

Databases

NetApp Solutions NetApp May 17, 2024

This PDF was generated from https://docs.netapp.com/us-en/netappsolutions/databases/automation_ora_aws-fsx_iscsi.html on May 17, 2024. Always check docs.netapp.com for the latest.

Table of Contents

etApp Enterprise Database Solutions	1
Oracle Database	1
Microsoft SQL Server	15
Open Source Databases	1
SnapCenter for Databases	1
DB Automation Toolkits	3
DB Sizing Toolkits	52

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

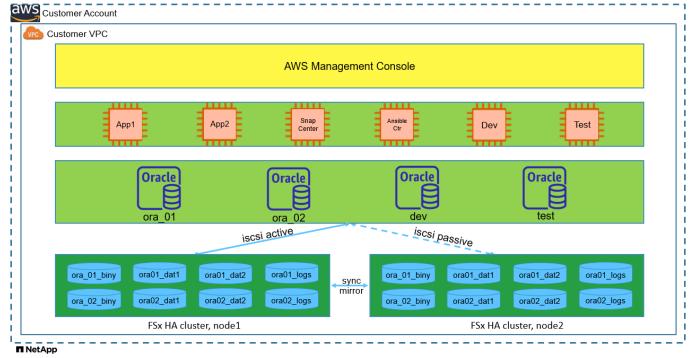
This solution is intended for the following people:

- 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.

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

Oracle database configuration in the lab environment

Server	Database	DB Storage
ora_01	NTAP1(NTAP1_PDB1,NTAP1_PD B2,NTAP1_PDB3)	iSCSI luns on Amazon FSx ONTAP file system
ora_02	NTAP2(NTAP2_PDB1,NTAP2_PD B2,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
```

 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"
```

- "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

1. Ansible target hosts file configuration:

```
# Enter Amazon FSx ONTAP management IP address
[ontap]
172.16.9.32
# Enter name for ec2 instance (not default IP address naming) to be
deployed one by one, follow by ec2 instance IP address, and ssh
private key of ec2-user for the instance.
[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
######
*****
# Enter the supported ONTAP platform: on-prem, aws-fsx.
ontap platform: aws-fsx
# Enter ONTAP cluster management user credentials
username: "fsxadmin"
password: "xxxxxxxx"
*****
###
          Linux env specific config variables
###
```

```
****
# Enter RHEL subscription to enable repo
redhat sub username: xxxxxxx
redhat sub password: "xxxxxxxx"
****
###
           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: xxxxxxx
```

