System Manager. Navigate to Cluster > Settings and click Peer Cluster to set up cluster peering with the CVO instance in the cloud.



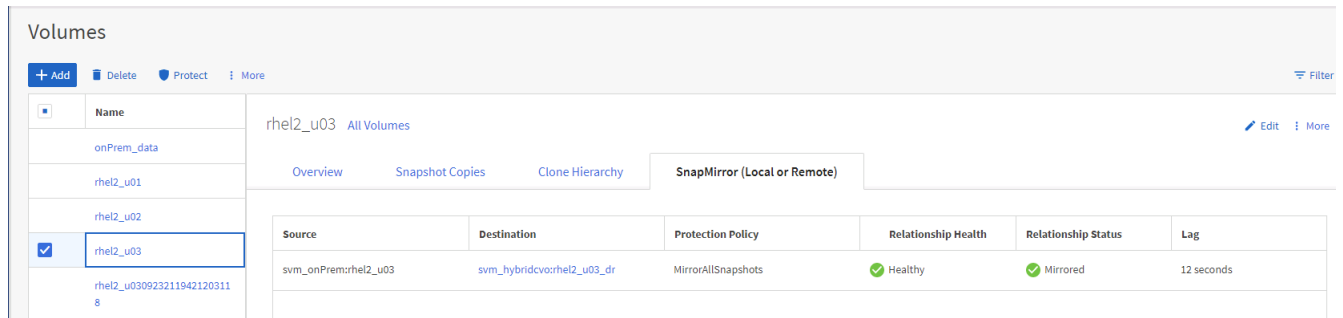
4. Go to the Volumes tab. Select the database volume to be replicated and click Protect.



5. Set the protection policy to Asynchronous. Select the destination cluster and storage SVM.




6. Validate that the volume is synced between the source and target and that the replication relationship is healthy.

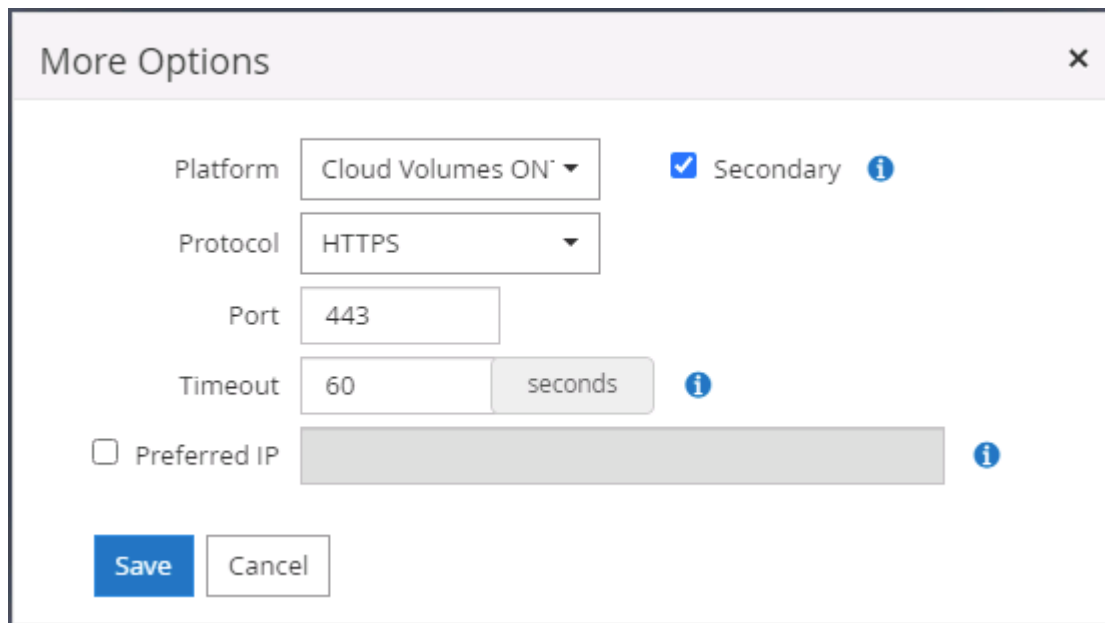


6. Add CVO database storage SVM to SnapCenter


1. Log into SnapCenter with a user ID with SnapCenterAdmin privileges.
2. Click the Storage System tab from the menu, and then click New to add a CVO storage SVM that hosts replicated target database volumes to SnapCenter. Enter the cluster management IP in the Storage System field, and enter the appropriate username and password.



3. Click More Options to open additional storage configuration options. In the Platform field, select Cloud Volumes ONTAP, check Secondary, and then click Save.



4. Assign the storage systems to SnapCenter database management user IDs as shown in 3. [SnapCenter host plugin installation](#).

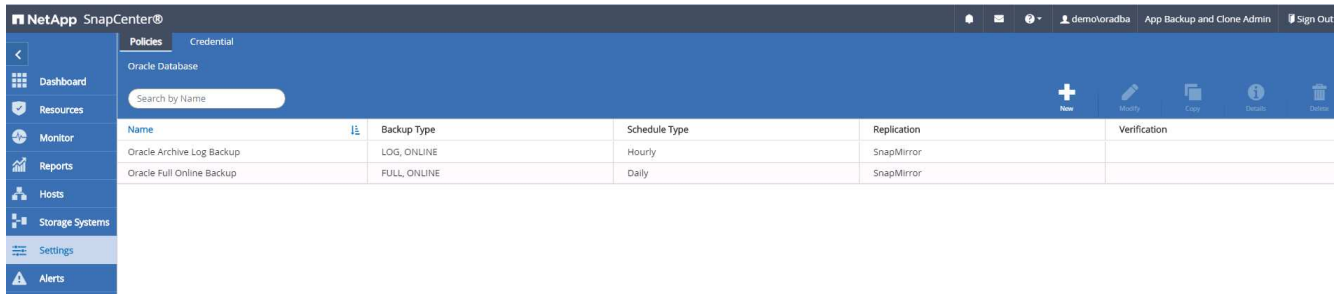


7. Setup database backup policy in SnapCenter

The following procedures demonstrates how to create a full database or log file backup policy. The policy can then be implemented to protect databases resources. The recovery point objective (RPO) or recovery time objective (RTO) dictates the frequency of database and/or log backups.

Create a full database backup policy for Oracle

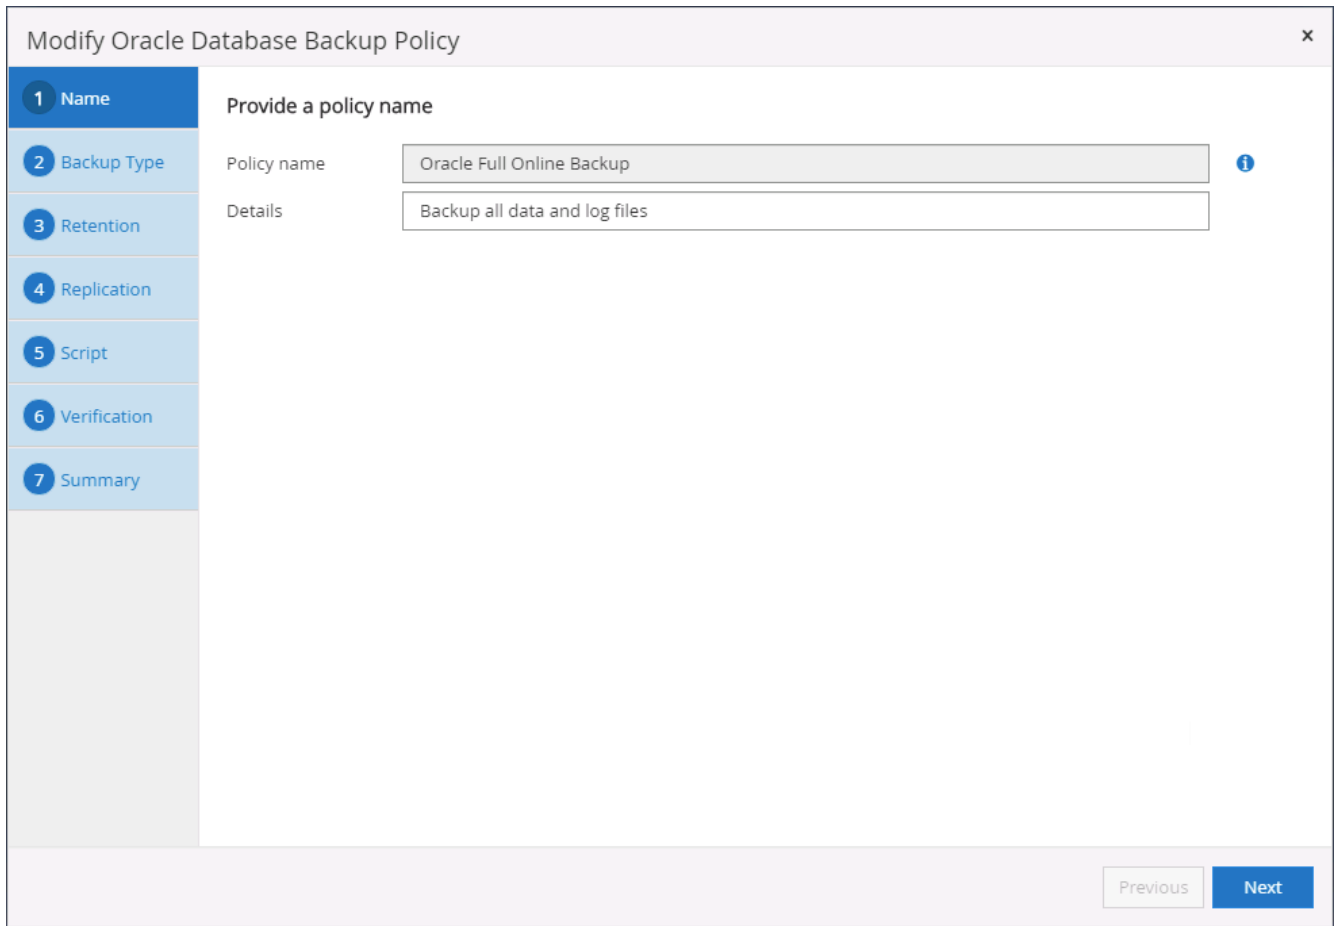
1. Log into SnapCenter as a database management user ID, click Settings, and then click Polices.



The screenshot shows the NetApp SnapCenter interface. The top navigation bar includes 'Policies' and 'Credential'. Below the navigation bar, there is a search box and a table of policies. The table has columns for Name, Backup Type, Schedule Type, Replication, and Verification. Two policies are listed: 'Oracle Archive Log Backup' and 'Oracle Full Online Backup'.

Name	Backup Type	Schedule Type	Replication	Verification
Oracle Archive Log Backup	LOG, ONLINE	Hourly	SnapMirror	
Oracle Full Online Backup	FULL, ONLINE	Daily	SnapMirror	

2. Click New to launch a new backup policy creation workflow or choose an existing policy for modification.



The screenshot shows the 'Modify Oracle Database Backup Policy' workflow screen. The screen is divided into a left sidebar with numbered steps (1-7) and a main content area. Step 1, 'Name', is selected. The main content area has a title 'Provide a policy name' and two input fields: 'Policy name' with the value 'Oracle Full Online Backup' and 'Details' with the value 'Backup all data and log files'. At the bottom right, there are 'Previous' and 'Next' buttons.

3. Select the backup type and schedule frequency.

Modify Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select Oracle database backup options

Choose backup type

Online backup

- Datafiles, control files, and archive logs
- Datafiles and control files
- Archive logs

Offline backup ?

- Mount
- Shutdown
- Save state of PDBs ?

Choose schedule frequency

Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling you to stagger your start times.

On demand

Hourly

Daily

Previous Next

4. Set the backup retention setting. This defines how many full database backup copies to keep.

Modify Oracle Database Backup Policy x

- 1 Name
- 2 Backup Type
- 3 Retention**
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Retention settings ?

Daily retention settings

Data backup retention settings ?

Total Snapshot copies to keep

Keep Snapshot copies for days

Archive Log backup retention settings

Total Snapshot copies to keep

Keep Snapshot copies for days

Previous Next

5. Select the secondary replication options to push local primary snapshots backups to be replicated to a secondary location in cloud.

Modify Oracle Database Backup Policy ×

- Name
- Backup Type
- Retention
- Replication**
- Script
- Verification
- Summary

Select secondary replication options ⓘ

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label ⓘ

Error retry count ⓘ

6. Specify any optional script to run before and after a backup run.

Modify Oracle Database Backup Policy x

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script**
- 6 Verification
- 7 Summary

Specify optional scripts to run before and after performing a backup job

Prescript full path

Prescript arguments

Postscript full path

Postscript arguments

Script timeout

7. Run backup verification if desired.

Modify Oracle Database Backup Policy x

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script
- 6 Verification**
- 7 Summary

Select the options to run backup verification

Run Verifications for following backup schedules

Select how often you want the schedules to occur in the policy. The specific verification times are set at backup job creation enabling you to stagger your verification start times.

Daily

Verification script commands

Script timeout secs

Prescript full path

Prescript arguments

Postscript full path

Postscript arguments

8. Summary.

Modify Oracle Database Backup Policy

1 Name	Summary	
2 Backup Type	Policy name	Oracle Full Online Backup
3 Retention	Details	Backup all data and log files
4 Replication	Backup type	Online backup
5 Script	Schedule type	Daily
6 Verification	RMAN catalog backup	Disabled
7 Summary	Archive log pruning	None
	On demand data backup retention	None
	On demand archive log backup retention	None
	Hourly data backup retention	None
	Hourly archive log backup retention	None
	Daily data backup retention	Delete Snapshot copies older than : 14 days
	Daily archive log backup retention	Delete Snapshot copies older than : 14 days
	Weekly data backup retention	None
	Weekly archive log backup retention	None
	Monthly data backup retention	None
	Monthly archive log backup retention	None
	Replication	SnapMirror enabled , Secondary policy label: Daily , Error retry count: 3

Previous Finish

Create a database log backup policy for Oracle

1. Log into SnapCenter with a database management user ID, click Settings, and then click Policies.
2. Click New to launch a new backup policy creation workflow, or choose an existing policy for modification.

New Oracle Database Backup Policy x

- 1 Name**
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Provide a policy name

Policy name i

Details

3. Select the backup type and schedule frequency.

New Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select Oracle database backup options

Choose backup type

- Online backup
 - Datafiles, control files, and archive logs
 - Datafiles and control files
 - Archive logs
- Offline backup i
 - Mount
 - Shutdown
 - Save state of PDBs i

Choose schedule frequency

Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling you to stagger your start times.

- On demand
- Hourly
- Daily

Previous Next

4. Set the log retention period.

New Oracle Database Backup Policy ✕

- 1 Name
- 2 Backup Type
- 3 Retention**
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Retention settings ⓘ

Hourly retention settings

Data backup retention settings ⓘ

Total Snapshot copies to keep

Keep Snapshot copies for days

Archive Log backup retention settings

Total Snapshot copies to keep

Keep Snapshot copies for days

Previous Next

5. Enable replication to a secondary location in the public cloud.

New Oracle Database Backup Policy ×

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication**
- 5 Script
- 6 Verification
- 7 Summary

Select secondary replication options ⓘ

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label: ⓘ

Error retry count: ⓘ

6. Specify any optional scripts to run before and after log backup.

New Oracle Database Backup Policy x

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Specify optional scripts to run before and after performing a backup job

Prescript full path

Prescript arguments

Postscript full path

Postscript arguments

Script timeout

7. Specify any backup verification scripts.

New Oracle Database Backup Policy ×

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script
- 6 Verification**
- 7 Summary

Select the options to run backup verification

Run Verifications for following backup schedules

Select how often you want the schedules to occur in the policy. The specific verification times are set at backup job creation enabling you to stagger your verification start times.

Verification script commands

Script timeout secs

Prescript full path

Prescript arguments

Postscript full path

Postscript arguments

8. Summary.

New Oracle Database Backup Policy

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary


Summary

Policy name	Oracle Archive Log Backup
Details	
Backup Oracle archive logs	
Backup type	Online backup
Schedule type	
Hourly	
RMAN catalog backup	Disabled
Archive log pruning	None
On demand data backup retention	None
On demand archive log backup retention	None
Hourly data backup retention	None
Hourly archive log backup retention	Delete Snapshot copies older than : 7 days
Daily data backup retention	None
Daily archive log backup retention	None
Weekly data backup retention	None
Weekly archive log backup retention	None
Monthly data backup retention	None
Monthly archive log backup retention	None
Replication	SnapMirror enabled , Secondary policy label: Hourly , Error retry count: 3

Previous Finish

Create a full database backup policy for SQL

1. Log into SnapCenter with a database management user ID, click Settings, and then click Policies.



2. Click New to launch a new backup policy creation workflow, or choose an existing policy for modification.

New SQL Server Backup Policy x

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Provide a policy name

Policy name i

Details

Backup all data and log files

PreviousNext

3. Define the backup option and schedule frequency. For SQL Server configured with an availability group, a preferred backup replica can be set.

New SQL Server Backup Policy x

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select SQL server backup options

Choose backup type

Full backup and log backup

Full backup

Log backup

Copy only backup i

Maximum databases backed up per Snapshot copy: i

Availability Group Settings v

Schedule frequency

Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling you to stagger your start times.

On demand

Hourly

Daily

Weekly

Monthly

4. Set the backup retention period.

New SQL Server Backup Policy x

- 1 Name
- 2 Backup Type
- 3 Retention**
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Retention settings

Retention settings for up-to-the-minute restore operation ⓘ

Keep log backups applicable to last full backups

Keep log backups applicable to last days

Full backup retention settings ⓘ

Daily

Total Snapshot copies to keep

Keep Snapshot copies for days

5. Enable backup copy replication to a secondary location in cloud.

New SQL Server Backup Policy x

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select secondary replication options i

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label i

Error retry count i

6. Specify any optional scripts to run before or after a backup job.

New SQL Server Backup Policy x

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script**
- 6 Verification
- 7 Summary

Specify optional scripts to run before performing a backup job

Prescript full path

Prescript arguments

Specify optional scripts to run after performing a backup job

Postscript full path

Postscript arguments

Script timeout

7. Specify the options to run backup verification.

New SQL Server Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select the options to run backup verification

Run verifications for the following backup schedules

Select how often you want the schedules to occur in the policy. The specific verification times are set at backup job creation enabling you to stagger your verification start times.

Daily

Database consistency checks options

Limit the integrity structure to physical structure of the database (PHYSICAL_ONLY)

Suppress all information message (NO_INFOMSGS)

Display all reported error messages per object (ALL_ERRORMSGs)

Do not check non-clustered indexes (NOINDEX)

Limit the checks and obtain the locks instead of using an internal database Snapshot copy (TABLOCK)

Log backup

Verify log backup. **i**

Verification script settings

Script timeout secs

Previous Next

8. Summary.

New SQL Server Backup Policy
✕

1 Name	Summary	
2 Backup Type	Policy name	SQL Server Full Backup
3 Retention	Details	Backup all data and log files
4 Replication	Backup type	Full backup and log backup
5 Script	Availability group settings	Backup only on preferred backup replica
6 Verification	Schedule Type	Daily
7 Summary	UTM retention	Total backup copies to retain : 7
	Daily Full backup retention	Total backup copies to retain : 7
	Replication	SnapMirror enabled , Secondary policy label: Daily , Error retry count: 3
	Backup prescript settings	undefined Prescript arguments:
	Backup postscript settings	undefined Postscript arguments:
	Verification for backup schedule type	none
	Verification prescript settings	undefined Prescript arguments:
	Verification postscript settings	undefined Postscript arguments:

Previous
Finish

Create a database log backup policy for SQL.

1. Log into SnapCenter with a database management user ID, click Settings > Policies, and then New to launch a new policy creation workflow.

New SQL Server Backup Policy x

1 Name Provide a policy name

2 Backup Type Policy name i

3 Retention Details

4 Replication

5 Script

6 Verification

7 Summary

2. Define the log backup option and schedule frequency. For SQL Server configured with a availability group, a preferred backup replica can be set.

New SQL Server Backup Policy x

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select SQL server backup options

Choose backup type

Full backup and log backup

Full backup

Log backup

Copy only backup i

Maximum databases backed up per Snapshot copy: i

Availability Group Settings v

Schedule frequency

Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling you to stagger your start times.

On demand

Hourly

Daily

Weekly

Monthly

3. SQL server data backup policy defines the log backup retention; accept the defaults here.

New SQL Server Backup Policy ×

- 1 Name
- 2 Backup Type
- 3 Retention**
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Log backup retention settings

Up-to-the-minute (UTM) retention settings retains log backups created as part of full backup and full and log backup operations. UTM retention settings also decides for how many full backups the log backups are to be retained. For example, if UTM retention settings is configured to retain log backups of the last 5 full backups, then the log backups of the last 5 full backups are retained and the rest are deleted.

[Previous](#) [Next](#)

4. Enable log backup replication to secondary in the cloud.

New SQL Server Backup Policy ×

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select secondary replication options ⓘ

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label: Hourly ⓘ

Error retry count: 3 ⓘ

Previous Next

5. Specify any optional scripts to run before or after a backup job.

New SQL Server Backup Policy ×

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script**
- 6 Verification
- 7 Summary

Specify optional scripts to run before performing a backup job

Prescript full path

Prescript arguments

Specify optional scripts to run after performing a backup job

Postscript full path

Postscript arguments

Script timeout

6. Summary.

New SQL Server Backup Policy ✕

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Summary

Policy name	SQL Server Log Backup
Details	
Backup SQL server log	
Backup type	Log transaction backup
Availability group settings	
Backup only on preferred backup replica	
Schedule Type	Hourly
Replication	
SnapMirror enabled , Secondary policy label: Hourly , Error retry count: 3	
Backup prescript settings	
undefined	
Prescript arguments:	
Backup postscript settings	
undefined	
Postscript arguments:	
Verification for backup schedule type	
none	
Verification prescript settings	
undefined	
Prescript arguments:	
Verification postscript settings	
undefined	
Postscript arguments:	


Previous
Finish

8. Implement backup policy to protect database


SnapCenter uses a resource group to backup a database in a logical grouping of database resources, such as multiple databases hosted on a server, a database sharing the same storage volumes, multiple databases supporting a business application, and so on. Protecting a single database creates a resource group of its own. The following procedures demonstrate how to implement a backup policy created in section 7 to protect Oracle and SQL Server databases.