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 -
Production
Version 19.18.0.0.0
SQL> select name, open mode, log mode from v$database;
NAME OPEN MODE
                    LOG_MODE
----- ------ ------ ------
NTAP1 READ WRITE
                           ARCHIVELOG
SQL> show pdbs
                                 OPEN MODE RESTRICTED
  CON_ID CON_NAME
2 PDB$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$logfile;
   MEMBER
   _____
                              _____
   +DATA/NTAP1/ONLINELOG/group 3.264.1155055531
   +LOGS/NTAP1/ONLINELOG/group 3.259.1155055539
   +DATA/NTAP1/ONLINELOG/group 2.263.1155055531
   +LOGS/NTAP1/ONLINELOG/group 2.257.1155055539
   +DATA/NTAP1/ONLINELOG/group 1.262.1155055531
   +LOGS/NTAP1/ONLINELOG/group 1.258.1155055539
   6 rows selected.
   SQL> exit
   Disconnected from Oracle Database 19c Enterprise Edition Release
   19.0.0.0 - Production
   Version 19.18.0.0.0
2. Validate Oracle listener.
   [oracle@ip-172-30-15-40 ~]$ 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 "0c03aafa7c6fd2e5e063280f1eacfbe0.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
                                        STABLE
ora.driver.afd
   1
       ONLINE ONLINE
                      ip-172-30-15-40
                                        STABLE
ora.evmd
       ONLINE ONLINE ip-172-30-15-40
   1
                                        STABLE
ora.ntap1.db
      ONLINE ONLINE ip-172-30-15-40
    1
```

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

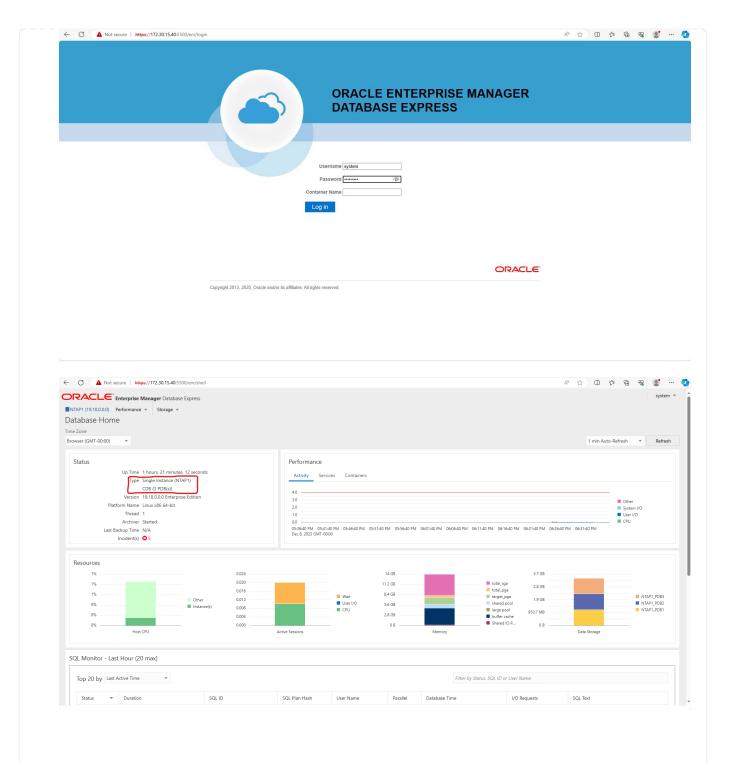
racle/product/19.0.0

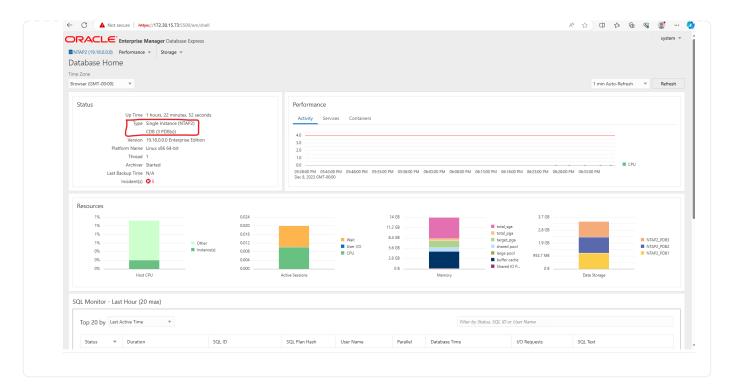
/NTAP1, STABLE

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
                                   512 4096 4194304
MOUNTED EXTERN N 512
                       0 155376
163840 155376
                                                      0
N DATA/
MOUNTED EXTERN N
                                    512 4096 4194304
                     512
                                 80972
81920 80972
                          0
                                                     0
N LOGS/
ASMCMD> 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
ASMCMD> afd state
ASMCMD-9526: The AFD state is 'LOADED' and filtering is 'ENABLED' on
host 'ip-172-30-15-40.ec2.internal'
ASMCMD> exit
```

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





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=AL!4422!3!467723097970!e!!g!!aws%20ec2&ef_id=Cj0KCQiA54 KfBhCKARIsAJzSrdqwQrghn6I71jiWzSeaT9Uh1-vY-VfhJixFxnv5rWwn2S7RqZOTQ0aAh7eEALw_wcB:G:s&s_kwcid=AL!4422!3!467723097970!e!!g!!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-

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

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

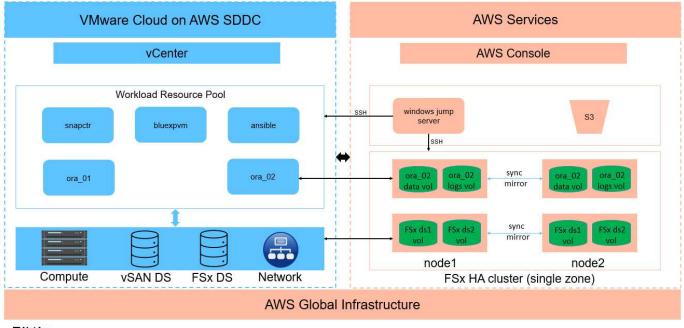
This solution is intended for the following people:

- · 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 pdb3)	2,cdb1_ VMDK datastore on FSx ONTAP
ora_01	cdb2(cdb2_pdb)	VMDK datastore on FSx ONTAP
ora_02	cdb3(cdb3_pdb1,cdb3_pdb2 pdb3)	2,cdb3_ 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.

< 11 <u>B</u> = Q			Networks Snapshots		
 venter.sddc.52.88.219.219 vmwarev SDDC-Datacenter Discovered virtual machine Discovered virtual machine Management VMs MSX-Edge-0 MSX-Edge-1 MSX-Manager-1 Workloads maible Biblexovm 	GUEST OS LAUNCH REMOTE CONSOLE	 Wirtual Machine Details Power Status Guest OS Vitware Tools DNS Name (1) IP Addresses (2) Encryption ▲ ▲ 	ACTIONS V II Powered On Red Hat Enterprise Linux 8 (64-bit) Running, version:11365 (Guest Managed) Ora_02 192:163:1132 fe80:250:561f.feb6:6295 Not encrypted	Usage It is the state of the state is the state of the state is the state of the st	1
贷 ora_01 贷 ora_02 贷 SCV-JH 贷 snapctr	Memory 16 Hard disk 1 5 Network adapter 1 0 CD/DVD drive 1 0	CPU(s), 413 MHz used 5 GB, I GB memory active 0 GB Thin Provision ① VorKloadDatastore dds-cgw-network-1 (connected) 0 50:56:66:295 onnected Φ × SXI 6.7 and later (VM version 14)	II PCI Devices II	Related Objects III Cluster III Cluster-1 Host To 56.0.68 Resource pool Compute-ResourcePool Networks Maddc-gw-network-1 Storage VorkloadDatastore IIII	

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 stated 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

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
                                Туре
                                        Size
Available Used%
_____ ____
_____ ____
     ora 02 biny aggr1 online RW
nim
                                        50GB
22.98GB
     51%
nim
    ora_02_data aggr1
                        online
                               RW
                                       100GB
18.53GB 80%
      ora 02 logs aggr1 online
                                RW
                                        50GB
nim
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:

	cec8dad373fd1::> mapping show oping 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.
- Discover the FSx iSCSI endpoints using either SVM iSCSI IP address. Change to your environmentspecific 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 -1 all