Create a resource group for full backup of Oracle

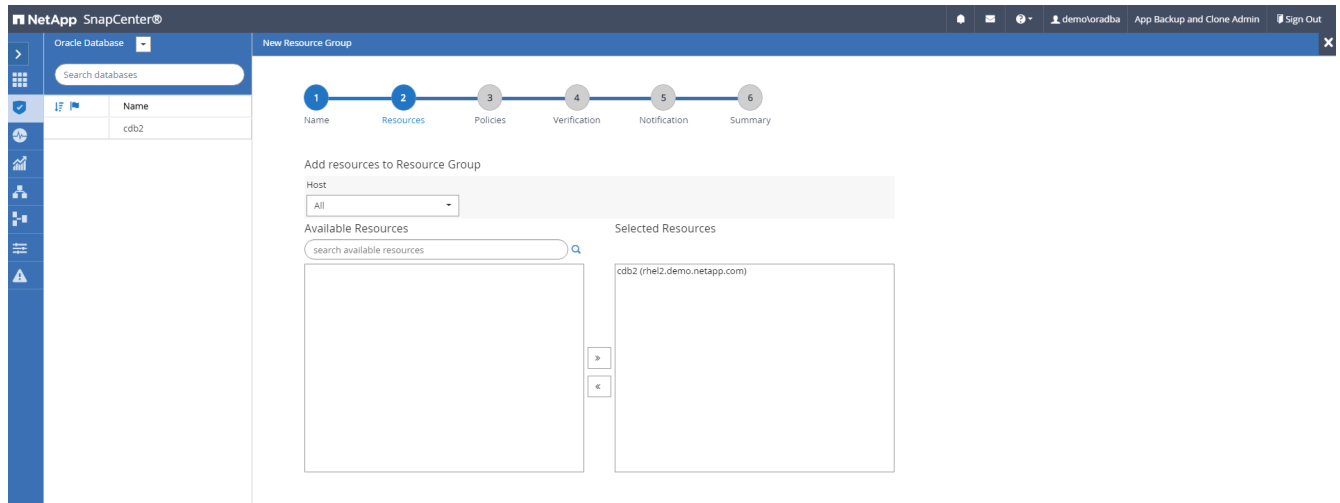
1. Log into SnapCenter with a database management user ID, and navigate to the Resources tab. In the View drop-down list, choose either Database or Resource Group to launch the resource group creation workflow.



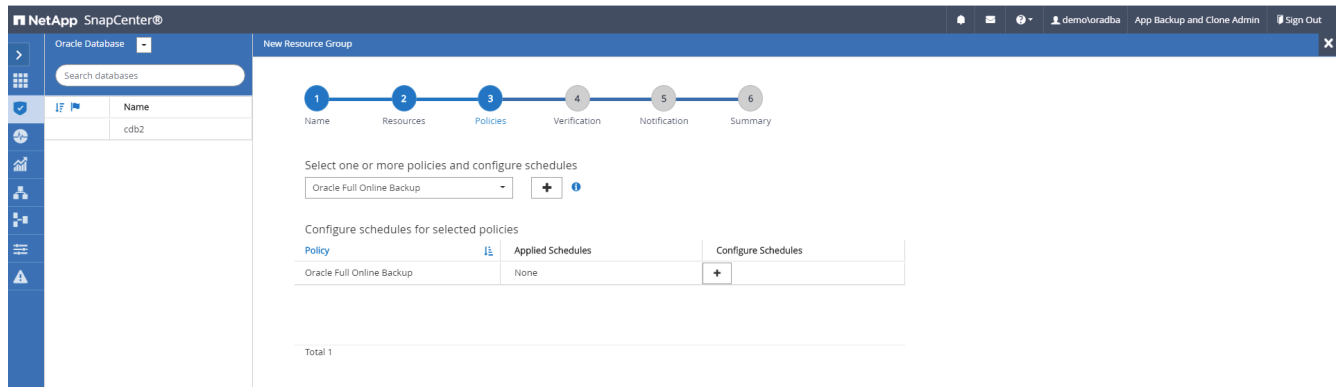
2. Provide a name and tags for the resource group. You can define a naming format for the Snapshot copy and bypass the redundant archive log destination if configured.




3. Add database resources to the resource group.



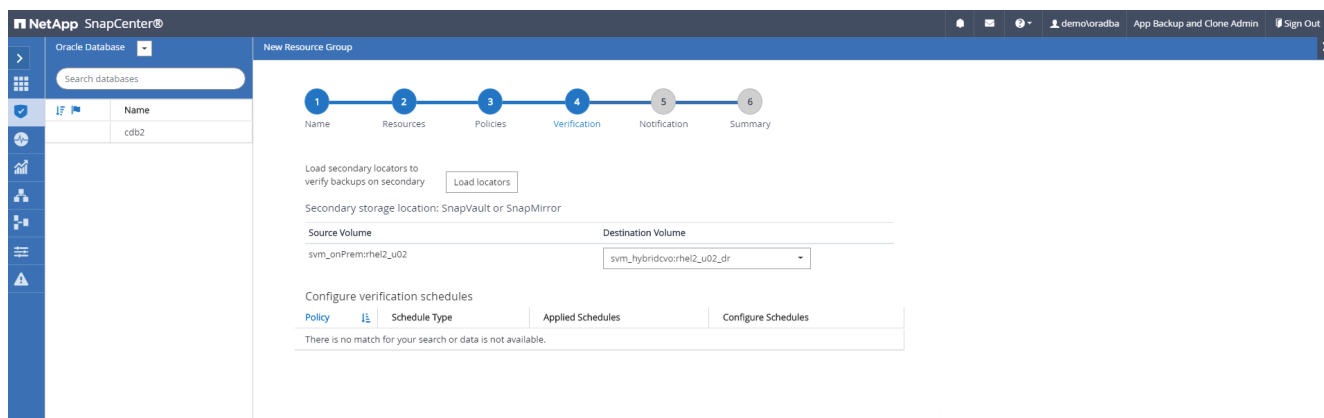
4. Select a full backup policy created in section 7 from the drop-down list.




5. Click the (+) sign to configure the desired backup schedule.




6. Click Load Locators to load the source and destination volume.



7. Configure the SMTP server for email notification if desired.

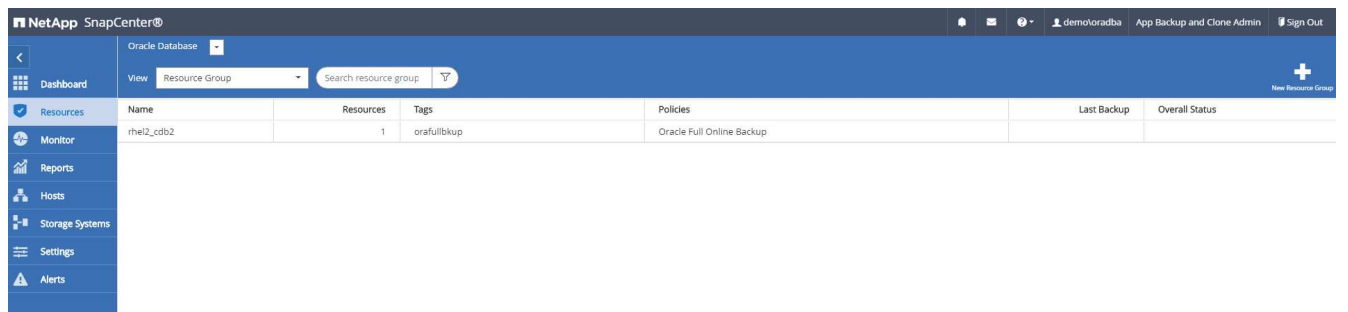


8. Summary.




Create a resource group for log backup of Oracle


1. Log into SnapCenter with a database management user ID, and navigate to the Resources tab. In the View drop-down list, choose either Database or Resource Group to launch the resource group creation workflow.



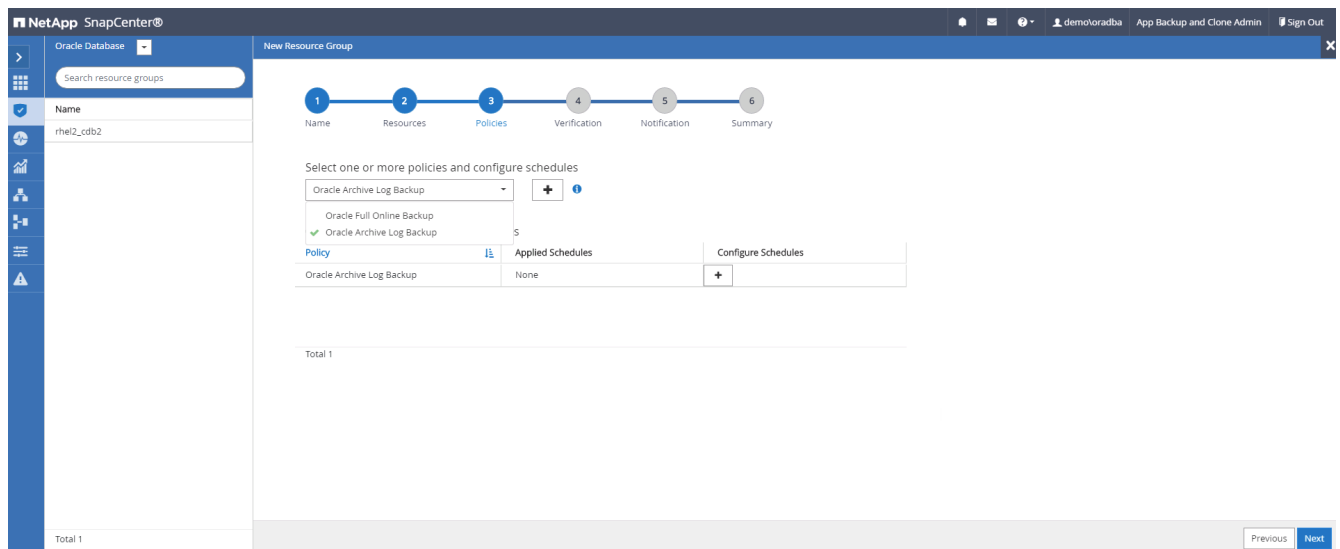
2. Provide a name and tags for the resource group. You can define a naming format for the Snapshot copy and bypass the redundant archive log destination if configured.




3. Add database resources to the resource group.



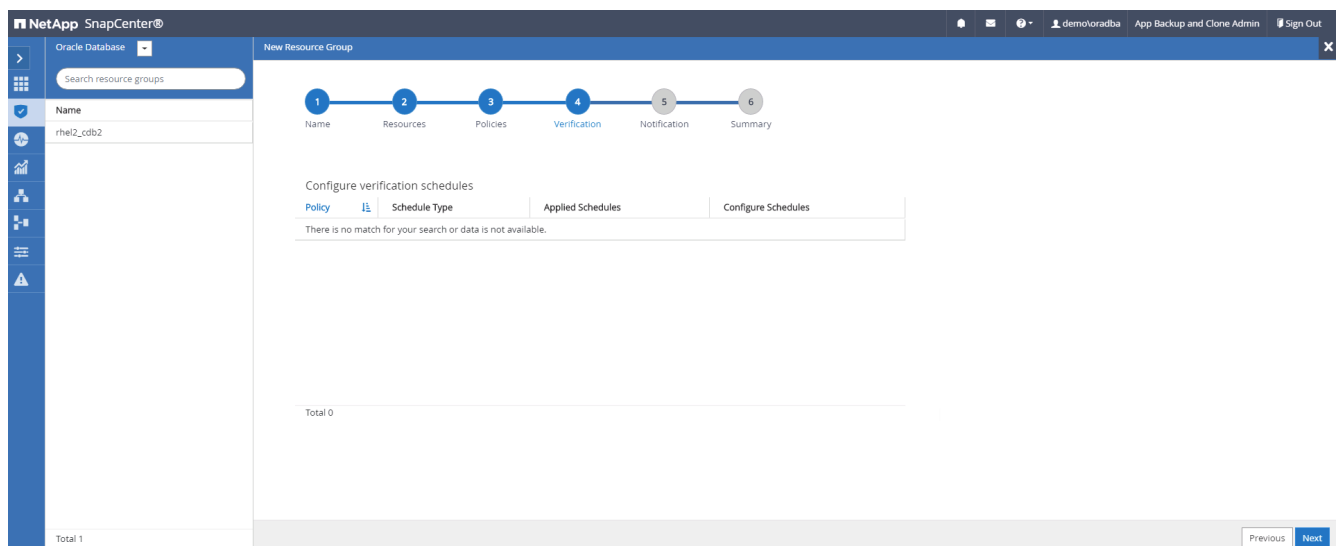
4. Select a log backup policy created in section 7 from the drop-down list.




5. Click on the (+) sign to configure the desired backup schedule.




6. If backup verification is configured, it displays here.



7. Configure an SMTP server for email notification if desired.




8. Summary.




Create a resource group for full backup of SQL Server

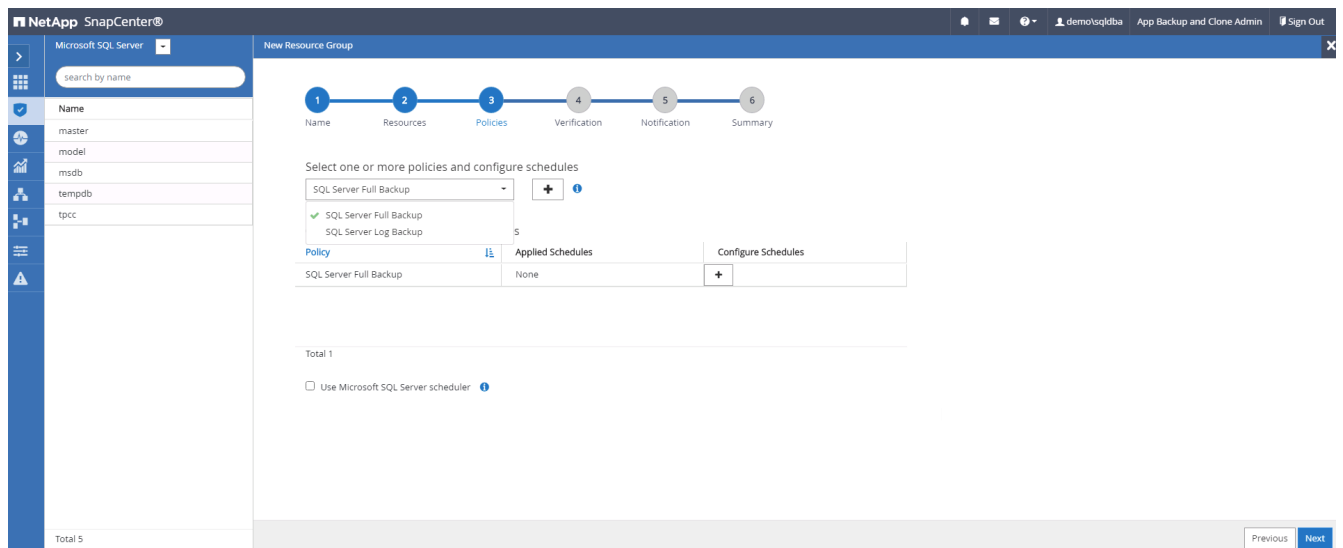
1. Log into SnapCenter with a database management user ID, and navigate to the Resources tab. In the View drop-down list, choose either a Database or Resource Group to launch the resource group creation workflow. Provide a name and tags for the resource group. You can define a naming format for the Snapshot copy.



2. Select the database resources to be backed up.



3. Select a full SQL backup policy created in section 7.



4. Add exact timing for backups as well as the frequency.

Add schedules for policy SQL Server Full Backup

Daily

Start date: 09/10/2021 6:20 PM

Expires on: 12/31/2021 6:20 PM

Repeat every: 1 days

i The schedules are triggered in the SnapCenter Server time zone.

Cancel OK

5. Choose the verification server for the backup on secondary if backup verification is to be performed. Click Load Locator to populate the secondary storage location.

NetApp SnapCenter

Microsoft SQL Server

New Resource Group

1 Name 2 Resources 3 Policies 4 Verification 5 Notification 6 Summary

Select the verification servers

Verification server: Select one or more servers

Load secondary locators to verify backups on secondary: Load locators

Secondary storage location: SnapVault or SnapMirror

Source Volume	Destination Volume
svm_onPrem:sql1_data	svm_hybridov:sql1_data_dr
svm_onPrem:sql1_log	svm_hybridov:sql1_log_dr


Configure verification schedules

Policy | Schedule Type | Applied Schedules | Configure Schedules

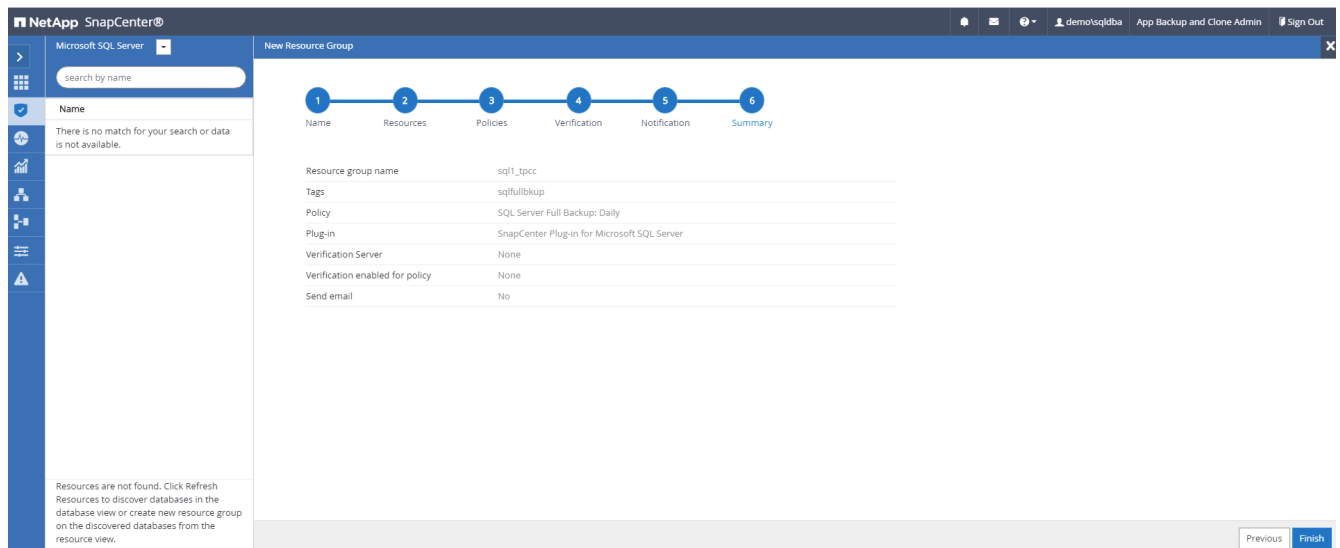
There is no match for your search or data is not available.

Previous Next

6. Configure the SMTP server for email notification if desired.




7. Summary.




Create a resource group for log backup of SQL Server


1. Log into SnapCenter with a database management user ID, and navigate to the Resources tab. In the View drop-down list, choose either a Database or Resource Group to launch the resource group creation workflow. Provide the name and tags for the resource group. You can define a naming format for the Snapshot copy.




2. Select the database resources to be backed up.




3. Select a SQL log backup policy created in section 7.




4. Add exact timing for the backup as well as the frequency.




5. Choose the verification server for the backup on secondary if backup verification is to be performed. Click the Load Locator to populate the secondary storage location.



6. Configure the SMTP server for email notification if desired.



7. Summary.



9. Validate backup

After database backup resource groups are created to protect database resources, the backup jobs runs according to the predefined schedule. Check the job execution status under the Monitor tab.

ID	Status	Name	Start date	End date	Owner
532	✓	Backup of Resource Group 'sql1_tpsc_log' with policy 'SQL Server Log Backup'	09/14/2021 8:35:01 PM	09/14/2021 8:37:10 PM	demo@sqldba
528	✓	Backup of Resource Group 'sql1_tpsc_log' with policy 'SQL Server Log Backup'	09/14/2021 7:35:01 PM	09/14/2021 7:37:09 PM	demo@sqldba
524	✓	Backup of Resource Group 'sql1_tpsc_log' with policy 'SQL Server Log Backup'	09/14/2021 6:35:01 PM	09/14/2021 6:37:08 PM	demo@sqldba
521	✓	Backup of Resource Group 'sql1_tpsc' with policy 'SQL Server Full Backup'	09/14/2021 6:25:01 PM	09/14/2021 6:27:14 PM	demo@sqldba
517	✓	Backup of Resource Group 'sql1_tpsc_log' with policy 'SQL Server Log Backup'	09/14/2021 5:35:01 PM	09/14/2021 5:37:09 PM	demo@sqldba
513	✓	Backup of Resource Group 'sql1_tpsc_log' with policy 'SQL Server Log Backup'	09/14/2021 4:35:01 PM	09/14/2021 4:37:08 PM	demo@sqldba
509	✓	Backup of Resource Group 'sql1_tpsc_log' with policy 'SQL Server Log Backup'	09/14/2021 3:35:01 PM	09/14/2021 3:37:10 PM	demo@sqldba
503	✓	Backup of Resource Group 'sql1_tpsc_log' with policy 'SQL Server Log Backup'	09/14/2021 2:35:01 PM	09/14/2021 2:37:09 PM	demo@sqldba

Go to the Resources tab, click the database name to view details of database backup, and toggle between Local copies and mirror copies to verify that Snapshot backups are replicated to a secondary location in the public cloud.

Backup Name	Count	Type	IF	End Date	Verified	Mounted	RMAN Cataloged	SCN
rhel2_cdb2_09-23-2021_14.35.03.3242_1	1	Log		09/23/2021 2:35:45 PM	Not Applicable	False	Not Cataloged	6872761
rhel2_cdb2_09-23-2021_14.35.03.3242_0	1	Data		09/23/2021 2:35:30 PM	Unverified	False	Not Cataloged	6872715
rhel2_cdb2_09-22-2021_14.35.02.0014_1	1	Log		09/22/2021 2:35:24 PM	Not Applicable	False	Not Cataloged	6737479
rhel2_cdb2_09-22-2021_14.35.02.0014_0	1	Data		09/22/2021 2:35:14 PM	Unverified	False	Not Cataloged	6737395
rhel2_cdb2_09-21-2021_14.35.02.1884_1	1	Log		09/21/2021 2:35:35 PM	Not Applicable	False	Not Cataloged	6598735

At this point, database backup copies in the cloud are ready to clone to run dev/test processes or for disaster recovery in the event of a primary failure.

Getting Started with AWS public cloud

This section describes the process of deploying Cloud Manager and Cloud Volumes ONTAP in AWS.

AWS public cloud



To make things easier to follow, we have created this document based on a deployment in AWS. However, the process is very similar for Azure and GCP.