The expected output from the command is:

```
[ec2-user@ip-172-30-15-58 ~]$ sudo iscsiadm --mode node -1 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:
10.49.0.12,3260] 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.

—		nlun lun show	
device	mode/E-Series host	lun	
	-	lun-pathname	
	adapter	protocol size product	
		/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
                           7 Mar 21 20:13
lrwxrwxrwx 1 root root
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:

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 {
                                 3600a09806c574235472455534e68512d
                wwid
                alias
                                 ora 02 biny 01
        }
        multipath {
                                 3600a09806c574235472455534e685141
                wwid
                alias
                                 ora 02 data 01
        }
        multipath {
                wwid
                                 3600a09806c574235472455534e685142
                alias
                                 ora 02 data 02
        }
        multipath {
                wwid
                                 3600a09806c574235472455534e685143
                alias
                                 ora 02 data 03
        }
        multipath {
                                 3600a09806c574235472455534e685144
                wwid
                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
                           7 Mar 21 20:41 ora 02 data 01 -> ../dm-
lrwxrwxrwx 1 root root
1
                            7 Mar 21 20:41 ora 02 data 02 -> ../dm-
lrwxrwxrwx 1 root root
2
lrwxrwxrwx 1 root root
                            7 Mar 21 20:41 ora 02 data 03 -> ../dm-
3
                            7 Mar 21 20:41 ora 02 data 04 -> ../dm-
lrwxrwxrwx 1 root root
4
lrwxrwxrwx 1 root root
                            7 Mar 21 20:41 ora 02 logs 01 -> ../dm-
5
                            7 Mar 21 20:41 ora 02 logs 02 -> ../dm-
lrwxrwxrwx 1 root root
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/m
apper/ora 02 data 02,/dev/mapper/ora 02 data 03,/dev/mapper/ora 02 d
ata 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 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 0 81780 81920 81780 0 N DATA/ 512 4096 4194304 MOUNTED EXTERN N 512 0 40960 40852 40852 0 N LOGS/ ASMCMD> afd state ASMCMD-9526: The AFD state is 'LOADED' and filtering is 'ENABLED' on host 'ora 02' ASMCMD> 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 <code>\$ORACLE HOME</code> and <code>\$ORACLE SID</code> 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.InstallEdition=EE oracle.install.db.OSDBA_GROUP=dba oracle.install.db.OSOPER_GROUP=oper oracle.install.db.OSDGDBA_GROUP=dba oracle.install.db.OSCMDBA_GROUP=dba oracle.install.db.OSRACDBA_GROUP=dba oracle.install.db.OSRACDBA_GROUP=dba

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.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
    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 _____ _____ READ WRITE CDB3 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.
- Login to each cdb as sysdba with sqlplus and set the DB recovery destination size to the +LOGS disk group size for both cdbs.

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

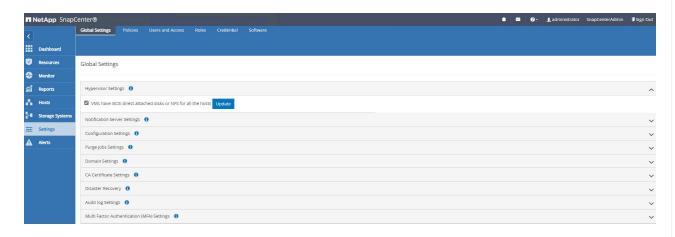
SnapCenter relies on a host-side plug-in on database VM to perform application-aware data protection management activities. For detailed information on NetApp SnapCenter plugin for Oracle, refer to this documentation What can you do with the Plug-in for Oracle Database. The following provides high level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

NetApp	SnapCenter®					. ■ . 0	- 👤 administr	ator SnapCe	enterAdmin	🖡 Sign Out
< Dashboan	Oracle Database							_		
Resources	Search by Name					÷	Modify	Copy	i Details	Delete
Monitor	Name	15	Backup Type	Schedule Type	Replication		Verificatio	1		
Reports	Oracle Archive Logs Backup		LOG, ONLINE	Hourly						
	Oracle Online Full Backup		FULL, ONLINE	Hourly						
Hosts										
- Storage Sy	stems									
Settings										
A Alerts										

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.

<		Global Settings			Credential	Software					
Dash	hboard	Search by Crede	ential Name	\supset						+ New	