1. Pre-flight check

Before deployment, make sure that the infrastructure is in place to allow for the deployment in the next stage. This includes the following:

- AWS account
- VPC in your region of choice
- Subnet with access to the public internet
- Permissions to add IAM roles into your AWS account
- A secret key and access key for your AWS user

2. Steps to deploy Cloud Manager and Cloud Volumes ONTAP in AWS



There are many methods for deploying Cloud Manager and Cloud Volumes ONTAP; this method is the simplest but requires the most permissions. If this method is not appropriate for your AWS environment, please consult the [NetApp Cloud Documentation](#).

Deploy the Cloud Manager connector

1. Navigate to [NetApp Cloud Central](#) and log in or sign up.



[Continue to Cloud Manager](#)


Log In to NetApp Cloud Central

Don't have an account yet? [Sign Up](#)

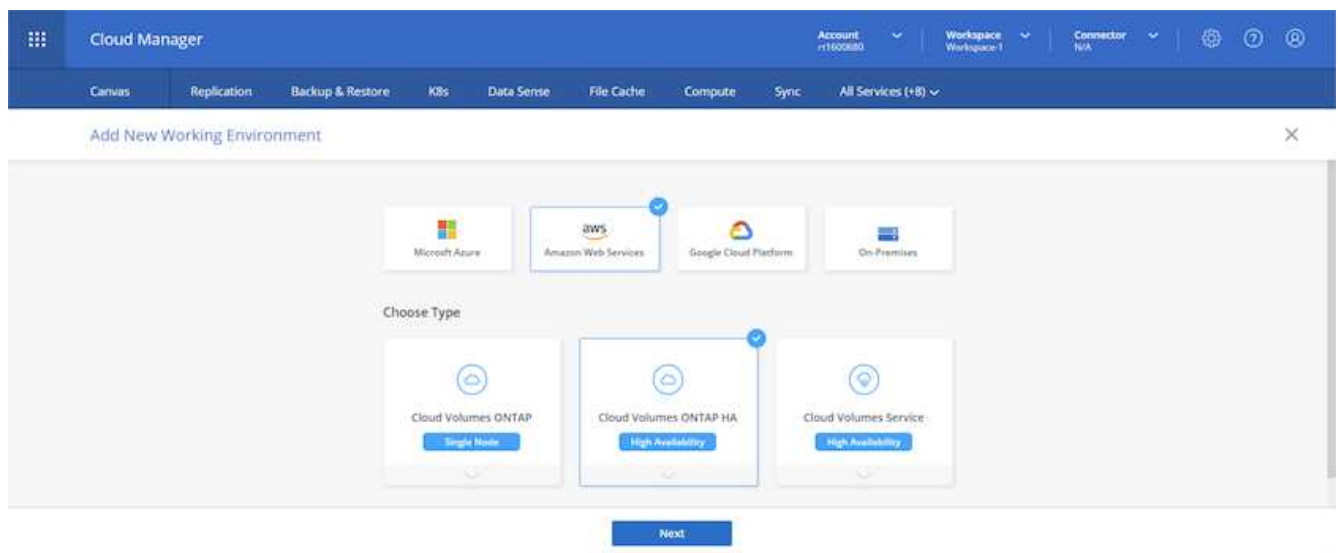
LOGIN

[Forgot your password?](#)

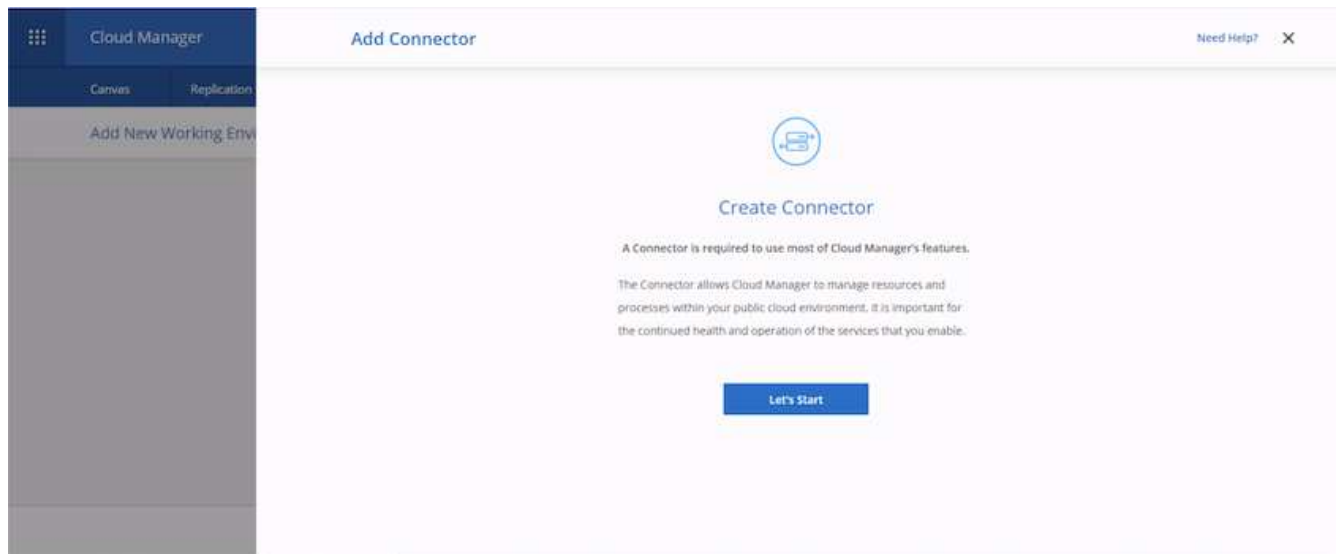
2. After you log in, you should be taken to the Canvas.




3. Click "Add Working Environment" and choose Cloud Volumes ONTAP in AWS. Here, you also choose whether you want to deploy a single node system or a high availability pair. I have chosen to deploy a high availability pair.




4. If no connector has been created, a pop-up appears asking you to create a connector.



5. Click Lets Start, and then choose AWS.



6. Enter your secret key and access key. Make sure that your user has the correct permissions outlined on the [NetApp policies page](#).



7. Give the connector a name and either use a predefined role as described on the [NetApp policies page](#) or ask Cloud Manager to create the role for you.

The screenshot shows the 'Add Connector' wizard in the AWS Cloud Manager console. The 'Details' step is active, showing the following fields and options:

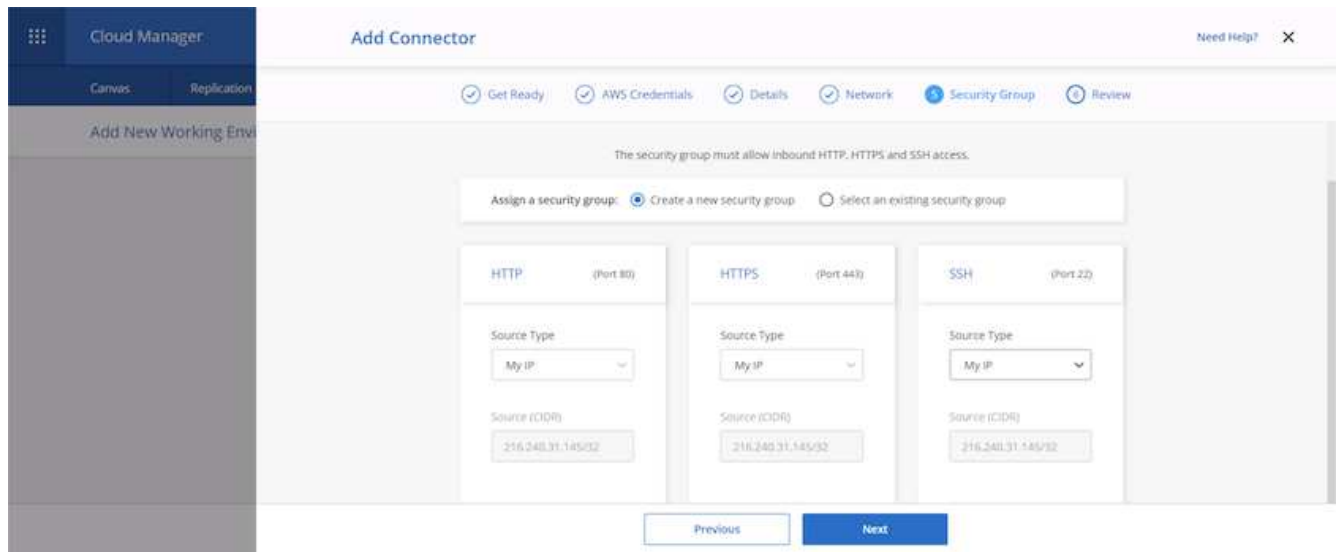
- Connector Instance Name:**
- Connector Role:** Create Role Select an existing Role
- Role Name:**
- Buttons:** 'Previous' and 'Next' buttons are visible at the bottom.

8. Give the networking information needed to deploy the connector. Verify that outbound internet access is enabled by:
- Giving the connector a public IP address
 - Giving the connector a proxy to work through
 - Giving the connector a route to the public internet through an Internet Gateway

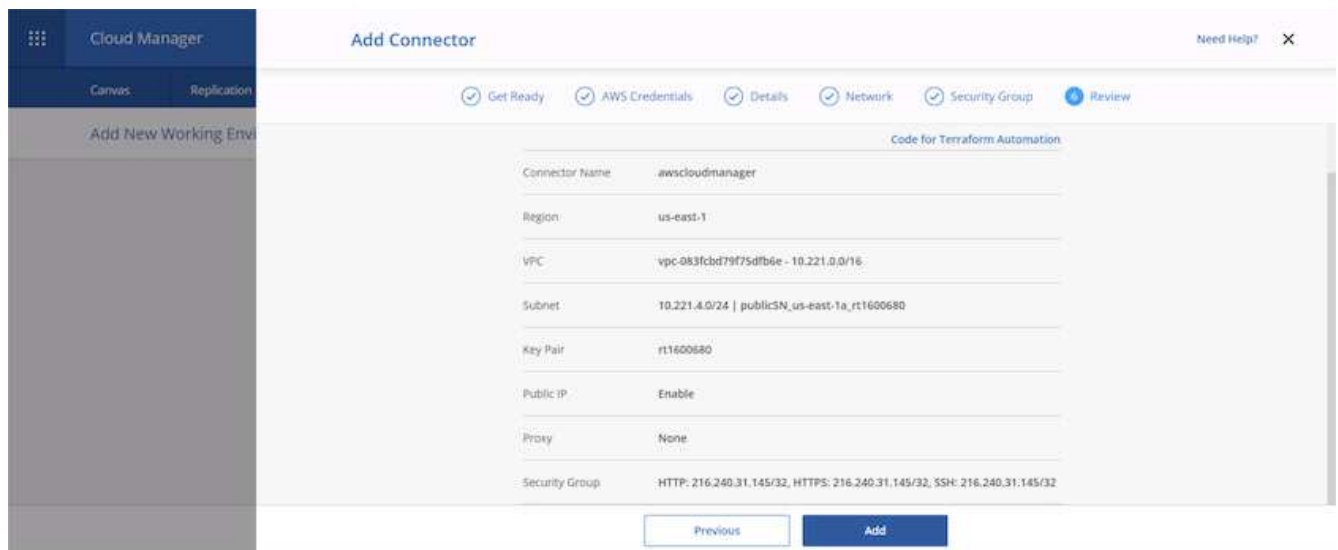
The screenshot shows the 'Add Connector' wizard in the AWS Cloud Manager console, now at the 'Network' step. The following fields and options are visible:

- Connectivity:**
 - VPC:**
 - Subnet:**
 - Key Pair:**
 - Public IP:**
- Proxy Configuration (Optional):**
 - HTTP Proxy:**
 - Define Credentials for this Proxy:**
 - Upload a root certificate:**
- Buttons:** 'Previous' and 'Next' buttons are visible at the bottom.


9. Provide communication with the connector via SSH, HTTP, and HTTPs by either providing a security group or creating a new security group. I have enabled access to the connector from my IP address only.




10. Review the information on the summary page and click Add to deploy the connector.



11. The connector now deploys using a cloud formation stack. You can monitor its progress from Cloud Manager or through AWS.




12. When the deployment is complete, a success page appears.

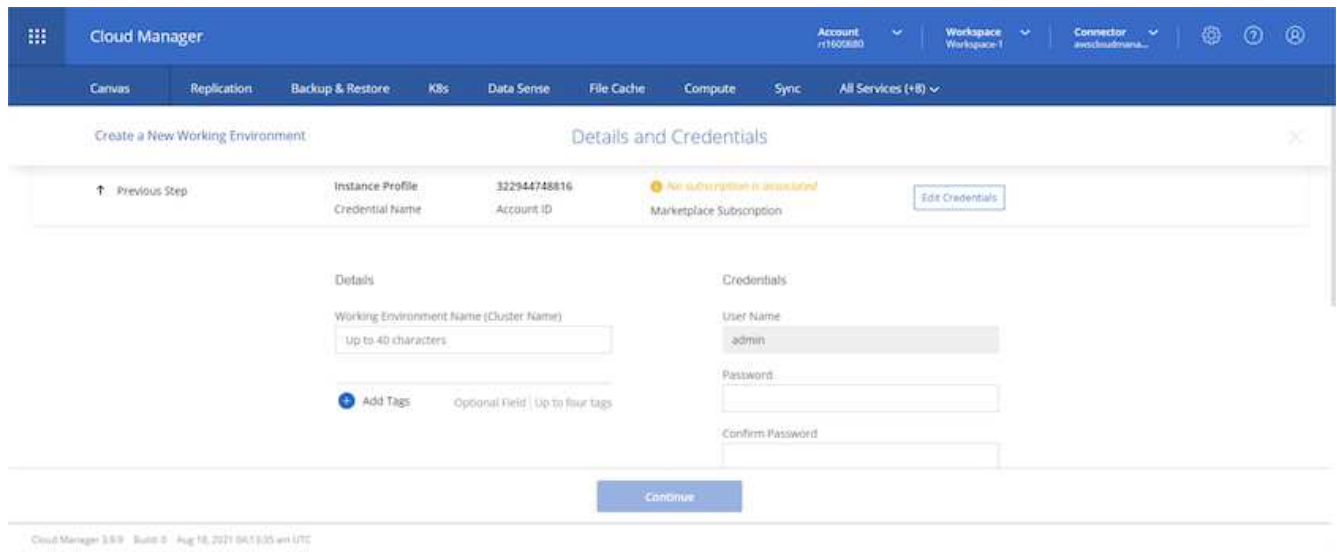


Deploy Cloud Volumes ONTAP

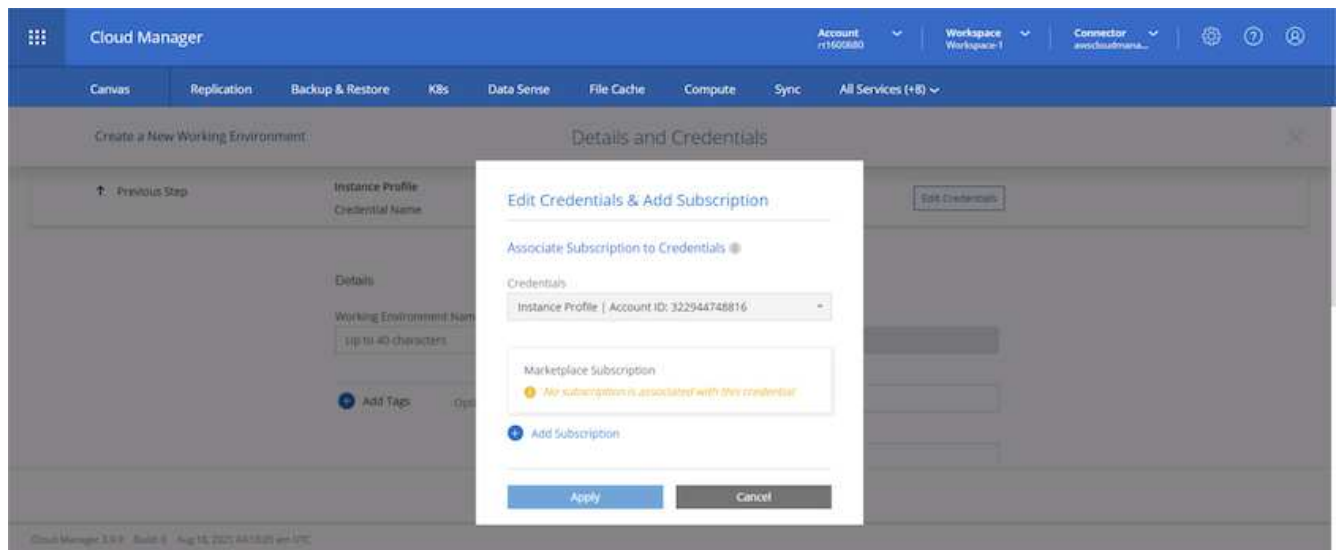
1. Select AWS and the type of deployment based on your requirements.



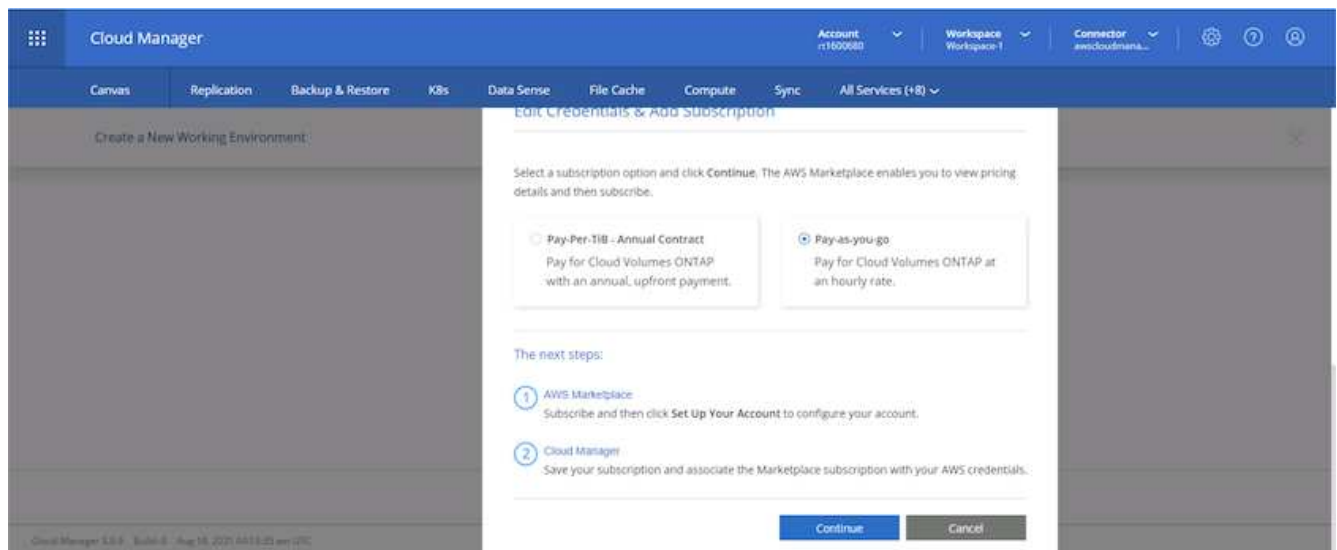
2. If no subscription has been assigned and you wish to purchase with PAYGO, choose Edit Credentials.




3. Choose Add Subscription.



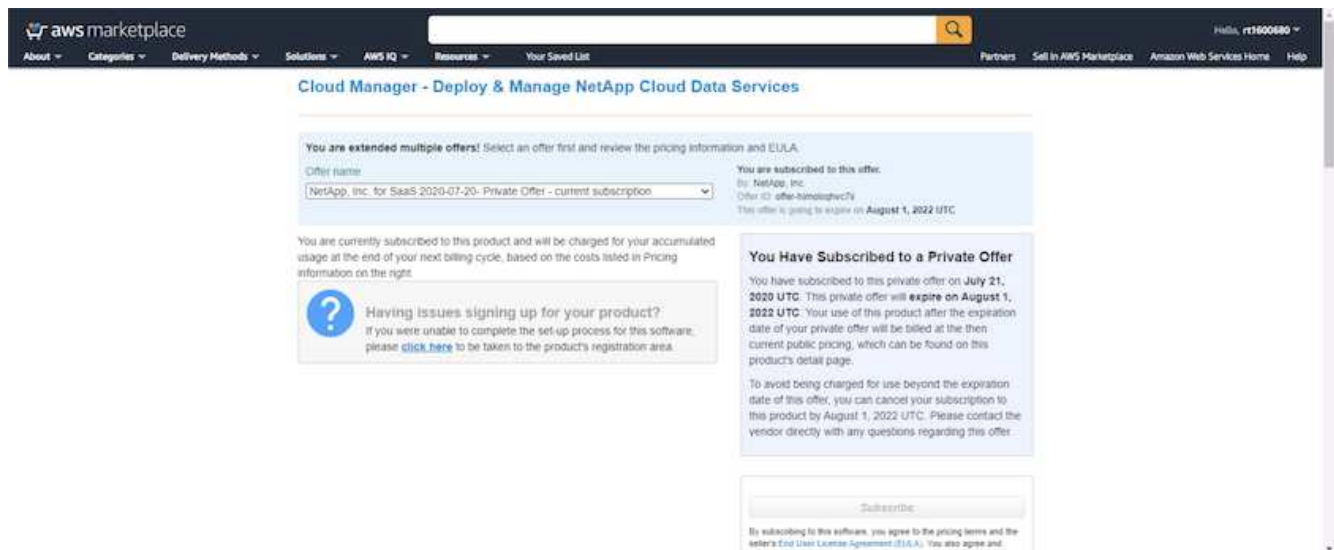
4. Choose the type of contract that you wish to subscribe to. I chose Pay-as-you-go.




5. You are redirected to AWS; choose Continue to Subscribe.



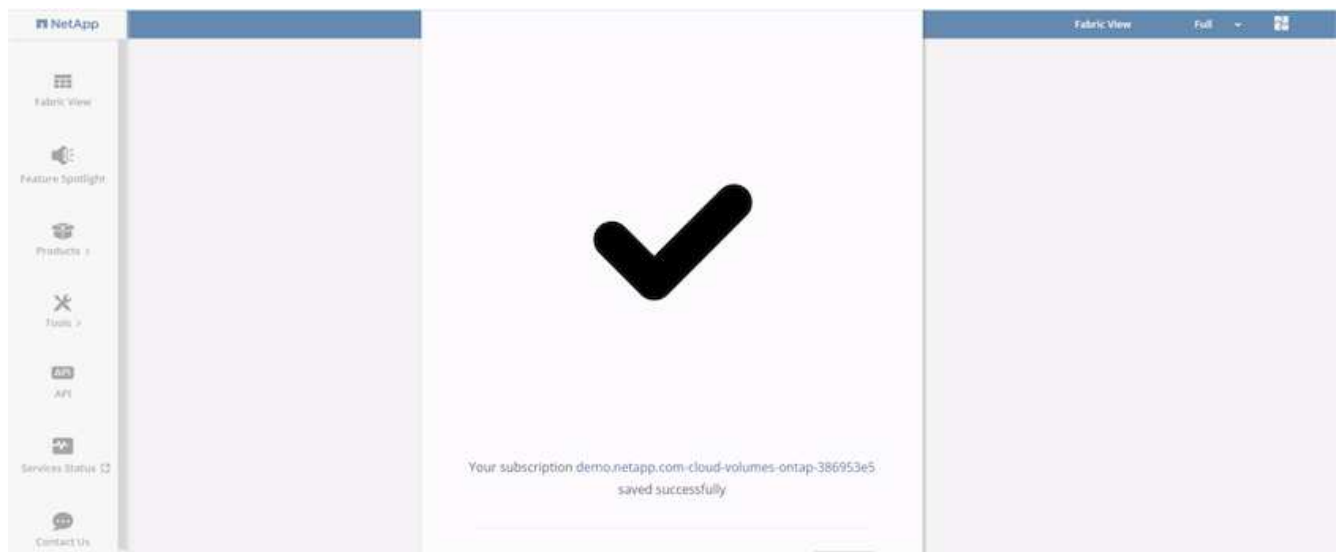
6. Subscribe and you are redirected back to NetApp Cloud Central. If you have already subscribed and don't get redirected, choose the "Click here" link.




7. You are redirected to Cloud Central where you must name your subscription and assign it to your Cloud Central account.



8. When successful, a check mark page appears. Navigate back to your Cloud Manager tab.



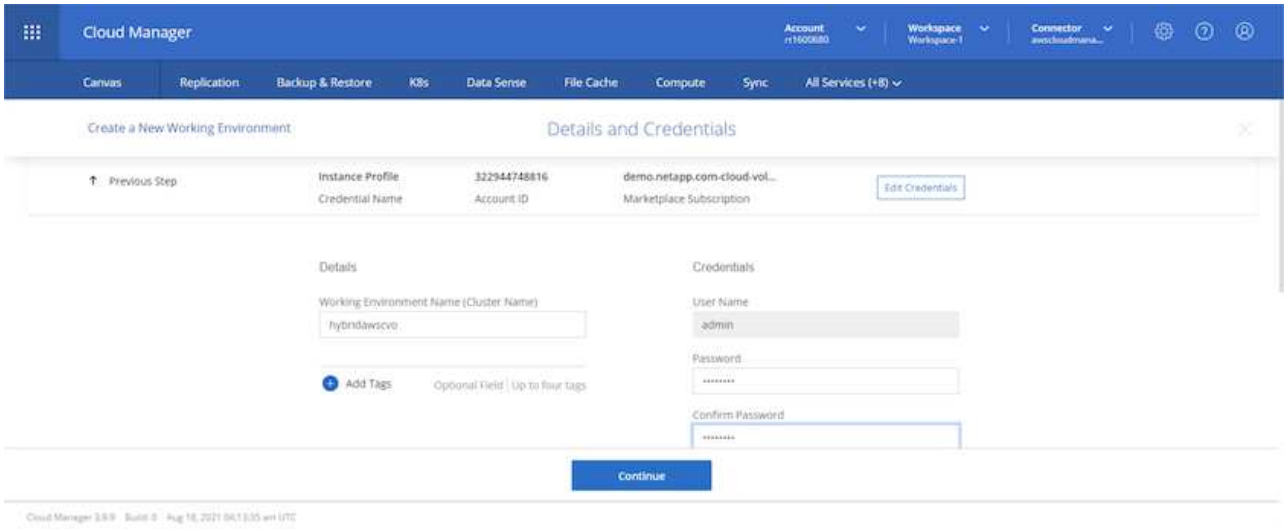
9. The subscription now appears in Cloud Central. Click Apply to continue.



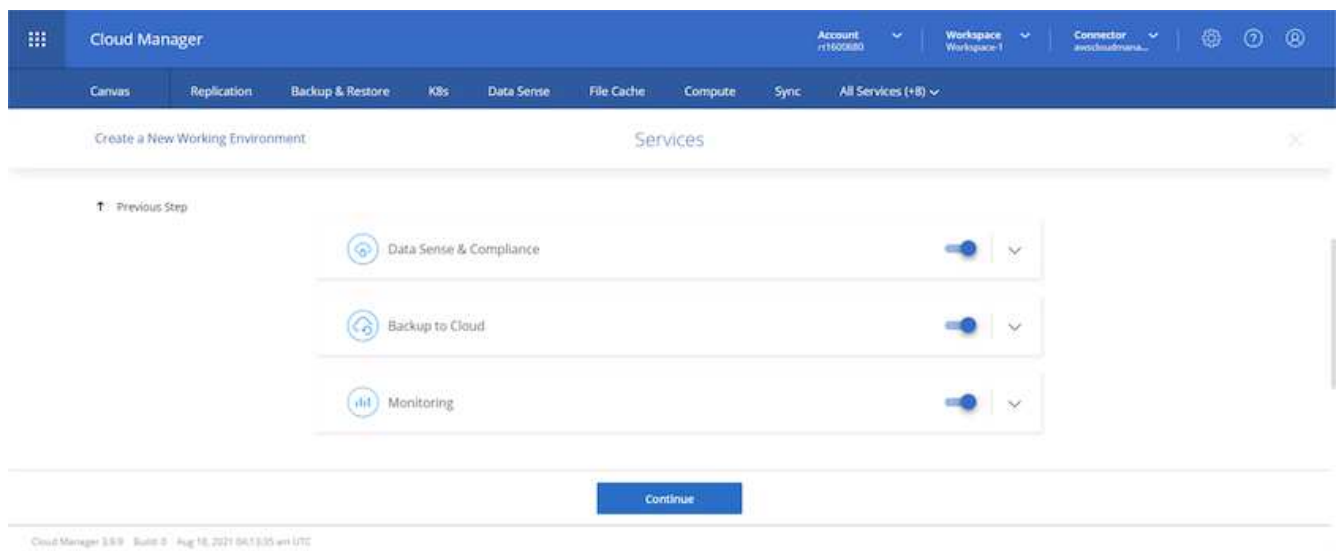
10. Enter the working environment details such as:

- a. Cluster name


- b. Cluster password
- c. AWS tags (Optional)




11. Choose which additional services you would like to deploy. To discover more about these services, visit the [NetApp Cloud Homepage](#).




12. Choose whether to deploy in multiple availability zones (requires three subnets, each in a different AZ), or a single availability zone. I chose multiple AZs.



13. Choose the region, VPC, and security group for the cluster to be deployed into. In this section, you also assign the availability zones per node (and mediator) as well as the subnets that they occupy.



14. Choose the connection methods for the nodes as well as the mediator.





The mediator requires communication with the AWS APIs. A public IP address is not required so long as the APIs are reachable after the mediator EC2 instance has been deployed.

1. Floating IP addresses are used to allow access to the various IP addresses that Cloud Volumes ONTAP uses, including cluster management and data serving IPs. These must be addresses that are not already routable within your network and are added to route tables in your AWS environment. These are required to enable consistent IP addresses for an HA pair during failover. More information about floating IP addresses can be found in the [NetApp Cloud Documentation](#).

Cloud Manager Account: rt1618549 Workspace: Workspace-1 Connector: awscloudmana...

Canvas Replication Backup & Restore K8s Data Sense File Cache Compute Sync All Services (+8)

Create a New Working Environment Floating IPs

↑ Previous Step

Floating IP addresses are required for cluster and SVM access and for NFS and CIFS data access. These floating IPs can migrate between HA nodes if failures occur. To access the data from outside the VPC, you can set up an AWS transit gateway.

You must specify IP addresses that are outside of the CIDR blocks for all VPCs in the selected AWS region.

Floating IP address for cluster management
10.222.0.200

Floating IP address 1 for NFS and CIFS data
10.222.0.201

Floating IP address 2 for NFS and CIFS data
10.222.0.202

Floating IP address for SVM management (Optional)
Enter Floating IP Address

Continue

2. Select which route tables the floating IP addresses are added to. These route tables are used by clients to communicate with Cloud Volumes ONTAP.

Cloud Manager Account: rt1600680 Workspace: Workspace-1 Connector: awscloudmana...

Canvas Replication Backup & Restore K8s Data Sense File Cache Compute Sync All Services (+8)

Create a New Working Environment Route Tables

↑ Previous Step

Select the route tables that should include routes to the floating IP addresses. This enables client access to the Cloud Volumes ONTAP HA pair. If you leave a route table unselected, clients that are associated with the route table cannot access the HA pair.

Additional information ⓘ


<input checked="" type="checkbox"/>	Name	Main	ID	Associate with Subnet	Tags
<input checked="" type="checkbox"/>	private_rt_rt1600680	No	rtb-08b42b88f5c826a5	3 Subnets	1 Tags
<input checked="" type="checkbox"/>	public_rt_rt1600680	Yes	rtb-0e46720d0da10c593	1 Subnets	1 Tags

2 Route Tables | The main route table is the default for the VPC


Continue

Cloud Manager 3.9.9 Build 0 Aug 18, 2021 06:13:35 am UTC


3. Choose whether to enable AWS managed encryption or AWS KMS to encrypt the ONTAP root, boot, and data disks.



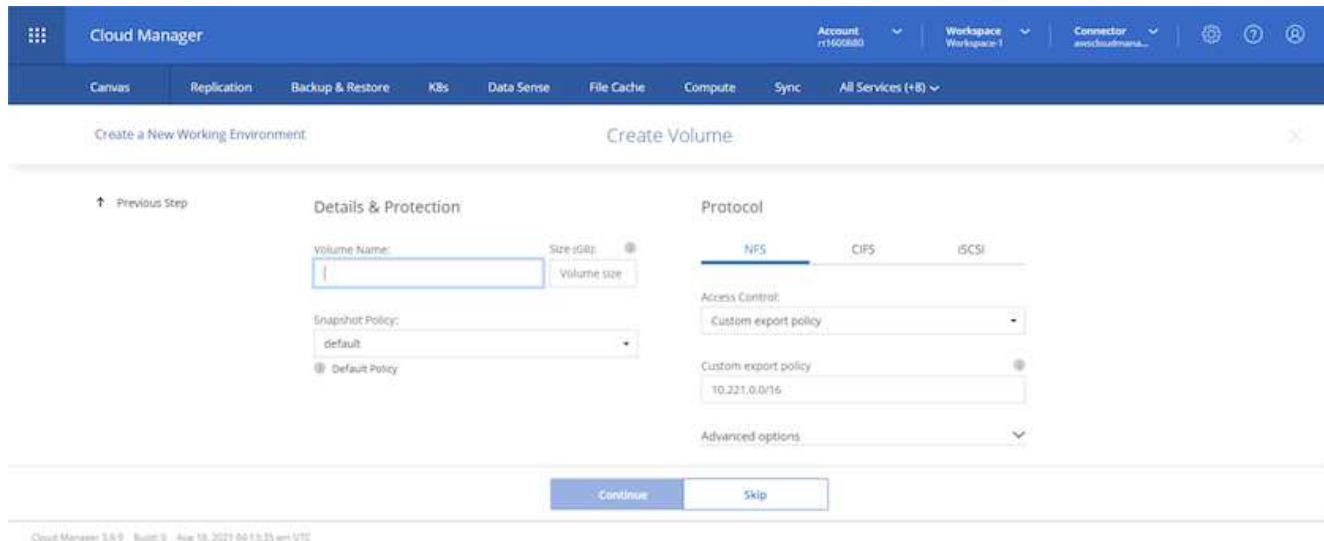
4. Choose your licensing model. If you don't know which to choose, contact your NetApp representative.



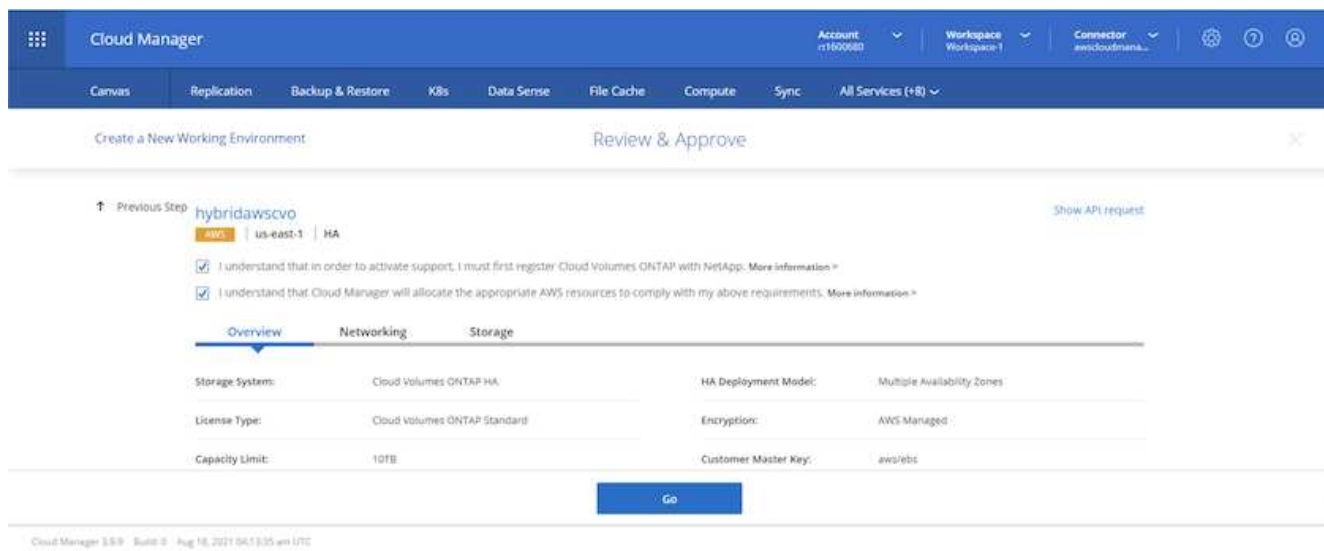
5. Select which configuration best suits your use case. This is related to the sizing considerations covered in the prerequisites page.




- Optionally, create a volume. This is not required, because the next steps use SnapMirror, which creates the volumes for us.



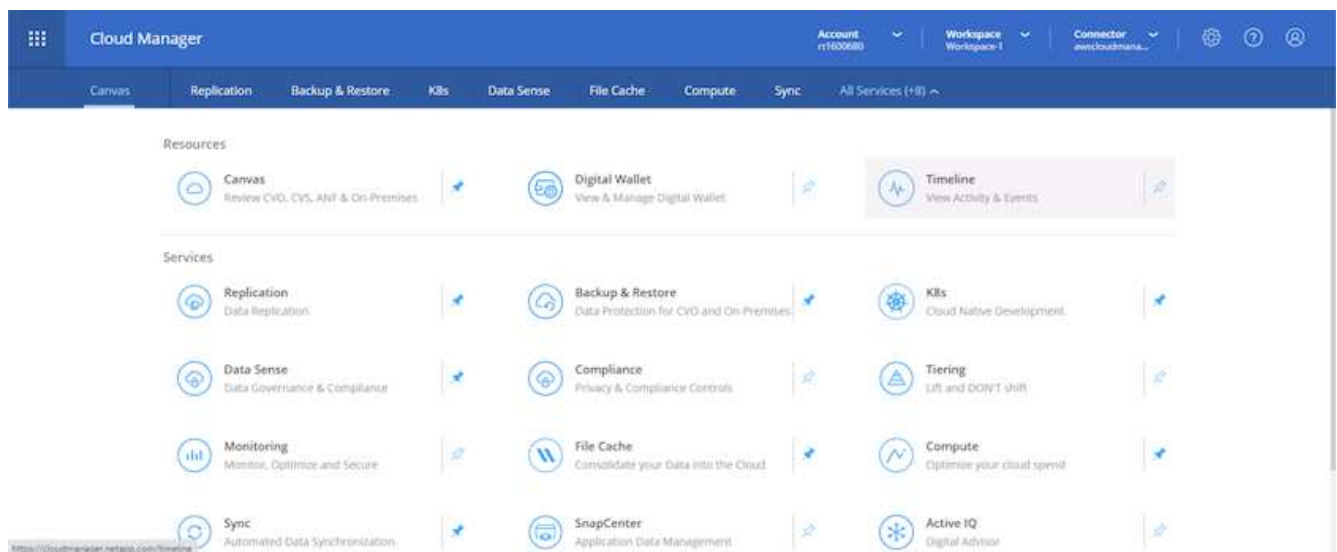
- Review the selections made and tick the boxes to verify that you understand that Cloud Manager deploys resources into your AWS environment. When ready, click Go.




- Cloud Volumes ONTAP now starts its deployment process. Cloud Manager uses AWS APIs and cloud formation stacks to deploy Cloud Volumes ONTAP. It then configures the system to your specifications, giving you a ready-to-go system that can be instantly utilized. The timing for this process varies depending on the selections made.



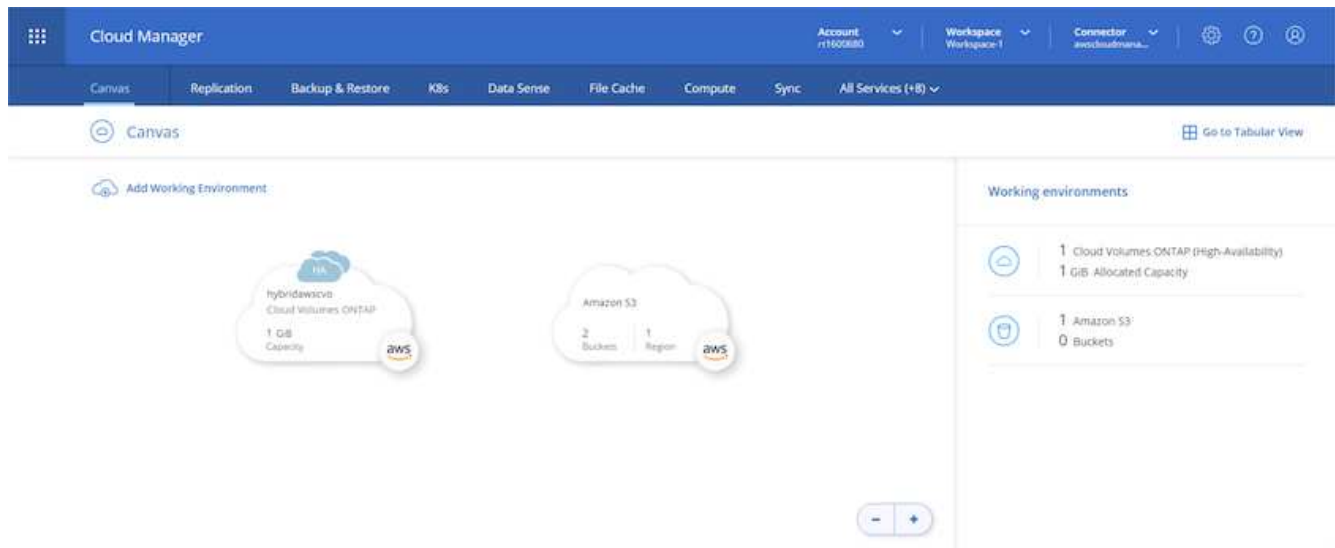
9. You can monitor the progress by navigating to the Timeline.



10. The Timeline acts as an audit of all actions performed in Cloud Manager. You can view all of the API calls that are made by Cloud Manager during setup to both AWS as well as the ONTAP cluster. This can also be effectively used to troubleshoot any issues that you face.



11. After deployment is complete, the CVO cluster appears on the Canvas, which the current capacity. The ONTAP cluster in its current state is fully configured to allow a true, out-of-the-box experience.




Configure SnapMirror from on-premises to cloud

Now that you have a source ONTAP system and a destination ONTAP system deployed, you can replicate volumes containing database data into the cloud.

For a guide on compatible ONTAP versions for SnapMirror, see the [SnapMirror Compatibility Matrix](#).

1. Click the source ONTAP system (on-premises) and either drag and drop it to the destination, select Replication > Enable, or select Replication > Menu > Replicate.

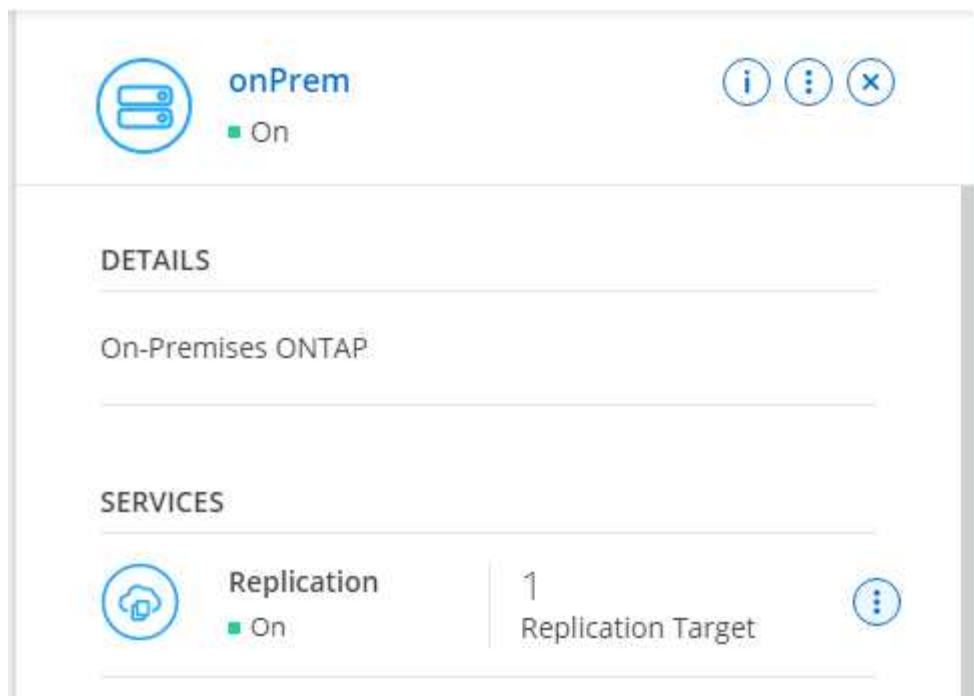


Select Enable.


SERVICES

	Replication ■ Off	<input type="button" value="Enable"/>	
---	-----------------------------	---------------------------------------	---


Or Options.




Replicate.




2. If you did not drag and drop, choose the destination cluster to replicate to.



3. Choose the volume that you'd like to replicate. We replicated the data and all log volumes.



4. Choose the destination disk type and tiering policy. For disaster recovery, we recommend an SSD as the disk type and to maintain data tiering. Data tiering tiers the mirrored data into low-cost object storage and saves you money on local disks. When you break the relationship or clone the volume, the data uses the fast, local storage.