Reso	ources	Credential Name					Authentication Mode	Details			
Moni	itor	admin					Linux	UserId:a	dmin		
🖬 Repo	orts										
Host	ts .										
-II Stora	age Systems										
E Setti	ings										
	he l										
Alert											

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

NetApp SnapC	lenter®					1	. 🛛 6	🔹 👤 administrati	or SnapCenterAdmin	🗊 Sign Out
	ONTAP Sto	prage								
Dashboard	Туре О	NTAP SVMs •	Search by Name						+ New	
Resources	ONTAP S	Storage Connections								
Monitor		Name	IL	IP	Cluster Name	User Name	Platform	Con	troller License	
Reports		nim			10.49.0.74		FSx	Not	applicable	
Hosts										
Storage Systems										
Settings										

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

Dashboard	Search by Name					+	
Resources	Name	1h	Туре	System	Plug-in	Version	Overall Status
Monitor	ora 02.vmc.local		Linux	Stand-alone	UNIX, Oracle Database	4.9	Running
Reports							
Hosts							
Storage System							
Ξ Settings							
Alerts							
Alerts							
Alerts							
Alerts							

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.

Resources IP Name Oracle Database Type Host/Cluster Resource Group Policies Last Badup Overall Montor cdb3 Single Instance (Multitement) ora_022mmclocal end end end host pro-	NetApp Snap(• • • •	1 administrator		
Postboard Name Oracle Database Type Host/Cluster Resource Group Policies Policies Last Badup Overall Resource cdba3 Single instance (Multitenant) oracle Database Type doot Columnation oracle Database Type Most Columnation Resource Group Policies Doot Columnation New Policies New Policies		Oracle Database								
Monitor cdb3 Single instance (Multitenant) ora_022/mclocal Not pro Monitor cdb4 Single instance (Multitenant) ora_022/mclocal Not pro Monitor Keports Single instance (Multitenant) ora_022/mclocal Not pro	Dashboard	View Database	Search database	s V					Refresh Resource	: 📩
Monthor cd/b4 Single Instance (Multitenant) ora_02/vmc/ocal Not profession Reports - Hosts	Resources	R Name	Oracle Database Type	Host/Cluster	Resource Group	Policies		Last Back	kup Overall Sta	tus
Reports Hosts B Storage Systems	Monitor	cdb3	Single Instance (Multitenant)	ora_02.vmc.local					Not protec	ted
Nosts Storage Systems	1 Passarta	cdb4	Single Instance (Multitenant)	ora_02.vmc.local					Not protec	ted
Storage Systems										
E Settings	Hosts									
Alerts	Storage Systems									

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.

Ne	tApp SnapCenter®				■ 0.	👤 administr	ator SnapCo	enterAdmin	🗊 Sign Out
	Oracle Database 🔹								
	Search resource groups	search				odity Resource Group	L Back up Now	Maintenance	Delete
	Name	Resource Name	Туре	Host					
	ora_02_data	cdb3	Oracle Database	ora_02.vmc.local					
2	ora_02_logs	cdb4	Oracle Database	ora_02.vmc.local					
•									
8									
Ş									

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.

Hourly						
Start date	10/07/2	2023 08:35 a	m	8		
Expires on	11/07/2	2023 08:35 a	m			
Repeat every	1	hours	0	mins	5	

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

Job D	etails	
Backu	ip of Resource Group 'ora_01_data' with policy 'Oracle Online Full Backup'	
√ v	Backup of Resource Group 'ora_01_data' with policy 'Oracle Online Full Backup'	
×	ora_01.vmc.local	
~	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	
A Tack	Name: ora_01.vmc.local Start Time: 10/07/2023 8:53:24 AM End Time: 10/07/2023 8:54:33 AM	
o lask	Mame, ora_on/vmc.local start nime, 10/07/2025 6.55:24 AM End nime; 10/07/2023 8:54:33 AM	
	View Logs Cancel Job C	Ilose

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

<			e Database						
	Dashboard		Database	Search databa	ases V				Refresh Resources Ad
9	Resources	19	Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Overall Status
٩	Monitor		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
îîÎ	Reports		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
Å.	Hosts Storage Systems		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
	Settings		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
A	Alerts								

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

\$	Oracle Data	ibase 🔸	cdb3 Topology							
	Search da	atabases						-	± 1	2
	17 14	Name						Backup to Object Stor		otect
٩		cdb1	Manage Copies							
a i		cdb2	22 Backups					Summary C	ard	