5. Select the destination volume name: we chose `[source_volume_name]_dr`.



6. Select the maximum transfer rate for the replication. This enables you to save bandwidth if you have a low bandwidth connection to the cloud such as a VPN.

Max Transfer Rate


You should limit the transfer rate. An unlimited rate might negatively impact the performance of other applications and it might impact your Internet performance.

- Limited to: MB/s
- Unlimited (recommended for DR only machines)

7. Define the replication policy. We chose a Mirror, which takes the most recent dataset and replicates that into the destination volume. You could also choose a different policy based on your requirements.


Replication Policy

Default Policies Additional Policies

 **Mirror**

Typically used for disaster recovery

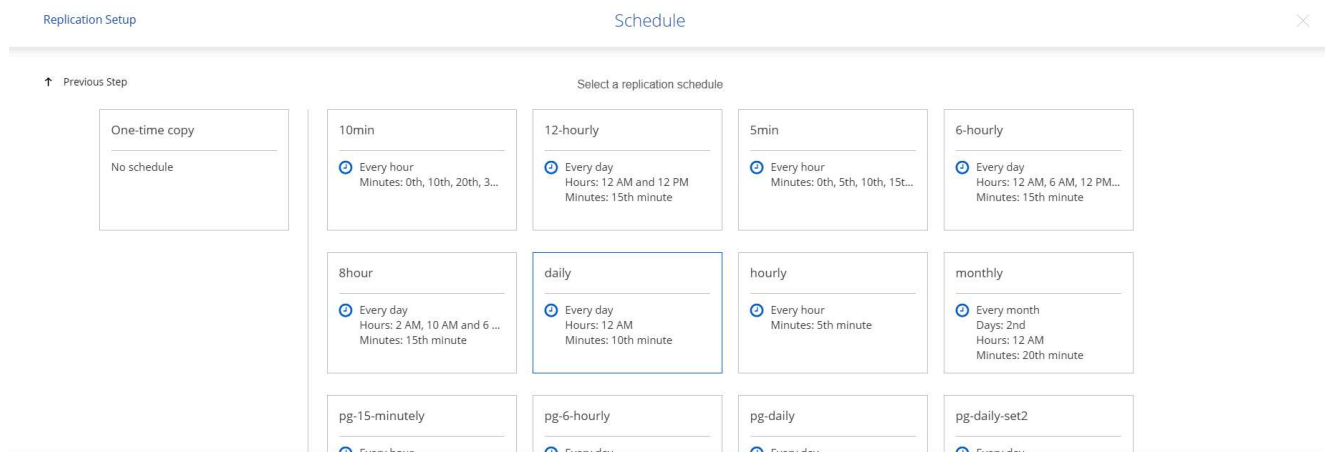
[More info](#)

 **Mirror and Backup (1 month retention)**

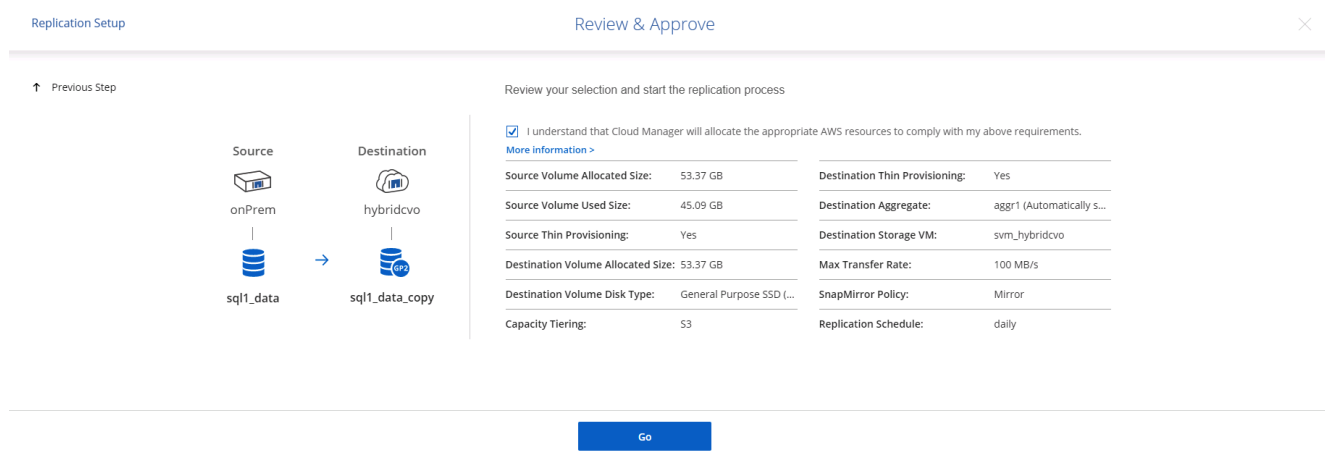
Configures disaster recovery and long-term retention of backups on the same destination volume

[More info](#)

8. Choose the schedule for triggering replication. NetApp recommends setting a "daily" schedule of for the data volume and an "hourly" schedule for the log volumes, although this can be changed based on requirements.

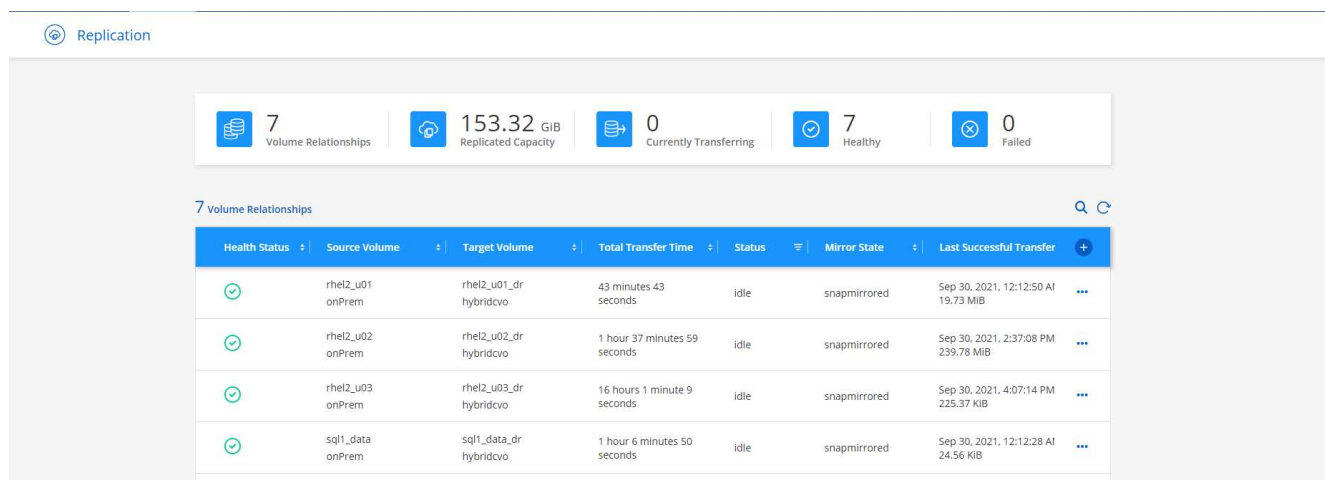


9. Review the information entered, click Go to trigger the cluster peer and SVM peer (if this is your first time replicating between the two clusters), and then implement and initialize the SnapMirror relationship.



10. Continue this process for data volumes and log volumes.

11. To check all of your relationships, navigate to the Replication tab inside Cloud Manager. Here you can manage your relationships and check on their status.



12. After all the volumes have been replicated, you are in a steady state and ready to move on to the disaster recovery and dev/test workflows.

3. Deploy EC2 compute instance for database workload

AWS has preconfigured EC2 compute instances for various workloads. The choice of instance type determines the number of CPU cores, memory capacity, storage type and capacity, and network performance. For the use cases, with the exception of the OS partition, the main storage to run database workload is allocated from CVO or the FSx ONTAP storage engine. Therefore, the main factors to consider are the choice of CPU cores, memory, and network performance level. Typical AWS EC2 instance types can be found here: [EC2 Instance Type](#).

Sizing the compute instance

1. Select the right instance type based on the required workload. Factors to consider include the number of business transactions to be supported, the number of concurrent users, data set sizing, and so on.
2. EC2 instance deployment can be launched through the EC2 Dashboard. The exact deployment procedures are beyond the scope of this solution. See [Amazon EC2](#) for details.

Linux instance configuration for Oracle workload

This section contain additional configuration steps after an EC2 Linux instance is deployed.

1. Add an Oracle standby instance to the DNS server for name resolution within the SnapCenter management domain.
2. Add a Linux management user ID as the SnapCenter OS credentials with sudo permissions without a password. Enable the ID with SSH password authentication on the EC2 instance. (By default, SSH password authentication and passwordless sudo is turned off on EC2 instances.)
3. Configure Oracle installation to match with on-premises Oracle installation such as OS patches, Oracle versions and patches, and so on.
4. NetApp Ansible DB automation roles can be leveraged to configure EC2 instances for database dev/test and disaster recovery use cases. The automation code can be download from the NetApp public GitHub site: [Oracle 19c Automated Deployment](#). The goal is to install and configure a database software stack on an EC2 instance to match on-premises OS and database configurations.

Windows instance configuration for SQL Server workload

This section lists additional configuration steps after an EC2 Windows instance is initially deployed.

1. Retrieve the Windows administrator password to log in to an instance via RDP.
2. Disable the Windows firewall, join the host to Windows SnapCenter domain, and add the instance to the DNS server for name resolution.
3. Provision a SnapCenter log volume to store SQL Server log files.
4. Configure iSCSI on the Windows host to mount the volume and format the disk drive.
5. Again, many of the previous tasks can be automated with the NetApp automation solution for SQL Server. Check the NetApp automation public GitHub site for newly published roles and solutions: [NetApp Automation](#).

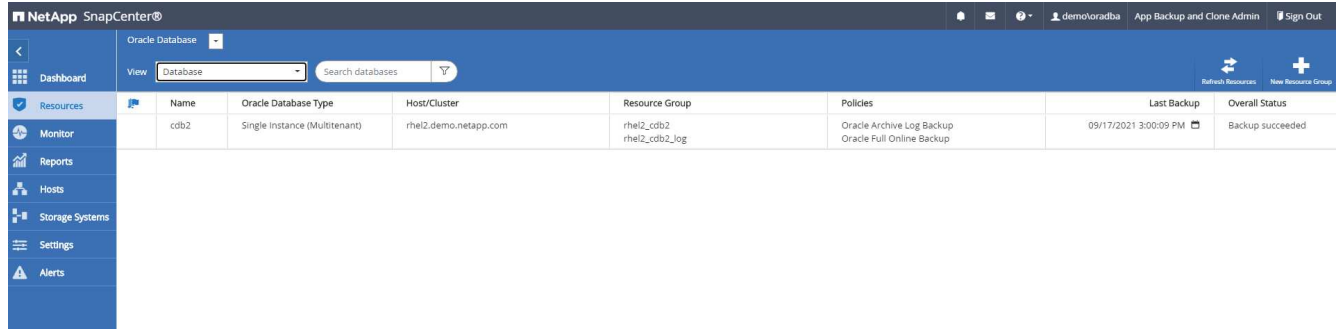
Workflow for dev/test bursting to cloud

The agility of the public cloud, the time to value, and the cost savings are all meaningful value propositions for enterprises adopting the public cloud for database application development and testing effort. There is no better tool than SnapCenter to make this a

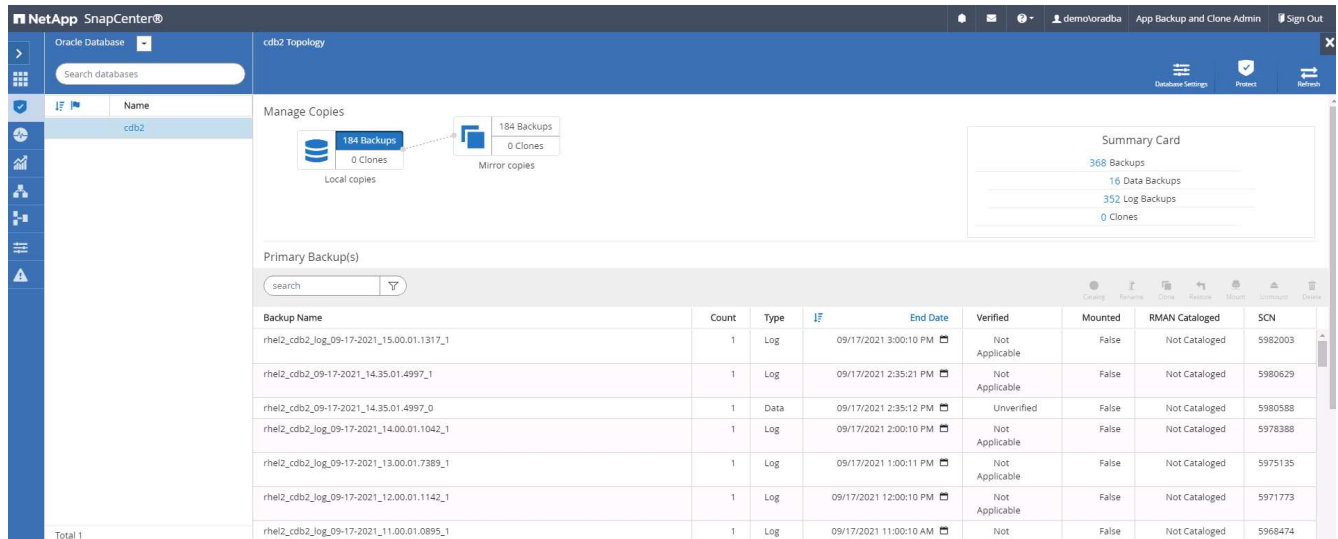
reality. SnapCenter can not only protect your production database on-premises, but can also it quickly clone a copy for application development or code testing in the public cloud while consuming very little extra storage. Following are details of the step-by-step processes for using this tool.

Clone an Oracle Database for dev/test from a replicated snapshot backup

1. Log into SnapCenter with a database management user ID for Oracle. Navigate to the Resources tab, which shows the Oracle databases being protected by SnapCenter.



2. Click the intended on-premises database name for the backup topology and the detailed view. If a secondary replicated location is enabled, it shows linked mirror backups.



3. Toggled to the mirrored backups view by clicking mirrored backups. The secondary mirror backup(s) is then displayed.

Backup Name	Count	Type	End Date	Verified	Mounted	RMAN Cataloged	SCN
rhel2_cdb2_log_09-17-2021_15.00.01.1317_1	1	Log	09/17/2021 3:00:10 PM	Not Applicable	False	Not Cataloged	5982003
rhel2_cdb2_09-17-2021_14.35.01.4997_1	1	Log	09/17/2021 2:35:21 PM	Not Applicable	False	Not Cataloged	5980629
rhel2_cdb2_09-17-2021_14.35.01.4997_0	1	Data	09/17/2021 2:35:12 PM	Unverified	False	Not Cataloged	5980588
rhel2_cdb2_log_09-17-2021_14.00.01.1042_1	1	Log	09/17/2021 2:00:10 PM	Not Applicable	False	Not Cataloged	5978388
rhel2_cdb2_log_09-17-2021_13.00.01.7389_1	1	Log	09/17/2021 1:00:11 PM	Not Applicable	False	Not Cataloged	5975135
rhel2_cdb2_log_09-17-2021_12.00.01.1142_1	1	Log	09/17/2021 12:00:10 PM	Not Applicable	False	Not Cataloged	5971773
rhel2_cdb2_log_09-17-2021_11.00.01.0895_1	1	Log	09/17/2021 11:00:10 AM	Not Applicable	False	Not Cataloged	5968474

- Choose a mirrored secondary database backup copy to be cloned and determine a recovery point either by time and system change number or by SCN. Generally, the recovery point should be trailing the full database backup time or SCN to be cloned. After a recovery point is decided, the required log file backup must be mounted for recovery. The log file backup should be mounted to target DB server where the clone database is to be hosted.

Mount backups

Choose the host to mount the backup:

Mount path: /var/opt/snapcenter/sco/backup_mount/rhel2_cdb2_09-17-2021_14.35.01.4997_1/cdb2

Secondary storage location : Snap Vault / Snap Mirror

Source Volume: svm_onPrem:rhel2_u03

Destination Volume:

Oracle Database | cdb2 Topology

Search databases

Manage Copies

184 Backups
0 Clones
Local copies

184 Backups
1 Clone
Mirror copies

Summary Card

368 Backups
16 Data Backups
352 Log Backups
1 Clone

Secondary Mirror Backup(s)

Backup Name	Count	Type	IF	End Date	Verified	Mounted	RMAN Cataloged	SCN
rhei2_cdb2_log_09-17-2021_16.00.01.2156_1	1	Log		09/17/2021 4:00:10 PM	Not Applicable	False	Not Cataloged	5985272
rhei2_cdb2_log_09-17-2021_15.00.01.1317_1	1	Log		09/17/2021 3:00:10 PM	Not Applicable	False	Not Cataloged	5982003
rhei2_cdb2_09-17-2021_14.35.01.4997_1	1	Log		09/17/2021 2:35:21 PM	Not Applicable	True	Not Cataloged	5980629
rhei2_cdb2_09-17-2021_14.35.01.4997_0	1	Data		09/17/2021 2:35:12 PM	Unverified	False	Not Cataloged	5980588
rhei2_cdb2_log_09-17-2021_14.00.01.1042_1	1	Log		09/17/2021 2:00:10 PM	Not Applicable	False	Not Cataloged	5978388



If log pruning is enabled and the recovery point is extended beyond the last log pruning, multiple archive log backups might need to be mounted.

- Highlight the full database backup copy to be cloned, and then click the clone button to start the DB clone Workflow.

cdb2 Topology

search

Clone

Backup Name	Count	Type	IF	End Date	Verified	Mounted	RMAN Cataloged	SCN
rhei2_cdb2_log_09-17-2021_16.00.01.2156_1	1	Log		09/17/2021 4:00:10 PM	Not Applicable	False	Not Cataloged	5985272
rhei2_cdb2_log_09-17-2021_15.00.01.1317_1	1	Log		09/17/2021 3:00:10 PM	Not Applicable	False	Not Cataloged	5982003
rhei2_cdb2_09-17-2021_14.35.01.4997_1	1	Log		09/17/2021 2:35:21 PM	Not Applicable	True	Not Cataloged	5980629
rhei2_cdb2_09-17-2021_14.35.01.4997_0	1	Data		09/17/2021 2:35:12 PM	Unverified	False	Not Cataloged	5980588
rhei2_cdb2_log_09-17-2021_14.00.01.1042_1	1	Log		09/17/2021 2:00:10 PM	Not Applicable	False	Not Cataloged	5978388

- Choose a proper clone DB SID for a complete container database or CDB clone.

Clone from cdb2
✕

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps
- 6 Notification
- 7 Summary

Complete Database Clone

Clone SID

Exclude PDBs

PDB Clone

Secondary storage location : Snap Vault / Snap Mirror

Data

Source Volume	Destination Volume
svm_onPrem:rhel2_u02	<input style="width: 90%;" type="text" value="svm_hybridcvo:rhel2_u02_dr"/>

Logs

Source Volume	Destination Volume
svm_onPrem:rhel2_u03	<input style="width: 90%;" type="text" value="svm_hybridcvo:rhel2_u03_dr"/>

7. Select the target clone host in the cloud, and datafile, control file, and redo log directories are created by the clone workflow.

Clone from cdb2
✕

1

2

3

4

5

6

7

Select the host to create a clone

Clone host

Datafile locations ⓘ

Reset

Control files ⓘ

<input type="text" value="/u02_cdb2test/cdb2test/control/control01.ctl"/>	✕		+
<input type="text" value="/u02_cdb2test/cdb2test/control/control02.ctl"/>	✕		Reset

Redo logs ⓘ

Group	Size	Unit	Number of files		
<input checked="" type="checkbox"/> RedoGroup 1 <input type="text" value="/u02_cdb2test/cdb2test/redolog/redolog03.log"/>	200	MB	1	✕	+
<input checked="" type="checkbox"/> RedoGroup 2	200	MB	1	✕	+

+
Reset

Previous
Next

8. The None credential name is used for OS-based authentication, which renders the database port irrelevant. Fill in the proper Oracle Home, Oracle OS User, and Oracle OS Group as configured in the target clone DB server.

Clone from cdb2 x

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps
- 6 Notification
- 7 Summary

Database Credentials for the clone

Credential name for sys user + ⓘ

Database port

Oracle Home Settings ⓘ

Oracle Home

Oracle OS User

Oracle OS Group

9. Specify the scripts to run before clone operation. More importantly, the database instance parameter can be adjusted or defined here.

Clone from cdb2
✕

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps
- 6 Notification
- 7 Summary

Specify scripts to run before clone operation ?

Prescript full path

Arguments

Script timeout secs

⊖ Database Parameter settings

processes	320	✕	▲
remote_login_passwordfile	EXCLUSIVE	✕	+
sga_target	4311744512	✕	▼
undo_tablespace	UNDOTBS1	✕	

10. Specify the recovery point either by the date and time or SCN. Until Cancel recovers the database up to the available archive logs. Specify the external archive log location from the target host where the archive log volume is mounted. If target server Oracle owner is different from the on-premises production server, verify that the archive log directory is readable by the target server Oracle owner.

Clone from cdb2

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps**
- 6 Notification
- 7 Summary

Recover Database

Until Cancel i
 Date and Time i
 Date-time format: MM/DD/YYYY hh:mm:ss
 Until SCN (System Change Number) i

Specify external archive log locations i

Create new DBID i
 Create tempfile for temporary tablespace i
 Enter SQL queries to apply when clone is created
 Enter scripts to run after clone operation i

```

oracle@ora-standby/tmp
[oracle@ora-standby tmp]$ ls /var/opt/snapcenter/sco/backup_mount/rhel2_cdb2_09-17-2021_14.35.01.4997_1/cdb2/1/orareco/CDB2/archivelog/
2021_08_26 2021_08_28 2021_08_30 2021_09_01 2021_09_03 2021_09_05 2021_09_07 2021_09_09 2021_09_11 2021_09_13 2021_09_15 2021_09_17
2021_08_27 2021_08_29 2021_08_31 2021_09_02 2021_09_04 2021_09_06 2021_09_08 2021_09_10 2021_09_12 2021_09_14 2021_09_16
[oracle@ora-standby tmp]$
  
```

11. Configure the SMTP server for email notification if desired.

Clone from cdb2

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps
- 6 Notification**
- 7 Summary

Provide email settings ?

Email preference:

From:

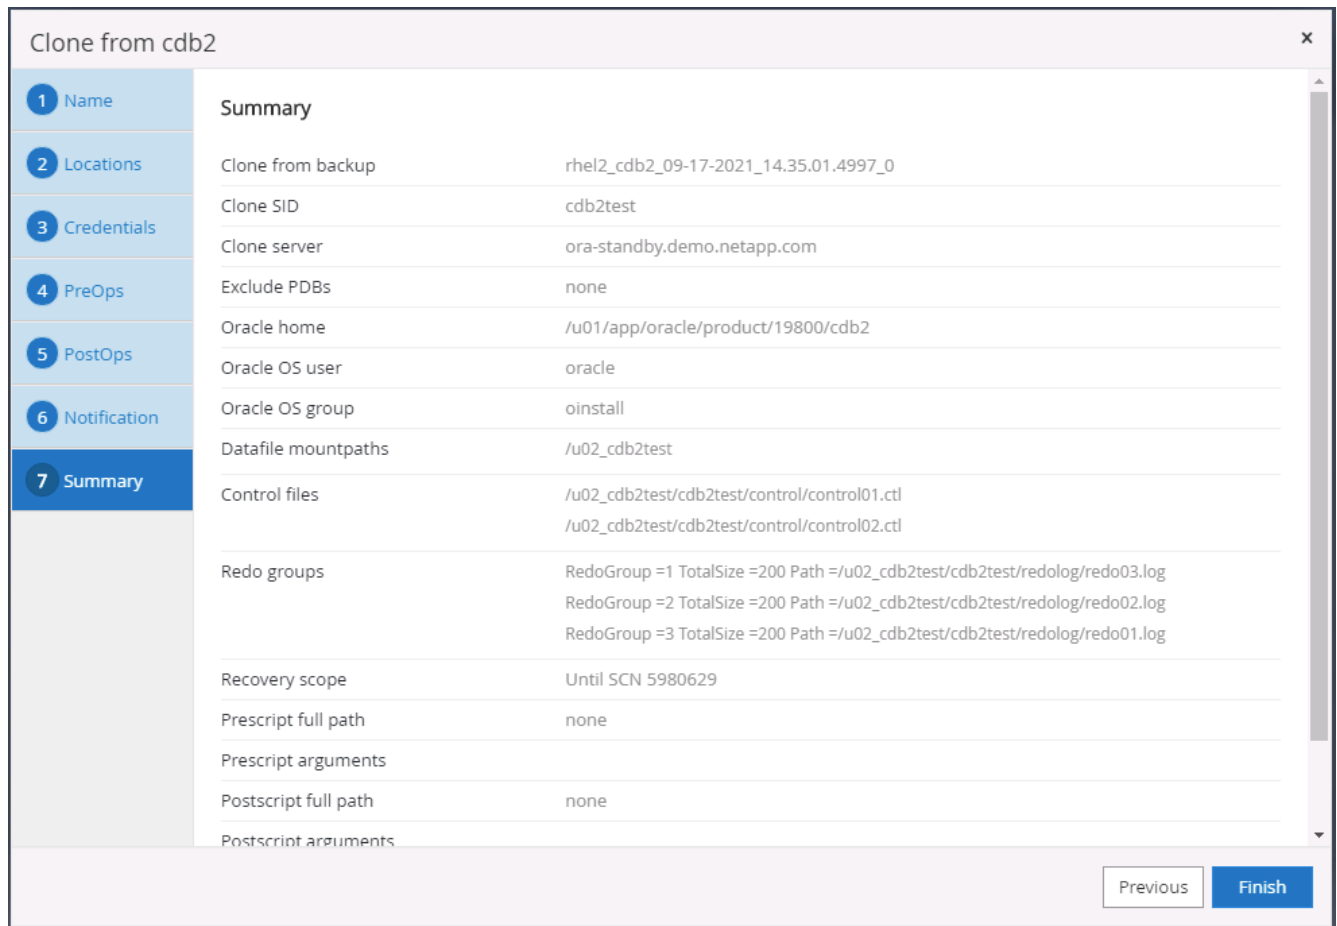
To:

Subject:

Attach job report

⚠ If you want to send notifications for Clone jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server.

12. Clone summary.



13. You should validate after cloning to make sure that the cloned database is operational. Some additional tasks, such as starting up the listener or turning off the DB log archive mode, can be performed on the dev/test database.

```

oracle@ora-standby/tmp
[oracle@ora-standby tmp]$ export ORACLE_SID=cdb2test
[oracle@ora-standby tmp]$ export ORACLE_HOME=/u01/app/oracle/product/19800/cdb2
[oracle@ora-standby tmp]$ export PATH=$PATH:$ORACLE_HOME/bin
[oracle@ora-standby tmp]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Fri Sep 17 17:49:29 2021
Version 19.3.0.0.0

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

Connected to:
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
Version 19.3.0.0.0

SQL> select name, log_mode from v$database;

NAME          LOG_MODE
-----
CDB2TEST     ARCHIVELOG

SQL> select instance_name, host_name from v$instance;

INSTANCE_NAME
HOST_NAME
-----
cdb2test
ora-standby.demo.netapp.com

SQL> show pdbs

CON_ID CON_NAME          OPEN MODE  RESTRICTED
-----
2 PDB$SEED          READ ONLY NO
3 CDB2_PDB1         READ WRITE NO
4 CDB2_PDB2         READ WRITE NO
5 CDB2_PDB3         READ WRITE NO

```

Clone a SQL database for dev/test from a replicated Snapshot backup

1. Log into SnapCenter with a database management user ID for SQL Server. Navigate to the Resources tab, which shows the SQL Sever user databases being protected by SnapCenter and a target standby SQL instance in the public cloud.

Name	Instance	Host	Last Backup	Overall Status	Type
master	sql1	sql1.demo.netapp.com		Not available for backup	System database
model	sql1	sql1.demo.netapp.com		Not available for backup	System database
msdb	sql1	sql1.demo.netapp.com		Not available for backup	System database
tempdb	sql1	sql1.demo.netapp.com		Not available for backup	System database
tpcc	sql1	sql1.demo.netapp.com	09/16/2021 7:35:05 PM	Backup succeeded	User database
master	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database
model	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database
msdb	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database
tempdb	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database

2. Click on the intended on-premises SQL Server user database name for the backups topology and detailed view. If a secondary replicated location is enabled, it shows linked mirror backups.

Backup Name	Count	Type	End Date	Verified
sql1_tpcc_09-16-2021_18.25.01.4024	1	Full backup	09/16/2021 6:25:05 PM	Unverified
sql1_tpcc_09-15-2021_18.25.01.4604	1	Full backup	09/15/2021 6:25:06 PM	Unverified
sql1_tpcc_09-14-2021_18.25.01.5233	1	Full backup	09/14/2021 6:25:05 PM	Unverified
sql1_tpcc_09-13-2021_18.25.01.4500	1	Full backup	09/13/2021 6:25:05 PM	Unverified
sql1_tpcc_09-12-2021_18.25.01.4016	1	Full backup	09/12/2021 6:25:05 PM	Unverified
sql1_tpcc_09-11-2021_18.25.01.3753	1	Full backup	09/11/2021 6:25:05 PM	Unverified
sql1_tpcc_09-10-2021_18.36.25.5430	1	Full backup	09/10/2021 6:36:29 PM	Unverified

3. Toggle to the Mirrored Backups view by clicking Mirrored Backups. Secondary Mirror Backup(s) are then displayed. Because SnapCenter backs up the SQL Server transaction log to a dedicated drive for recovery, only full database backups are displayed here.

Backup Name	Count	Type	End Date	Verified
sql1_tpcc_09-16-2021_18.25.01.4024	1	Full backup	09/16/2021 6:25:05 PM	Unverified
sql1_tpcc_09-15-2021_18.25.01.4604	1	Full backup	09/15/2021 6:25:06 PM	Unverified
sql1_tpcc_09-14-2021_18.25.01.5233	1	Full backup	09/14/2021 6:25:05 PM	Unverified
sql1_tpcc_09-13-2021_18.25.01.4500	1	Full backup	09/13/2021 6:25:05 PM	Unverified
sql1_tpcc_09-12-2021_18.25.01.4016	1	Full backup	09/12/2021 6:25:05 PM	Unverified
sql1_tpcc_09-11-2021_18.25.01.3753	1	Full backup	09/11/2021 6:25:05 PM	Unverified
sql1_tpcc_09-10-2021_18.36.25.5430	1	Full backup	09/10/2021 6:36:29 PM	Unverified

4. Choose a backup copy, and then click the Clone button to launch the Clone from Backup workflow.

Backup Name	Count	Type	IF	End Date	Verified
sql1_tpcc_09-19-2021_18.25.01.4134	1	Full backup		09/19/2021 6:25:05 PM	Unverified
sql1_tpcc_09-18-2021_18.25.01.3963	1	Full backup		09/18/2021 6:25:05 PM	Unverified
sql1_tpcc_09-17-2021_18.25.01.4218	1	Full backup		09/17/2021 6:25:05 PM	Unverified
sql1_tpcc_09-16-2021_18.25.01.4024	1	Full backup		09/16/2021 6:25:05 PM	Unverified
sql1_tpcc_09-15-2021_18.25.01.4604	1	Full backup		09/15/2021 6:25:06 PM	Unverified
sql1_tpcc_09-14-2021_18.25.01.5233	1	Full backup		09/14/2021 6:25:05 PM	Unverified
sql1_tpcc_09-13-2021_18.25.01.4500	1	Full backup		09/13/2021 6:25:05 PM	Unverified

Clone from backup

1 Clone Options

Clone settings

Clone server: Choose

Clone instance: Nothing selected

Clone name: tpcc

Choose mount option

Auto assign mount point

Auto assign volume mount point under path: full file path

Secondary storage location : Snap Vault / Snap Mirror

Source Volume	Destination Volume
svm_onPrem:sql1_data	svm_hybridcvo:sql1_data_dr
svm_onPrem:sql1_log	svm_hybridcvo:sql1_log_dr

Previous Next

5. Select a cloud server as the target clone server, clone instance name, and clone database name. Choose either an auto-assign mount point or a user-defined mount point path.

×
Clone from backup

1 Clone Options

2 Logs

3 Script

4 Notification

5 Summary

Clone settings

Clone server

sql-standby.demo.netapp.com

i

Clone instance

sql-standby

i

Clone name

tpcc_clone

Choose mount option

Auto assign mount point i

Auto assign volume mount point under path

full file path

i

Secondary storage location : Snap Vault / Snap Mirror

Source Volume	Destination Volume
svm_onPrem:sql1_data	<div style="border: 1px solid #ccc; padding: 2px 5px; border-radius: 3px;">svm_hybridcvo:sql1_data_dr</div>
svm_onPrem:sql1_log	<div style="border: 1px solid #ccc; padding: 2px 5px; border-radius: 3px;">svm_hybridcvo:sql1_log_dr</div>

Previous

Next

6. Determine a recovery point either by a log backup time or by a specific date and time.

Clone from backup x

1 Clone Options

2 Logs

3 Script

4 Notification

5 Summary

Choose logs

All log backups

By log backups until

By specific date until

None

7. Specify optional scripts to run before and after the cloning operation.

Clone from backup

- 1 Clone Options
- 2 Logs
- 3 Script**
- 4 Notification
- 5 Summary

Specify optional scripts to run before and after performing a clone from backup job

Prescript full path

Prescript arguments

Postscript full path

Postscript arguments

Script timeout

8. Configure an SMTP server if email notification is desired.

Clone from backup ✕

- 1 Clone Options
- 2 Logs
- 3 Script
- 4 Notification**
- 5 Summary

Provide email settings ?

Email preference

From


To

Subject

Attach Job Report

⚠ If you want to send notifications for Clone jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server. ✕

9. Clone Summary.



10. Monitor the job status and validate that the intended user database has been attached to a target SQL instance in the cloud clone server.

ID	Status	Name	Start date	End date	Owner
766	✓	Clone from backup 'sql1_tpcc_09-16-2021_18.25.01.4024'	09/16/2021 8:05:25 PM	09/16/2021 8:06:17 PM	demo\$sqldba
763	✓	Discover resources for all hosts	09/16/2021 7:56:49 PM	09/16/2021 7:56:54 PM	demo\$sqldba
761	✓	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/16/2021 7:35:00 PM	09/16/2021 7:37:08 PM	demo\$sqldba
760	⚠	Discover resources for all hosts	09/16/2021 7:19:05 PM	09/16/2021 7:19:09 PM	demo\$sqldba
759	⚠	Discover resources for all hosts	09/16/2021 7:18:43 PM	09/16/2021 7:18:48 PM	demo\$sqldba
756	⚠	Discover resources for all hosts	09/16/2021 6:59:51 PM	09/16/2021 6:59:56 PM	demo\$sqldba
753	✓	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/16/2021 6:35:00 PM	09/16/2021 6:37:07 PM	demo\$sqldba
750	✓	Backup of Resource Group 'sql1_tpcc' with policy 'SQL Server Full Backup'	09/16/2021 6:25:01 PM	09/16/2021 6:27:14 PM	demo\$sqldba
749	✓	Discover resources for host 'sql-standby.demo.netapp.com'	09/16/2021 6:19:00 PM	09/16/2021 6:19:05 PM	DemoAdministrator
745	✓	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/16/2021 5:35:00 PM	09/16/2021 5:37:08 PM	demo\$sqldba

Post-clone configuration

1. An Oracle production database on-premises is usually running in log archive mode. This mode is not necessary for a development or test database. To turn off log archive mode, log into the Oracle DB as sysdba, execute a log mode change command, and start the database for access.
2. Configure an Oracle listener, or register the newly cloned DB with an existing listener for user access.
3. For SQL Server, change the log mode from Full to Easy so that the SQL Server dev/test log file can be readily shrunk when it is filling up the log volume.

Refresh clone database

1. Drop cloned databases and clean up the cloud DB server environment. Then follow the previous procedures to clone a new DB with fresh data. It only takes few minutes to clone a new database.
2. Shutdown the clone database, run a clone refresh command by using the CLI. See the following SnapCenter documentation for details: [Refresh a clone](#).

Where to go for help?

If you need help with this solution and use cases, join the [NetApp Solution Automation community support Slack channel](#) and look for the solution-automation channel to post your questions or inquires.

Disaster recovery workflow

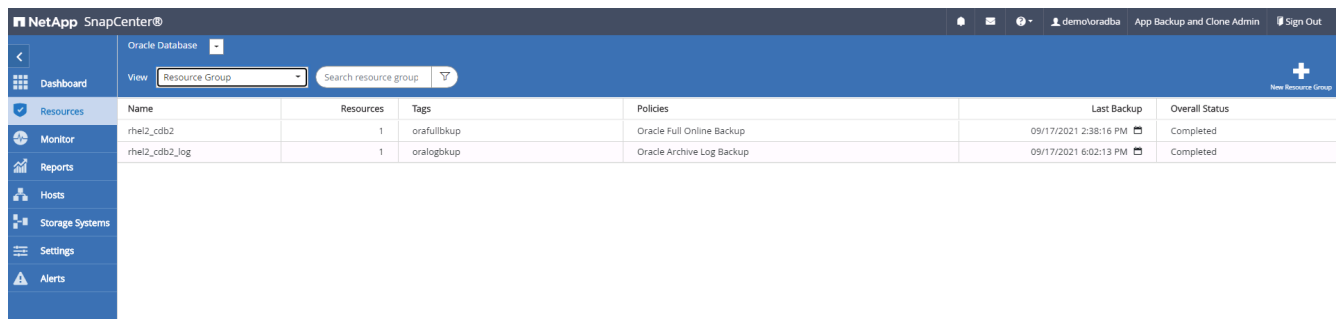
Enterprises have embraced the public cloud as a viable resource and destination for disaster recovery. SnapCenter makes this process as seamless as possible. This disaster recovery workflow is very similar to the clone workflow, but database recovery runs through the last available log that was replicated to cloud to recover all the business transactions possible. However, there are additional pre-configuration and post-configuration steps specific to disaster recovery.

Clone an on-premises Oracle production DB to cloud for DR



1. To validate that the clone recovery runs through last available log, we created a small test table and inserted a row. The test data would be recovered after a full recovery to last available log.

```
oracle@rhel2~$
SQL> create table dr_test(
  2 id integer,
  3 event varchar(200),
  4 dt timestamp);
Table created.
SQL> insert into dr_test values(1, 'testing DB clone for DR and roll forward DB to last available log', sysdate);
1 row created.
SQL> select * from dr_test;
-----
ID
-----
EVENT
-----
DT
-----
1
testing DB clone for DR and roll forward DB to last available log
17-SEP-21 02.12.13.000000 PM
SQL> commit;
Commit complete.
SQL>
```

2. Log into SnapCenter as a database management user ID for Oracle. Navigate to the Resources tab, which shows the Oracle databases being protected by SnapCenter.



3. Select the Oracle log resource group and click Backup Now to manually run an Oracle log backup to flush the latest transaction to the destination in the cloud. In a real DR scenario, the last transaction recoverable depends on the database log volume replication frequency to the cloud, which in turn depends on the RTO or RPO policy of the company.



Asynchronous SnapMirror loses data that has not made it to the cloud destination in the database log backup interval in a disaster recovery scenario. To minimize data loss, more frequent log backup can be scheduled. However there is a limit to the log backup frequency that is technically achievable.

4. Select the last log backup on the Secondary Mirror Backup(s), and mount the log backup.

NetApp SnapCenter® Oracle Database cdb2 Topology

Search databases

Manage Copies

Local copies: 185 Backups, 0 Clones

Mirror copies: 185 Backups, 2 Clones

Summary Card

- 370 Backups
- 16 Data Backups
- 354 Log Backups
- 2 Clones

Secondary Mirror Backup(s)

Backup Name	Count	Type	IF	End Date	Verified	Mounted	RMAN Cataloged	SCN
rhel2_cdb2_log_09-17-2021_18.20.04.1177_1	1	Log		09/17/2021 6:20:13 PM	Not Applicable	False	Not Cataloged	5994710
rhel2_cdb2_log_09-17-2021_18.00.01.2424_1	1	Log		09/17/2021 6:00:09 PM	Not Applicable	False	Not Cataloged	5992079
rhel2_cdb2_log_09-17-2021_17.00.01.1566_1	1	Log		09/17/2021 5:00:20 PM	Not Applicable	False	Not Cataloged	5988842

Mount backups

Choose the host to mount the backup:

Mount path: /var/opt/snapcenter/sco/backup_mount/rhel2_cdb2_log_09-17-2021_18.20.04.1177_1/cdb2

Secondary storage location : Snap Vault / Snap Mirror

Source Volume: svm_onPrem:rhel2_u03

Destination Volume:

5. Select the last full database backup and click Clone to initiate the clone workflow.

NetApp SnapCenter® Oracle Database cdb2 Topology

Search databases

Manage Copies

Local copies: 185 Backups, 0 Clones

Mirror copies: 185 Backups, 2 Clones

Summary Card

- 370 Backups
- 16 Data Backups
- 354 Log Backups
- 2 Clones

Secondary Mirror Backup(s)

Backup Name	Count	Type	IF	End Date	Verified	Mounted	RMAN Cataloged	SCN
rhel2_cdb2_log_09-17-2021_18.20.04.1177_1	1	Log		09/17/2021 6:20:13 PM	Not Applicable	True	Not Cataloged	5994710
rhel2_cdb2_log_09-17-2021_18.00.01.2424_1	1	Log		09/17/2021 6:00:09 PM	Not Applicable	False	Not Cataloged	5992079
rhel2_cdb2_log_09-17-2021_17.00.01.1566_1	1	Log		09/17/2021 5:00:20 PM	Not Applicable	False	Not Cataloged	5988842
rhel2_cdb2_log_09-17-2021_16.00.01.2156_1	1	Log		09/17/2021 4:00:10 PM	Not Applicable	False	Not Cataloged	5985272
rhel2_cdb2_log_09-17-2021_15.00.01.1317_1	1	Log		09/17/2021 3:00:10 PM	Not Applicable	False	Not Cataloged	5982003
rhel2_cdb2_09-17-2021_14.35.01.4997_1	1	Log		09/17/2021 2:35:21 PM	Not Applicable	False	Not Cataloged	5980629
rhel2_cdb2_09-17-2021_14.35.01.4997_0	1	Data		09/17/2021 2:35:12 PM	Unverified	False	Not Cataloged	5980588

Total 3

6. Select a unique clone DB ID on the host.

Clone from cdb2

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Complete Database Clone

Clone SID:

Exclude PDBs:

PDB Clone

Secondary storage location : Snap Vault / Snap Mirror

Data

Source Volume: svm_onPrem:rhel2_u02

Destination Volume:

Logs

Source Volume: svm_onPrem:rhel2_u03

Destination Volume:

Previous Next

7. Provision a log volume and mount it to the target DR server for the Oracle flash recovery area and online logs.

ONTAP System Manager

Search actions, objects, and pages

DASHBOARD

STORAGE

Overview

Applications

Volumes

LUNs

Shares

Qtrees

Quotas

Storage VMs

Tiers

NETWORK

EVENTS & JOBS

PROTECTION

HOSTS

Volumes

+ Add More

Name	Storage VM	Status	Capacity
ora_standby_u01	svm_hybridcvo	Online	12.3 GB used / 17.7 GB available / 31.6 GB
rhel2_u01_dr	svm_hybridcvo	Online	
rhel2_u02_dr	svm_hybridcvo	Online	
rhel2_u02_dr0917211608119360	svm_hybridcvo	Online	
rhel2_u02_dr0917211703534863	svm_hybridcvo	Online	
rhel2_u03_dr	svm_hybridcvo	Online	
rhel2_u03_dr0917211824574775	svm_hybridcvo	Online	

Add Volume

NAME:

CAPACITY:

More Options Cancel Save


```

ec2-user@ora-standby/tmp
[ec2-user@ora-standby tmp]$ sudo mkdir /u03_cdb2dr
[ec2-user@ora-standby tmp]$ chown oracle:oinstall /u03_cdb2dr
chown: changing ownership of '/u03_cdb2dr': Operation not permitted
[ec2-user@ora-standby tmp]$ sudo chown oracle:oinstall /u03_cdb2dr
[ec2-user@ora-standby tmp]$ sudo mount -t nfs 10.221.1.6:/ora_standby_u03 /u03_cdb2dr
[ec2-user@ora-standby tmp]$ df -h
Filesystem                Size      Used Avail Use% Mounted on
devtmpfs                  7.6G         0   7.6G   0% /dev
tmpfs                     7.6G         0   7.6G   0% /dev/shm
tmpfs                     7.6G        17M   7.6G   1% /run
tmpfs                     7.6G         0   7.6G   0% /sys/fs/cgroup
/dev/nvme0n1p2            10G         9.0G   1.1G  90% /
10.221.1.6:/ora_standby_u01 21G        13G   8.6G  42% /u01
tmpfs                     1.6G         0   1.6G   0% /run/user/1000
10.221.1.6:/Sc28182452-3fa8-448c-9e4a-c5a9e465f353 100G       3.1G   97G   4% /u02_cdb2dev
tmpfs                     1.6G         0   1.6G   0% /run/user/54921
10.221.1.6:/Sc39c05df8-4b00-4b3a-853c-9d6d338e5df7 100G       3.7G   97G   4% /u02_cdb2test
10.221.1.6:/Sccf886a5c-3273-475e-ad97-472b2a8dcccce 100G       3.8G   97G   4% /var/opt/snapcenter/sco/backup_mount/rhel12_cdb2_log_09-17-2021_18.20.04.1177_1/cdb2/1
10.221.1.6:/ora_standby_u03 21G        320K   20G   1% /u03_cdb2dr
[ec2-user@ora-standby tmp]$

```



The Oracle clone procedure does not create a log volume, which needs to be provisioned on the DR server before cloning.