		cdb3 cdb4	0 Clones					22 Backups		
<u> </u>		COD4	Local copies					8 Data Back		
80 B								14 Log Backu 0 Clones	ps	
##								0		
A			Primary Backup(s)							
			(search)					• İ Catalog Rena		
			Backup Name	Count	Туре	l₹ End Date	Verified	Mounted	RMAN Cataloged	SCN
			ora_02_10-07-2023_08.05.02.4105_1	Ť	Log	10/07/2023 8:05:26 AM 📋	Not Applicable	False	Not Cataloged	292
			ora_02_10-07-2023_07.50.02.4250_1	1	Log	10/07/2023 7:50:27 AM 🛱	Not Applicable	False	Not Cataloged	292
			ora_02_10-07-2023_07.45.02.4192_1	T	Log	10/07/2023 7:45:49 AM 🛱	Not Applicable	False	Not Cataloged	292
			ora_02_10-07-2023_07.45.02.4192_0	1	Data	10/07/2023 7:45:31 AM 🛱	Unverified	False	Not Cataloged	2923
			ora_02_10-07-2023_07.35.02.3846_1	1	Log	10/07/2023 7:35:25 AM 🛱	Not Applicable	False	Not Cataloged	2920
			ora_02_10-07-2023_07.20.02.3803_1	1	Log	10/07/2023 7:20:25 AM 🛱	Not Applicable	False	Not Cataloged	2925
			ora_02_10-07-2023_07.05.02.3948_1	î	Log	10/07/2023 7:05:26 AM 🛱	Not Applicable	False	Not Cataloged	2924
			ora_02_10-07-2023_06.50.02.3786_1	1	Log	10/07/2023 6:50:26 AM 🛱	Not	False	Not Cataloged	292

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$database;
NAME OPEN MODE
_____
       READ WRITE
CDB3
SQL> show pdbs
   CON ID CON NAME
                                OPEN MODE RESTRICTED
2 PDB$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

3. 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.

Oracle Database 👻	cdb3 Topology							
Search databases							÷ 🔽	=
17 Mame						Backup to Object Store D	utabase Settings Protect	Refres
cdb3	Manage Copies							
cdb4	6 Backups					Summary Card	t	
	0 Clones					6 Backups		
	Local copies					2 Data Backups 4 Log Backups		
						0 Clones		
	Primary Backup(s)							
	Frinary backup(s)							
	(search)					• I Catalog Rename	Gone Restore Mount	
		Count	Туре	17 End Date	Verified			
	(search)	Count	Type Log	LF End Date	Verified Not Applicable	Catalog Rename	Cone Restore Mount	SCN
	(search) Backup Name	Count 1			Not	Catalog Rename Mounted	Clone Restore Mount RMAN Cataloged	5CN 2795205
	Search V Backup Name ora.02.10-06-2023.14.22.59.0383_1	1	Log	10/06/2023 2:23:43 PM 🛱	Not Applicable	Catalog Rename Mounted False	Cone Restore Mount RMAN Cataloged Not Cataloged	2795203 2795113
	Search V Backup Name ora.02.10-06-2023_1422.59.0383_1 ora.02.10-06-2023_1422.59.0383_0	1	Log Data	10/06/2023 2:23:43 PM	Not Applicable Unverified Not	Caulog Rename Mounted False False	Cone Restore Mount RMAN Cataloged Not Cataloged Not Cataloged	2795205 2795205 2795113 2794928
	Teach Teach Backup Name ora.02_10-06-2023_14:22:59.0383_1 ora.02_10-06-2023_14:22:59.0383_0 ora.02_10-06-2023_14:22:59.0383_0 ora.02_10-06-2023_14:22:59.0383_0 ora.02_10-06-2023_14:20:01.8472_1	1	Log Data Log	10/06/2023 2:23:43 PM	Not Applicable Unverified Not Applicable Not	Caulog Rename Mounted False False	Clore Pestore Mount RMAN Cataloged Not Cataloged Not Cataloged Not Cataloged	Uwbouni D

4. Choose restore scope.

Restore cdb3		×
1 Restore Scope	Restore Scope 🚯	
2 Recovery Scope	All Datafiles	
3 PreOps	 Pluggable databases (PDBs) Pluggable database (PDB) tablespaces 	
4 PostOps		2
5 Notification	Database State	
6 Summary	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	

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

Restore cdb3		×
1 Restore Scope	Choose Recovery Scope	
2 Recovery Scope	All Logs Until SCN (System Change Number)	
3 PreOps	SCN 2795205	
4 PostOps	O Date and Time	
5 Notification		
6 Summary	Specify external archive log files locations 🛛 🖸 🧻 🕕	
<i>i</i> After the ope	eration is complete, it is recommended to create a full backup of the Oracle database.	×
		Previous Next

6. Specify any optional pre-scripts to run.

Restore cdb3			×
1 Restore Scope	Specify optional so	cripts to run before performing a restore job 0	
2 Recovery Scope	Prescript full path	/var/opt/snapcenter/spl/scripts/ Enter Prescript path	
3 PreOps	Arguments		
4 PostOps	Script timeout	60 secs	
5 Notification			
6 Summary			
		Previous Nex	đ

7. Specify any optional after-script to run.

Restore cdb3					×
1 Restore Scope	Specify optional sc	ripts to run after performing a	restore job 🕕		
2 Recovery Scope	Postscript full path	/var/opt/snapcenter/spl/scripts/	Enter Postscript path		
3 PreOps	Arguments				
4 PostOps	Open the databas	e or container database in READ-W	RITE mode after recovery		
5 Notification					
6 Summary					
				Previous	t

8. Send a job report if desired.

Restore cdb3			
1 Restore Scope	Provide email set	tings 0	
2 Recovery Scope	Email preference	Never •	
3 PreOps	From	From email	
4 PostOps	То	Email to	
PostOps	Subject	Notification	
5 Notification	🗌 Attach job report		
6 Summary			
		Previous	Next

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

Restore cdb3		
Restore Scope	Summary	
Recovery Scope	Backup name	ora_02_10-06-2023_14.22.59.0383_0
PreOps	Backup date	10/06/2023 2:23:27 PM
	Restore scope	All DataFiles
PostOps	Recovery scope	Until SCN 2795205
Notification	Auxiliary destination	
Summary	Options	Change database state if necessary , Open the database or container database in READ-WRITE mode after recovery
Summary	Prescript full path	None
	Prescript arguments	
	Postscript full path	None
	Postscript arguments	
	Send email	No
		Previous Finis
		Previous Prins

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

ne	Target	State	Server	State details
al Resource	s			
a.DATA.dg				10-0
	ONLINE	ONLINE	ora_02	STABLE
A.LISTENER.1		THEFT	1222 Box	
	ONLINE	INTERMEDIATE	ora_02	Not All Endpoints Re gistered, STABLE
.LOGS.dg				
	ONLINE	ONLINE	ora_02	STABLE
LOGS_CDB3_				
	ONLINE	ONLINE	ora_02	STABLE
.asm		A177 7117		
	ONLINE	ONLINE	ora_02	Started,STABLE
.ons	OPPT THE	OPPT THE		073 57 5
	OFFLINE	OFFLINE	ora_02	STABLE
	5492128			
ster Resour	C C C S			
		INTEDMENTATE	ora 02	Dismounted Mount In:
.cdb3.db 1		INTERMEDIATE	ora_02	
		INTERMEDIATE	ora_02	tiated,HOME=/u01/app
		INTERMEDIATE	ora_02	<pre>tiated,HOME=/u01/app /oracle/product/19.0</pre>
1		INTERMEDIATE	ora_02	tiated,HOME=/u01/app
l .cdb4.db	ONLINE			<pre>tiated,HOME=/u01/app /oracle/product/19.0 .0/cdb3,STABLE</pre>
1			ora_02 ora_02	<pre>tiated,HOME=/u01/app /oracle/product/19.0 .0/cdb3,STABLE Open,HOME=/u01/app/0</pre>
l .cdb4.db	ONLINE			<pre>tiated,HOME=/u01/app /oracle/product/19.0 .0/cdb3,STABLE Open,HOME=/u01/app/o racle/product/19.0.0</pre>
l .cdb4.db l	ONLINE			<pre>tiated,HOME=/u01/app /oracle/product/19.0 .0/cdb3,STABLE Open,HOME=/u01/app/0</pre>
l .cdb4.db l	ONLINE	ONLINE	ora_02	<pre>tiated,HOME=/u01/app /oracle/product/19.0 .0/cdb3,STABLE Open,HOME=/u01/app/o racle/product/19.0.0</pre>
l .cdb4.db l .cssd l	ONLINE			<pre>tiated,HOME=/u01/app /oracle/product/19.0 .0/cdb3,STABLE Open,HOME=/u01/app/o racle/product/19.0.0 /cdb4,STABLE</pre>
l .cdb4.db l .cssd l	ONLINE ONLINE ONLINE	ONLINE	ora_02	<pre>tiated,HOME=/u01/app /oracle/product/19.0 .0/cdb3,STABLE Open,HOME=/u01/app/o racle/product/19.0.0 /cdb4,STABLE</pre>
l .cdb4.db l .cssd l.diskmon l	ONLINE ONLINE ONLINE OFFLINE	ONLINE	ora_02	<pre>tiated,HOME=/u01/app /oracle/product/19.0 .0/cdb3,STABLE Open,HOME=/u01/app/d racle/product/19.0.0 /cdb4,STABLE STABLE</pre>
l .cdb4.db l .cssd l.diskmon l	ONLINE ONLINE ONLINE OFFLINE	ONLINE	ora_02 ora_02	<pre>tiated,HOME=/u01/app /oracle/product/19.0 .0/cdb3,STABLE Open,HOME=/u01/app/0 racle/product/19.0.0 /cdb4,STABLE STABLE</pre>
a.cdb4.db l a.cssd l a.diskmon l a.driver.afd	ONLINE ONLINE ONLINE OFFLINE	ONLINE ONLINE OFFLINE	ora_02	Open,HOME=/u01/app/o racle/product/19.0.0 /cdb4,STABLE STABLE STABLE

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

Restore	e 'ora_02.vmc.local\cdb3'	
~ .	Restore 'ora_02.vmc.local\cdb3'	
~	ora_02.vmc.local	
~	Prescripts	
~	Mount log backups	
~	Pre Restore	
~	► Restore	
4	Post Restore	
~	Unmount log backups	
~	Postscripts	
~	Post Restore Cleanup	
~	▶ Data Collection	
~	Send EMS Messages	
Task Na	ame: ora_02.vmc.local Start Time: 10/06/2023 3:29:27 PM End Time: 10/06/2023 3:48:0	4 PM
	View Logs	ancel Job Close

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 -
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
_____
DТ
_____
_____
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.

Oracle Database	•	cdb3 Topology							
Search databas	ses								V Xect P
15 let 1	Name	_					Backup to Object Ston		xect
	:db3	Manage Copies							
c	:db4	19 Backups					Summary C	ard	
		0 Clones					19 Backups		
		Local copies					6 Data Backu 13 Log Backup		
							0 Clones		
		Primary Backup(s)							
		search 🛛					• ž Catalog Renar	e Clone Restore Mour	d Unmani
		Backup Name	Count	Туре	17 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	28 <mark>1</mark> 4539