8. Select the target clone host and location to place the data files, control files, and redo logs.

Clone from cdb2

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps
- 6 Notification
- 7 Summary

Select the host to create a clone

Clone host:

Datafile locations

Reset

Control files

✕

✕ Reset

Redo logs

Group	Size	Unit	Number of files
RedoGroup 1	200	MB	1
<input type="text" value="/u03_cdb2dr/cdb2dr/redolog/redo03.log"/>			
RedoGroup 2	200	MB	1

✕ + Reset

Previous
Next

9. Select the credentials for the clone. Fill in the details of the Oracle home configuration on the target server.

Clone from cdb2 x

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps
- 6 Notification
- 7 Summary

Database Credentials for the clone

Credential name for sys user: + ⓘ

Database port:

Oracle Home Settings ⓘ

Oracle Home:

Oracle OS User:

Oracle OS Group:

10. Specify the scripts to run before cloning. Database parameters can be adjusted if needed.

Clone from cdb2
✕

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps
- 6 Notification
- 7 Summary

Specify scripts to run before clone operation ❗

Prescript full path

Arguments

Script timeout secs

⊖ Database Parameter settings

audit_file_dest	/u01/app/oracle/admin/cdb2dr/adump	✕	<input style="width: 20px; height: 20px; border: 1px solid #ccc;" type="button" value="+"/> <input style="width: 40px; height: 20px; border: 1px solid #ccc;" type="button" value="Reset"/>
audit_trail	DB	✕	
open_cursors	300	✕	
pga_aggregate_target	1432354816	✕	

11. Select Until Cancel as the recovery option so that the recovery runs through all available archive logs to recoup the last transaction replicated to the secondary cloud location.

Clone from cdb2

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps**
- 6 Notification
- 7 Summary

Recover Database

Until Cancel ⓘ

Date and Time ⓘ
Date-time format: MM/DD/YYYY hh:mm:ss

Until SCN (System Change Number) ⓘ

Specify external archive log locations ⓘ ⓘ ⓘ

`/var/opt/snapcenter/sco/backup_mount/rhel2_cdb2_log_09-17-2021_18.20.04.1177_1/cdb2/1/orareco/CDB2/archivelog/`

Create new DBID ⓘ

Create tempfile for temporary tablespace ⓘ

Enter SQL queries to apply when clone is created

Enter scripts to run after clone operation ⓘ

Previous Next

12. Configure the SMTP server for email notification if needed.

Clone from cdb2

- 1 Name
- 2 Locations
- 3 Credentials
- 4 PreOps
- 5 PostOps
- 6 Notification**
- 7 Summary

Provide email settings ?

Email preference:

From:


To:

Subject:

Attach job report

⚠ If you want to send notifications for Clone jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server.

13. DR clone summary.



14. Cloned DBs are registered with SnapCenter immediately after clone completion and are then available for backup protection.

Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Overall Status
cdb2	Single Instance (Multitenant)	rhe12.demo.netapp.com	rhe12_cdb2 rhe12_cdb2_log	Oracle Archive Log Backup Oracle Full Online Backup	09/17/2021 7:00:10 PM	Backup succeeded
cdb2dev	Single Instance (Multitenant)	ora-standby.demo.netapp.com				Not protected
cdb2dr	Single Instance (Multitenant)	ora-standby.demo.netapp.com				Not protected
cdb2test	Single Instance (Multitenant)	ora-standby.demo.netapp.com				Not protected

Post DR clone validation and configuration for Oracle

1. Validate the last test transaction that has been flushed, replicated, and recovered at the DR location in the cloud.

```

oracle@ora-standby:/u01/app/oracle/product/19800/cdb2/dbs
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
Version 19.3.0.0.0

SQL> set lin 200
SQL> select instance_name, host_name from v$instance;

INSTANCE_NAME      HOST_NAME
-----
cdb2dr              ora-standby.demo.netapp.com

SQL> alter pluggable database cdb2_pdb1 open;

Pluggable database altered.

SQL> alter session set container=cdb2_pdb1;

Session altered.

SQL> select * from pdbadmin.dr_test;

      ID
-----
EVENT
-----
DT
-----
1
testing DB clone for DR and roll forward DB to last available log
17-SEP-21 02.12.13.000000 PM

SQL>

```

2. Configure the flash recovery area.

```

oracle@ora-standby:/u01/app/oracle/product/19800/cdb2/dbs
[oracle@ora-standby dbs]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Fri Sep 17 22:07:11 2021
Version 19.3.0.0.0

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

Connected to:
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
Version 19.3.0.0.0

SQL> show parameter db_recovery_file_dest

NAME                                 TYPE          VALUE
-----
db_recovery_file_dest                string        /u03_cdb2dr/cdb2dr
db_recovery_file_dest_size           big integer   17208M
SQL> alter system set db_recovery_file_dest='/u03_cdb2dr/cdb2dr' scope=both;

System altered.

SQL> show parameter db_recovery_file_dest

NAME                                 TYPE          VALUE
-----
db_recovery_file_dest                string        /u03_cdb2dr/cdb2dr
db_recovery_file_dest_size           big integer   17208M
SQL>

```

3. Configure the Oracle listener for user access.
4. Split the cloned volume off of the replicated source volume.
5. Reverse replication from the cloud to on-premises and rebuild the failed on-premises database server.



Clone split may incur temporary storage space utilization that is much higher than normal operation. However, after the on-premises DB server is rebuilt, extra space can be released.

Clone an on-premises SQL production DB to cloud for DR

1. Similarly, to validate that the SQL clone recovery ran through last available log, we created a small test table and inserted a row. The test data would be recovered after a full recovery to the last available log.

```


Administrator Command Prompt - sqlcmd - SQLCMD
C:\Users\administrator.DEMO>sqlcmd
1> select host_name()
2> go

-----
SQL1
(1 rows affected)
1> use tpcc
2> go
Changed database context to 'tpcc'.
1> insert into snap_sync values ('test snap mirror DR for SQL', getdate())
2> go


(1 rows affected)
1> select * from snap_sync
2> go
event                                                    dt
-----
test snap mirror DR for SQL                            2021-09-20 14:23:04.533
(1 rows affected)
1>

```


2. Log into SnapCenter with a database management user ID for SQL Server. Navigate to the Resources tab, which shows the SQL Server protection resources group.




3. Manually run a log backup to flush the last transaction to be replicated to secondary storage in the public cloud.



4. Select the last full SQL Server backup for the clone.



- Set the clone setting such as the Clone Server, Clone Instance, Clone Name, and mount option. The secondary storage location where cloning is performed is auto-populated.



- Select all log backups to be applied.

Clone from backup x

1 Clone Options

2 Logs

3 Script

4 Notification

5 Summary

Choose logs

All log backups

By log backups until

By specific date until

None

7. Specify any optional scripts to run before or after cloning.

Clone from backup

- 1 Clone Options
- 2 Logs
- 3 Script**
- 4 Notification
- 5 Summary

Specify optional scripts to run before and after performing a clone from backup job

Prescript full path

Prescript arguments

Postscript full path

Postscript arguments

Script timeout

8. Specify an SMTP server if email notification is desired.

Clone from backup

- 1 Clone Options
- 2 Logs
- 3 Script
- 4 Notification**
- 5 Summary

Provide email settings ?

Email preference:

From:

To:

Subject:

Attach Job Report

⚠ If you want to send notifications for Clone jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server.

9. DR clone summary. Cloned databases are immediately registered with SnapCenter and available for backup protection.

Clone from backup

- 1 Clone Options
- 2 Logs
- 3 Script
- 4 Notification
- 5 Summary

Summary

Clone server: sql-standby.demo.netapp.com

Clone instance: sql-standby

Clone name: tpcc_dr

Mount option: Auto Mount

Prescript full path: None

Prescript arguments: None

Postscript full path: None

Postscript arguments: None

Send email: No

Previous Finish

NetApp SnapCenter®

Microsoft SQL Server

View Database search by name

Resources	Name	Instance	Host	Last Backup	Overall Status	Type
Monitor	master	sql1	sql1.demo.netapp.com		Not available for backup	System database
Reports	model	sql1	sql1.demo.netapp.com		Not available for backup	System database
Hosts	msdb	sql1	sql1.demo.netapp.com		Not available for backup	System database
Storage Systems	tempdb	sql1	sql1.demo.netapp.com		Not available for backup	System database
Settings	tpcc	sql1	sql1.demo.netapp.com	09/22/2021 5:35:08 PM	Backup failed, Schedules on hold	User database
Alerts	master	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database
	model	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database
	msdb	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database
	tempdb	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database
	tpcc_clone	sql-standby	sql-standby.demo.netapp.com		Not protected	User database
	tpcc_dev	sql-standby	sql-standby.demo.netapp.com		Not protected	User database
	tpcc_dr	sql-standby	sql-standby.demo.netapp.com		Not protected	User database

Post DR clone validation and configuration for SQL

1. Monitor clone job status.

NetApp SnapCenter®

Jobs Schedules Events Logs

search by name

Jobs - Filter	ID	Status	Name	Start date	End date	Owner
	1052	✓	Clone from backup 'sql1_tpcc_09-19-2021_18.25.01.4134'	09/20/2021 2:36:17 PM	09/20/2021 2:37:06 PM	demo/sqlqdba
	1047	✓	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/20/2021 2:35:01 PM	09/20/2021 2:37:08 PM	demo/sqlqdba
	1045	✓	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/20/2021 2:28:17 PM	09/20/2021 2:30:25 PM	demo/sqlqdba
	1044	✓	Clone from backup 'sql1_tpcc_09-17-2021_18.25.01.4218'	09/20/2021 1:39:24 PM	09/20/2021 1:40:09 PM	demo/sqlqdba
	1042	✓	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/20/2021 1:35:01 PM	09/20/2021 1:37:08 PM	demo/sqlqdba
	1040	✓	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/20/2021 12:35:01 PM	09/20/2021 12:37:08 PM	demo/sqlqdba

2. Validate that last transaction has been replicated and recovered with all log file clones and recovery.

```
Administrator: Command Prompt - sqlcmd - SQLCMD
C:\Users\administrator.DEMO>sqlcmd
1> select host_name()
2> go
-----
SQL-STANDBY
(1 rows affected)
1> use tpcc_dr
2> go
Changed database context to 'tpcc_dr'.
1> select * from snap_sync
2> go
event                                     dt
-----
test_snap_mirror DR for SQL                2021-09-20 14:23:04.533
(1 rows affected)
1> select getdate()
2> go
-----
2021-09-20 14:39:19.937
(1 rows affected)
1>
```

3. Configure a new SnapCenter log directory on the DR server for SQL Server log backup.
4. Split the cloned volume off of the replicated source volume.
5. Reverse replication from the cloud to on-premises and rebuild the failed on-premises database server.

Where to go for help?

If you need help with this solution and use cases, please join the [NetApp Solution Automation community support Slack channel](#) and look for the solution-automation channel to post your questions or inquiries.

DB Automation Toolkits

SnapCenter Oracle Clone Lifecycle Automation

Allen Cao, Niyaz Mohamed, NetApp

This solution provides an Ansible based automation toolkit for configuring Oracle database High Availability and Disaster Recovery (HA/DR) with AWS FSx ONTAP as Oracle database storage and EC2 instances as the compute instances in AWS.

Purpose

Customers love the FlexClone feature of NetApp ONTAP storage for databases with significant storage cost savings. This Ansible based toolkit automates the setup, cloning, and refreshing of cloned Oracle databases on schedule using the NetApp SnapCenter command line utilities for streamlined lifecycle management. The toolkit is applicable to Oracle databases deployed to ONTAP storage either on-premises or public cloud and managed by NetApp SnapCenter UI tool.

This solution addresses the following use cases:

- Setup Oracle database clone-specification configuration file.
- Create and refresh clone Oracle database on user defined schedule.

Audience

This solution is intended for the following people:

- A DBA who manages Oracle databases with SnapCenter.
- A storage administrator who manages ONTAP storage with SnapCenter.

- An application owner who has access to SnapCenter UI.

License

By accessing, downloading, installing or using the content in this GitHub repository, you agree the terms of the License laid out in [License file](#).



There are certain restrictions around producing and/or sharing any derivative works with the content in this GitHub repository. Please make sure you read the terms of the License before using the content. If you do not agree to all of the terms, do not access, download or use the content in this repository.

Solution deployment

Prerequisites for deployment

Deployment requires the following prerequisites.

```
Ansible controller:  
  Ansible v.2.10 and higher  
  ONTAP collection 21.19.1  
  Python 3  
Python libraries:  
  netapp-lib  
  xmltodict  
  jmespath
```

```
SnapCenter server:  
  version 5.0  
  backup policy configured  
  Source database protected with a backup policy
```

```
Oracle servers:  
  Source server managed by SnapCenter  
  Target server managed by SnapCenter  
  Target server with identical Oracle software stack as source server  
  installed and configured
```

Download the toolkit

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

Ansible target hosts file configuration

The toolkit includes a hosts file which define the targets that an Ansible playbook running against. Usually, it is the target Oracle clone hosts. Following is an example file. A host entry includes target host IP address as well as ssh key for an admin user access to the host to execute clone or refresh command.

#Oracle clone hosts

```
[clone_1]  
ora_04.cie.netapp.com ansible_host=10.61.180.29  
ansible_ssh_private_key_file=ora_04.pem
```

```
[clone_2]
```

```
[clone_3]
```

Global variables configuration

The Ansible playbooks take variable inputs from several variable files. Below is an example global variable file vars.yml.

```
# ONTAP specific config variables
```

```
# SnapCtr specific config variables
```

```
snapctr_usr: xxxxxxxx  
snapctr_pwd: 'xxxxxxx'
```

```
backup_policy: 'Oracle Full offline Backup'
```

```
# Linux specific config variables
```

```
# Oracle specific config variables
```

Host variables configuration

Host variables are defined in host_vars directory named as {{ host_name }}.yml. Below is an example of target Oracle host variable file ora_04.cie.netapp.com.yml that shows typical configuration.

```
# User configurable Oracle clone db host specific parameters
```

```
# Source database to clone from  
source_db_sid: NTAP1  
source_db_host: ora_03.cie.netapp.com
```

```
# Clone database  
clone_db_sid: NTAP1DEV
```

```
snapctr_obj_id: '{{ source_db_host }}\{{ source_db_sid }}'
```

Additional clone target Oracle server configuration

Clone target Oracle server should have the same Oracle software stack as source Oracle server installed and patched. Oracle user `.bash_profile` has `$ORACLE_BASE`, and `$ORACLE_HOME` configured. Also, `$ORACLE_HOME` variable should match with source Oracle server setting. Following is an example.

```
# .bash_profile
```

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

```
# User specific environment and startup programs
export ORACLE_BASE=/u01/app/oracle
export ORACLE_HOME=/u01/app/oracle/product/19.0.0/NTAP1
```

Playbook execution

There are total of three playbooks to execute Oracle database clone lifecycle with SnapCenter CLI utilities.

1. Install Ansible controller prerequisites - one time only.

```
ansible-playbook -i hosts ansible_requirements.yml
```

2. Setup clone specification file - one time only.

```
ansible-playbook -i hosts clone_1_setup.yml -u admin -e
@vars/vars.yml
```

3. Create and refresh clone database regularly from crontab with a shell script to call a refresh playbook.

```
0 */4 * * * /home/admin/na_oracle_clone_lifecycle/clone_1_refresh.sh
```

For an additional clone database, create a separate `clone_n_setup.yml` and `clone_n_refresh.yml`, and `clone_n_refresh.sh`. Configure the Ansible target hosts and `hostname.yml` file in `host_vars` directory accordingly.

Where to find additional information

To learn more about the NetApp solution automation, review the following website [NetApp Solution Automation](#)

Automated Oracle Migration

NetApp Solutions Engineering Team

This solution provides an Ansible based automation toolkit for migrating Oracle database using PDB relocation with maximum availability methodology. The migration can be any combinations of on-premises and cloud as either source or target.

Purpose

This toolkit automates Oracle database migration from on-premises to AWS cloud with FSx ONTAP storage and EC2 compute instance as target infrastructure. It assumes the customer already has an on-premises Oracle database deployed in the CDB/PDB model. The toolkit will allow the customer to relocate a named PDB from a container database on an Oracle host using the Oracle PDB relocation procedure with a maximum availability option. That means the source PDB on any on-premises storage array relocates to a new container database with minimal service interruption. The Oracle relocation procedure will move the Oracle data files while database is online. It subsequently reroutes user sessions from on-premises to the relocated database services at the time of switching over when all data files move over to AWS cloud. The underlined technology is proven Oracle PDB hot clone methodology.



Although the migration toolkit is developed and validated on AWS cloud infrastructure, it builds on Oracle application-level solutions. Therefore, the toolkit is applicable to other public cloud platforms, such as Azure, GCP, etc.

This solution addresses the following use cases:

- Create migration user and grant required privileges at on-prem source DB server.
- Relocate a PDB from on-premises CDB to a target CDB in cloud while the source PDB is online until switch over.

Audience

This solution is intended for the following people:

- A DBA who migrates Oracle databases from on-premises to AWS cloud.
- A database solution architect who is interested in Oracle database migration from on-premises to AWS cloud.
- A storage administrator who manages AWS FSx ONTAP storage that supports Oracle databases.
- An application owner who likes to migrate Oracle database from on-premises to AWS cloud.

License

By accessing, downloading, installing or using the content in this GitHub repository, you agree the terms of the License laid out in [License file](#).



There are certain restrictions around producing and/or sharing any derivative works with the content in this GitHub repository. Please make sure you read the terms of the License before using the content. If you do not agree to all of the terms, do not access, download or use the content in this repository.

Solution deployment

Prerequisites for deployment

Deployment requires the following prerequisites.

```
Ansible v.2.10 and higher
ONTAP collection 21.19.1
Python 3
Python libraries:
  netapp-lib
  xmltodict
  jmespath
```

```
Source Oracle CDB with PDBs on-premises
Target Oracle CDB in AWS hosted on FSx and EC2 instance
Source and target CDB on same version and with same options installed
```

```
Network connectivity
  Ansible controller to source CDB
  Ansible controller to target CDB
  Source CDB to target CDB on Oracle listener port (typical 1521)
```

Download the toolkit

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

Host variables configuration

Host variables are defined in `host_vars` directory named as `{{ host_name }}`.yml. An example host variable file `host_name.yml` is included to demonstrate typical configuration. Following are key considerations:

```
Source Oracle CDB - define host specific variables for the on-prem CDB
ansible_host: IP address of source database server host
source_oracle_sid: source Oracle CDB instance ID
source_pdb_name: source PDB name to migrate to cloud
source_file_directory: file directory of source PDB data files
target_file_directory: file directory of migrated PDB data files
```

```
Target Oracle CDB - define host specific variables for the target CDB
including some variables for on-prem CDB
ansible_host: IP address of target database server host
target_oracle_sid: target Oracle CDB instance ID
target_pdb_name: target PDB name to be migrated to cloud (for max
availability option, the source and target PDB name must be the same)
source_oracle_sid: source Oracle CDB instance ID
source_pdb_name: source PDB name to be migrated to cloud
source_port: source Oracle CDB listener port
source_oracle_domain: source Oracle database domain name
source_file_directory: file directory of source PDB data files
target_file_directory: file directory of migrated PDB data files
```

DB server host file configuration

AWS EC2 instance use IP address for host naming by default. If you use different name in hosts file for Ansible, setup host naming resolution in `/etc/hosts` file for both source and target server. Following is an example.

```
127.0.0.1    localhost localhost.localdomain localhost4
localhost4.localhost4
::1         localhost localhost.localdomain localhost6
localhost6.localhost6
172.30.15.96 source_db_server
172.30.15.107 target_db_server
```

Playbook execution - executed in sequence

1. Install Ansible controller prerequisites.

```
ansible-playbook -i hosts requirements.yml
```

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

2. Execute pre-migration tasks against on-prem server - assuming admin is ssh user for connection to on-prem Oracle host with sudo permission.

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

3. Execute Oracle PDB relocation from on-prem CDB to target CDB in AWS EC2 instance - assuming ec2-user for EC2 DB instance connection, and db1.pem with ec2-user ssh key pairs.

```
ansible-playbook -i hosts ora_pdb_relocate.yml -u ec2-user --private  
-key db1.pem -t ora_pdb_relo_primary
```

Where to find additional information

To learn more about the NetApp solution automation, review the following website [NetApp Solution Automation](#)

Automated Oracle HA/DR in AWS FSx ONTAP

NetApp Solutions Engineering Team

This solution provides an Ansible based automation toolkit for configuring Oracle database High Availability and Disaster Recovery (HA/DR) with AWS FSx ONTAP as Oracle database storage and EC2 instances as the compute instances in AWS.

Purpose

This toolkit automates the tasks of setting up and managing a High Availability and Disaster Recovery (HR/DR) environment for Oracle database deployed in AWS cloud with FSx for ONTAP storage and EC2 compute instances.

This solution addresses the following use cases:

- Setup HA/DR target host - kernel configuration, Oracle configuration to match up with source server host.
- Setup FSx ONTAP - cluster peering, vserver peering, Oracle volumes snapmirror relationship setup from source to target.

- Backup Oracle database data via snapshot - execute off crontab
- Backup Oracle database archive log via snapshot - execute off crontab
- Run failover and recovery on HA/DR host - test and validate HA/DR environment
- Run resync after failover test - re-establish database volumes snapmirror relationship in HA/DR mode

Audience

This solution is intended for the following people:

- A DBA who set up Oracle database in AWS for high availability, data protection, and disaster recovery.
- A database solution architect who is interested in storage level Oracle HA/DR solution in the AWS cloud.
- A storage administrator who manages AWS FSx ONTAP storage that supports Oracle databases.
- An application owner who like to stand up Oracle database for HA/DR in AWS FSx/EC2 environment.

License

By accessing, downloading, installing or using the content in this GitHub repository, you agree the terms of the License laid out in [License file](#).



There are certain restrictions around producing and/or sharing any derivative works with the content in this GitHub repository. Please make sure you read the terms of the License before using the content. If you do not agree to all of the terms, do not access, download or use the content in this repository.

Solution deployment

Prerequisites for deployment

Deployment requires the following prerequisites.

```
Ansible v.2.10 and higher
ONTAP collection 21.19.1
Python 3
Python libraries:
  netapp-lib
  xmltodict
  jmespath
```

```
AWS FSx storage as is available
```

```
AWS EC2 Instance
  RHEL 7/8, Oracle Linux 7/8
  Network interfaces for NFS, public (internet) and optional management
  Existing Oracle environment on source, and the equivalent Linux
  operating system at the target
```

Download the toolkit

```
git clone https://github.com/NetApp/na_ora_hadr_failover_resync.git
```

Global variables configuration

The Ansible playbooks are variable driven. An example global variable file `fsx_vars_example.yml` is included to demonstrate typical configuration. Following are key considerations:

ONTAP - retrieve FSx storage parameters using AWS FSx console for both source and target FSx clusters.

cluster name: source/destination

cluster management IP: source/destination

inter-cluster IP: source/destination

vserver name: source/destination

vserver management IP: source/destination

NFS lifs: source/destination

cluster credentials: fsxadmin and vsadmin pwd to be updated in `roles/ontap_setup/defaults/main.yml` file

Oracle database volumes - they should have been created from AWS FSx console, volume naming should follow strictly with following standard:

Oracle binary: `{{ host_name }}_bin`, generally one lun/volume

Oracle data: `{{ host_name }}_data`, can be multiple luns/volume, add additional line for each additional lun/volume in variable such as `{{ host_name }}_data_01`, `{{ host_name }}_data_02` ...

Oracle log: `{{ host_name }}_log`, can be multiple luns/volume, add additional line for each additional lun/volume in variable such as `{{ host_name }}_log_01`, `{{ host_name }}_log_02` ...

host_name: as defined in hosts file in root directory, the code is written to be specifically matched up with host name defined in host file.

Linux and DB specific global variables - keep it as is.

Enter redhat subscription if you have one, otherwise leave it black.

Host variables configuration

Host variables are defined in `host_vars` directory named as `{{ host_name }}`.yml. An example host variable file `host_name.yml` is included to demonstrate typical configuration. Following are key considerations:

```
Oracle - define host specific variables when deploying Oracle in
multiple hosts concurrently
  ansible_host: IP address of database server host
  log_archive_mode: enable archive log archiving (true) or not (false)
  oracle_sid: Oracle instance identifier
  pdb: Oracle in a container configuration, name pdb_name string and
number of pdbs (Oracle allows 3 pdbs free of multitenant license fee)
  listener_port: Oracle listener port, default 1521
  memory_limit: set Oracle SGA size, normally up to 75% RAM
  host_datastores_nfs: combining of all Oracle volumes (binary, data,
and log) as defined in global vars file. If multi luns/volumes, keep
exactly the same number of luns/volumes in host_var file
```

```
Linux - define host specific variables at Linux level
  hugepages_nr: set hugepage for large DB with large SGA for
performance
  swap_blocks: add swap space to EC2 instance. If swap exist, it will
be ignored.
```

DB server host file configuration

AWS EC2 instance use IP address for host naming by default. If you use different name in hosts file for Ansible, setup host naming resolution in `/etc/hosts` file for both source and target servers. Following is an example.

```
127.0.0.1  localhost localhost.localdomain localhost4
localhost4.localdomain4
::1       localhost localhost.localdomain localhost6
localhost6.localdomain6
172.30.15.96 db1
172.30.15.107 db2
```

Playbook execution - executed in sequence

1. Install Ansible controller prerequisites.

```
ansible-playbook -i hosts requirements.yml
```

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

2. Setup target EC2 DB instance.

```
ansible-playbook -i hosts ora_dr_setup.yml -u ec2-user --private-key  
db2.pem -e @vars/fsx_vars.yml
```

3. Setup FSx ONTAP snapmirror relationship between source and target database volumes.

```
ansible-playbook -i hosts ontap_setup.yml -u ec2-user --private-key  
db2.pem -e @vars/fsx_vars.yml
```

4. Backup Oracle database data volumes via snapshot from crontab.

```
10 * * * * cd /home/admin/na_ora_hadr_failover_resync &&  
/usr/bin/ansible-playbook -i hosts ora_replication_cg.yml -u ec2-  
user --private-key db1.pem -e @vars/fsx_vars.yml >>  
logs/snap_data_`date +%Y-%m%d-%H%M%S`.log 2>&1
```

5. Backup Oracle database archive log volumes via snapshot from crontab.

```
0,20,30,40,50 * * * * cd /home/admin/na_ora_hadr_failover_resync &&  
/usr/bin/ansible-playbook -i hosts ora_replication_logs.yml -u ec2-  
user --private-key db1.pem -e @vars/fsx_vars.yml >>  
logs/snap_log_`date +%Y-%m%d-%H%M%S`.log 2>&1
```

6. Run failover and recover Oracle database on target EC2 DB instance - test and validate HA/DR configuration.

```
ansible-playbook -i hosts ora_recovery.yml -u ec2-user --private-key  
db2.pem -e @vars/fsx_vars.yml
```

7. Run resync after failover test - re-establish database volumes snapmirror relationship in replication mode.

```
ansible-playbook -i hosts ontap_ora_resync.yml -u ec2-user --private  
-key db2.pem -e @vars/fsx_vars.yml
```

Where to find additional information

To learn more about the NetApp solution automation, review the following website [NetApp Solution Automation](#)

AWS FSx ONTAP Cluster and EC2 Instance Provision

NetApp Solutions Engineering Team

This solution provides a Terraform based automation toolkit for provisioning of FSx ONTAP cluster and EC2 compute instance.

Purpose

This toolkit automates the tasks of provisioning of an AWS FSx ONTAP storage cluster and an EC2 compute instance, which can be subsequently used for database deployment.

This solution addresses the following use cases:

- Provision an EC2 compute instance in AWS cloud in a predefined VPC subnet and set ssh key for EC2 instance access as ec2-user.
- Provision an AWS FSx ONTAP storage cluster in desired availability zones and configure a storage SVM and set cluster admin user fsxadmin password.

Audience

This solution is intended for the following people:

- A DBA who manages databases in AWS EC2 environment.
- A database solution architect who is interested in database deployment in AWS EC2 ecosystem.
- A storage administrator who manages AWS FSx ONTAP storage that supports databases.
- An application owner who likes to standup database in AWS EC2 ecosystem.

License

By accessing, downloading, installing or using the content in this GitHub repository, you agree the terms of the License laid out in [License file](#).



There are certain restrictions around producing and/or sharing any derivative works with the content in this GitHub repository. Please make sure you read the terms of the License before using the content. If you do not agree to all of the terms, do not access, download or use the content in this repository.

Solution deployment

Prerequisites for deployment

Deployment requires the following prerequisites.

```
An Organization and AWS account has been setup in AWS public cloud
An user to run the deployment has been created
IAM roles has been configured
IAM roles granted to user to permit provisioning the resources
```

VPC and security configuration

```
A VPC has been created to host the resources to be provisioned
A security group has been configured for the VPC
A ssh key pair has been created for EC2 instance access
```

Network configuration

```
Subnets has been created for VPC with network segments assigned
Route tables and network ACL configured
NAT gateways or internet gateways configured for internet access
```

Download the toolkit

```
git clone https://github.com/NetApp/na_aws_fsx_ec2_deploy.git
```

Connectivity and authentication

The toolkit is supposed to be executed from an AWS cloud shell. AWS cloud shell is a browser-based shell that makes it easy to securely manage, explore, and interact with your AWS resources. CloudShell is pre-authenticated with your console credentials. Common development and operations tools are pre-installed, so no local installation or configuration is required.

Terraform provider.tf and main.tf files configuration

The provider.tf defines the provider that Terraform is provisioning resources from via API calls. The main.tf defines the resources and attributes of resources that are to be provisioned. Following are some details:

```
provider.tf:
terraform {
  required_providers {
    aws = {
      source = "hashicorp/aws"
      version = "~> 4.54.0"
    }
  }
}
```

```
main.tf:
resource "aws_instance" "ora_01" {
  ami                = var.ami
  instance_type      = var.instance_type
  subnet_id          = var.subnet_id
  key_name           = var.ssh_key_name
  root_block_device {
    volume_type      = "gp3"
    volume_size      = var.root_volume_size
  }
  tags = {
    Name              = var.ec2_tag
  }
}
.....
```

Terraform variables.tf and terraform.tfvars configuration

The variables.tf declares the variables to be used in main.tf. The terraform.tfvars contains the actual values for the variables. Following are some examples:

```
variables.tf:
  ### EC2 instance variables ###
```

```
variable "ami" {
  type      = string
  description = "EC2 AMI image to be deployed"
}
```

```
variable "instance_type" {
  type      = string
  description = "EC2 instance type"
}
....
```

```
terraform.tfvars:
  # EC2 instance variables
```

```
ami              = "ami-06640050dc3f556bb" //RedHat 8.6 AMI
instance_type    = "t2.micro"
ec2_tag          = "ora_01"
subnet_id        = "subnet-04f5fe7073ff514fb"
ssh_key_name     = "sufi_new"
root_volume_size = 30
....
```

Step by step procedures - executed in sequence

1. Install Terraform in AWS cloud shell.

```
git clone https://github.com/tfutils/tfenv.git ~/.tfenv
```

```
mkdir ~/bin
```

```
ln -s ~/.tfenv/bin/* ~/bin/
```

```
tfenv install
```

```
tfenv use 1.3.9
```

2. Download the toolkit from NetApp GitHub public site

```
git clone https://github.com/NetApp-  
Automation/na_aws_fsx_ec2_deploy.git
```

3. Run init to initialize terraform

```
terraform init
```

4. Output the execution plan

```
terraform plan -out=main.plan
```

5. Apply the execution plan

```
terraform apply "main.plan"
```

6. Run destroy to remove the resources when done

```
terraform destroy
```


Where to find additional information

To learn more about the NetApp solution automation, review the following website [NetApp Solution Automation](#)

DB Sizing Toolkits

Oracle Sizing Guidance for Azure NetApp Files

Allen Cao, Niyaz Mohamed, NetApp

This solution provides an useful toolkit for sizing compute and storage for Oracle deployment on ANF in Azure cloud.

Purpose

Moving existing Oracle workload from one platform to another, such as from on-prem to public cloud, needs sizing compute and storage in the target platform to meet performance and service level requirements. This documentation demonstrates a simple toolkit to accomplish that goal.

Unlike a new database application, which may grow over time, an existing Oracle workload has established workload patterns in compute and storage requirements, which are recorded in an Oracle Workload Repository or AWR. This toolkit utilizes an HTML parser to retrieve relevant information from Oracle AWR. The results are supplemented by additional sizing information obtained via SQL scripts against the database to provide meaningful compute and storage guidance when relocating the Oracle database.

This solution addresses the following use cases:

- Provide sizing guidance for Oracle database server compute when relocating database from on-prem to Microsoft Azure cloud.
- Provide sizing guidance for Oracle database server storage when relocating database from on-prem to Microsoft Azure NetApp Files.

Audience

This solution is intended for the following people:

- A DBA who manages Oracle databases in on-prem private data center or Microsoft Azure cloud environment.
- A storage administrator who manages on-prem storage or Microsoft Azure NetApp Files storage that supports Oracle databases.
- An application owner who likes to migrate Oracle database from on-prem to Microsoft Azure cloud.

License

By accessing, downloading, installing or using the content in this toolkit repository, you agree the terms of the License laid out in [License file](#).



There are certain restrictions around producing and/or sharing any derivative works with the content in this toolkit repository. Please make sure you read the terms of the License before using the content. If you do not agree to all of the terms, do not access, download or use the content in this repository.

Solution deployment

Prerequisites for deployment

Deployment requires the following prerequisites.

- Oracle AWR reports that capture the snapshots of database activities during peak application workload.
- Access to Oracle database to execute SQL scripts with DBA privilege.

Download the toolkit

Retrieve the toolkit from repository [Oracle Sizing Guidance for ANF](#)


How to use the toolkit?

The toolkit consists of a web-based HTML parser and two SQL scripts to gather Oracle database information. The output is then input into an Excel template to generate sizing guidance of computing and storage for the Oracle database server.


- Use an [HTML parser](#) AWR module to retrieve sizing information of a current Oracle database from an AWR report.
- Execute ora_db_data_size.sql as a DBA to retrieve physical Oracle data file size from database.
- Execute ora_db_logs_size.sql as a DBA to retrieve Oracle archived logs size with desired archive logs retention window (days).
- Input sizing information obtained above into excel template file oracle_db_sizing_template_anf.xlsx to create a sizing guidance on compute and storage for Oracle DB server.

Toolkit usage demonstration

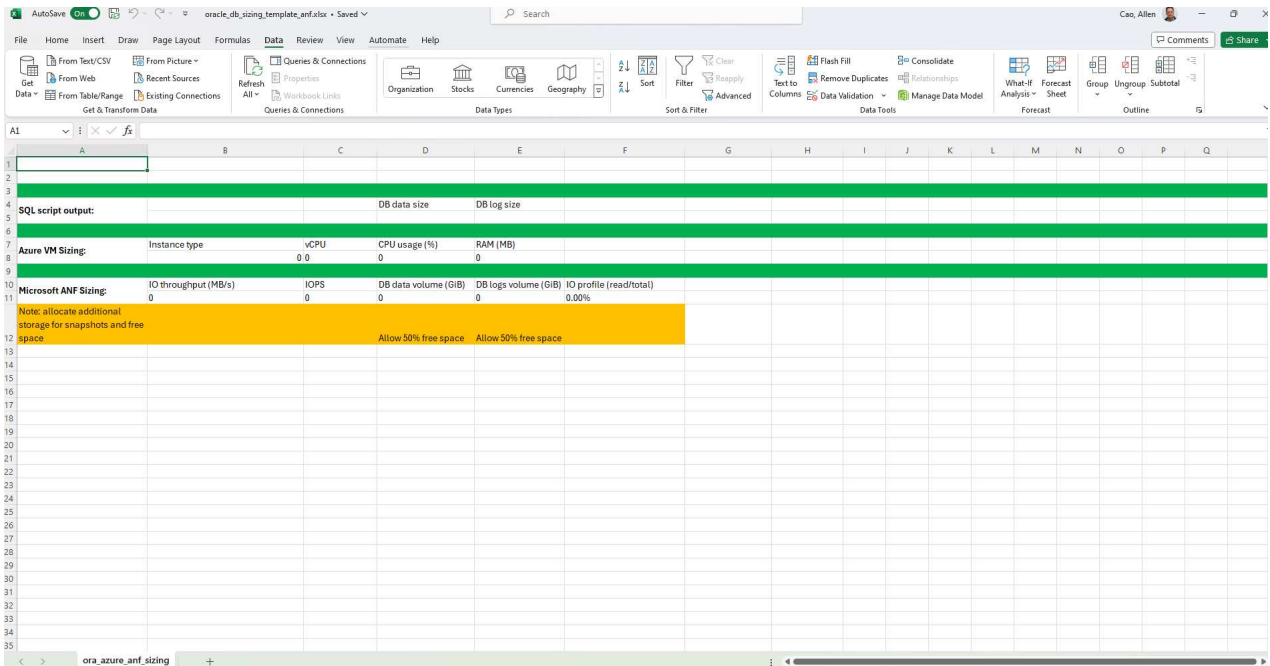
1. Open HTML parser AWR module.




2. Check output format as .csv and click Upload files to upload awr report. The parser returns results in a HTML page with a table summary as well as an output.csv file in Download folder.



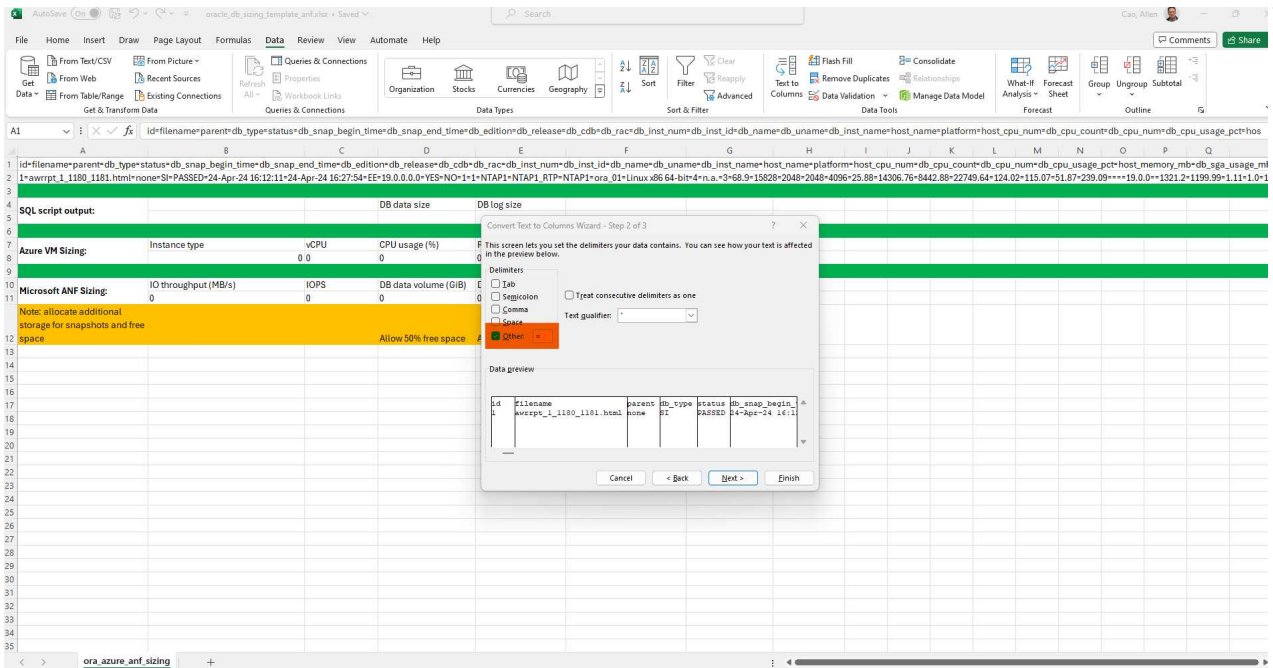
3. Open the excel template file and copy paste the csv content into column A and cell 1 to generate the DB server sizing information.




4. Highlight column A and fields 1 and 2, click on Data, then Text to Columns to open the Text Wizard. Choose Delimited, then Next to next screen.



5. Check Other, then enter '=' as Delimiters. Click on Next to next screen.



6. Click on **Finish** to complete the string conversion into readable column format. Note the VM and ANF sizing fields have been populated with data retrieved from the Oracle AWR report.



A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q			
1	id	filename	parent	db_type	status	db_snap_begin_time	db_snap_end_time	db_edition	db_releas	db_cdb	db_rac	db_inst_ni	db_inst_ic	db_name	db_unam	db_inst_n	host_nam	platform	
2		1 awrrpt_1_1180_1181.html	none	SI	PASSED	4/24/2024 16:12	4/24/2024 16:27	EE	19.0.0.0.0	YES	NO	1		1	NTAP1	NTAP1_RT	NTAP1	ora_01	Linux x86
4	SQL script output:			DB data size	DB log size														
7	Azure VM Sizing:																		
	Instance type	vCPU	CPU usage (%)	RAM (MB)															
	SI	4	68.9	15828															
10	Microsoft ANF Sizing:																		
	IO throughput (MB/s)	IOPS	DB data volume (GiB)	DB logs volume (GiB)	IO profile (read/total)														
	239.09	22749.64	0	0	62.89%														
12	Note: allocate additional storage for snapshots and free space			Allow 50% free space	Allow 50% free space														

7. Execute script ora_db_data_size.sql, ora_db_logs_size.sql as a DBA in sqlplus to retrieve existing Oracle database data size and archived logs size with the number of days of retention window.

```

[oracle@ora_01 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Tue Mar 5 15:25:27 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> @/home/oracle/ora_db_data_size.sql;

Aggregate DB File Size, GiB Aggregate DB File RW, GiB Aggregate DB
File RO, GiB
-----
-----
                                159.05                                159.05
0

```

```

SQL> @/home/oracle/ora_db_logs_size.sql;
Enter value for archivelog_retention_days: 14
old 6:      where first_time >= sysdate -
&archivelog_retention_days
new 6:      where first_time >= sysdate - 14

Log Size, GiB
-----
          93.83

SQL>

```



The database sizing information retrieved using above scripts is the sum of actual size of all physical database data files or log files. It does not factor into the free space that may be available inside each data file.

8. Input the result into excel file to complete the sizing guidance output.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q		
1	id	filename	parent	db_type	status	db_snap_begin_time	db_snap_end_time	db_edition	db_releas	db_cdb	db_rac	db_inst_ni	db_inst_ic	db_name	db_unam	db_inst_ni	host_nam	platform	
2		1 awrrpt_1_1180_1181.html	none	SI	PASSED	4/24/2024 16:12	4/24/2024 16:27	EE	19.0.0.0.0	YES	NO	1		1	NTAP1	NTAP1_RT	NTAP1	ora_01	Linux x86
3				DB data size	DB log size														
4				159.05	93.83														
5	SQL script output:																		
6																			
7	Azure VM Sizing:			Instance type	vCPU	CPU usage (%)	RAM (MB)												
8				SI	4	66.9	15828												
9	Microsoft ANF Sizing:																		
10	IO throughput (MB/s)		IOPS	DB data volume (GiB)	DB logs volume (GiB)	IO profile (read/total)													
11	239.09		22749.64	318.1	187.66	62.89%													
12	Note: allocate additional storage for snapshots and free space			Allow 50% free space	Allow 50% free space														
13																			
14																			
15																			
16																			
17																			
18																			
19																			
20																			
21																			
22																			
23																			
24																			
25																			
26																			
27																			
28																			
29																			
30																			
31																			
32																			
33																			
34																			
35																			

9. ANF uses a three-tier service level (Standard, Premium, Ultra) to manage database volume throughput limit. Refer to [Service levels for Azure NetApp Files](#) for details. Based on sizing guidance output, choose an ANF service level that provides throughput that meet the requirement for the database.

Where to find additional information

To learn more about the NetApp database solutions, review the following website [NetApp Enterprise Database Solutions](#)

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.