		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	ĩ	Log	10/06/2023 4:50:25 PM 🛱	Not Applicable	False	Not Cataloged	281238
		ora_02_10-06-2023_16.45.02.2685_1	1	Log	10/06/2023 4:45:45 PM 🛱	Not Applicable	False	Not Cataloged	281204
		ora_02_10+06+2023_16:45:02:2685_0	1	Data	10/06/2023 4:45:30 PM 🛱	Unverified	False	Not Cataloged	281199

2. Name the clone database SID.

Clone from cdb	3		×
1 Name	Complete Databas	e Clone	
2 Locations	Clone SID	cdb3tst	
3 Credentials	Exclude PDBs	Type to find PDBs	
4 PreOps	O PDB Clone		
(5) PostOps			
6 Notification			
(7) Summary			
		Previous	Next

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

Name								
Name	Select the host to create	a clone						
2 Locations	Clone host ora_	02.vmc <mark>.lo</mark> cal			•			
3 Credentials	O Datafile locations ()							
4 PreOps	+SC_2090922_cdb3tst	+SC_2090922_cdb3tst						
5 PostOps							ja <u>1</u>	
5 Notification	⊙ Control files ①							
7 Summary		+SC_2090922_cdb3tst/cdb3tst/control/control01.ctl × +SC_2090922_cdb3tst/cdb3tst/control/control02.ctl ×						
	⊙ Redo logs 🚯							
	Group		Size	Unit	Number of files			
	 RedoGroup 1 	×	200	MB	2		+	
	RedoGroup 2	×	200	MB	2	+	Reset	
	 RedoGroup 3 	×	200	MB	2	+		

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

Clone from cd	b3		×
1 Name	Database Credentials fo	r the clone	
2 Locations	Credential name for sys user	None • • •	
3 Credentials	ASM instance Credential name	None • + ()	
4 PreOps	Database port	1521	
5 PostOps	ASM Port	1521	
6 Notification	Oracle Home Settings	0	
7 Summary	Oracle Home	/u01/app/oracle/product/19.0.0/cdb4	
	Oracle OS User	oracle	
	Oracle OS Group	oinstall	
		Previous	

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

Clone from cd	b3				×
1 Name	Specify scripts to r	run before clone oper	ation 🕕		
 2 Locations 3 Credentials 4 PreOps 	Prescript full path Arguments Script timeout	/var/opt/snapcenter/s	pl/scripts/ Enter Prescript path		
	🖸 Database Parame	ter settings			
5 PostOps	remote_login_passi	wordfile	320 EXCLUSIVE	× *	
6 Notification	sga_target		2048M	× Reset	
(7) Summary	undo_tablespace		UNDOTBS1	×	
				Previous	đ

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

1 Name	🖉 Recover Database
 2 Locations 3 Credentials 4 PreOps 5 PostOps 6 Notification 7 Summary 	 Until Cancel Date and Time Date-time format: MM/DD/YYYY hh:mm:ss Until SCN (System Change Number) Specify external archive log locations Image: Image: /li>
	 Create new DBID Create tempfile for temporary tablespace Enter SQL queries to apply when clone is created Enter scripts to run after clone operation

7. Review the summary and launch the clone job.

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

8. Monitor the clone job execution from Monitor tab.

lione	from backup 'ora_02_10-06-2023_16.45.	U2.2065_U
✓	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

9. Cloned database is immediately registered in SnapCenter.

		atabase 👻						
Dashboard	rd View	Database	Search databases	V				Refresh Resources
Resources	s 17 🏴	Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Overall Status
Monitor		cdb1	Single Instance (Multitenant)	ora_01.vmc.local				Not protected
		cdb2	Single Instance (Multitenant)	ora_01.vmc.local				Not protected
Reports Hosts		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 succeed
I Storage Sy	ystems	cdb3tst	Single Instance (Multitenant)	ora_02.vmc.local				Not protected
Settings		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 succeed
Alerts								

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 O2 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 -
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

https://aws.amazon.com/fsx/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

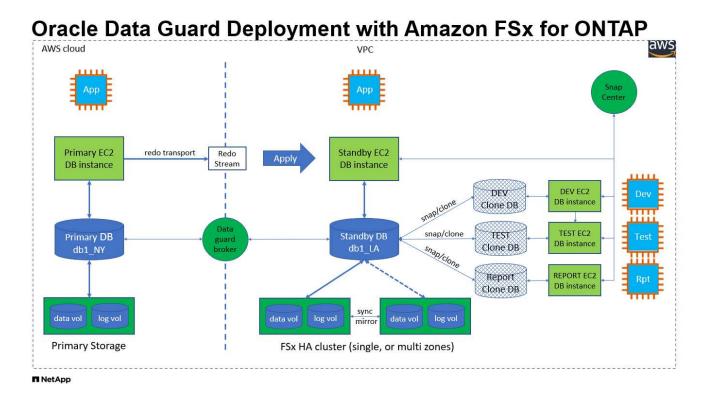
This solution is intended for the following people:

- 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



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 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 alway 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 actually 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.dq
           ONLINE ONLINE
                         ip-172-30-15-45
                                             STABLE
ora.LISTENER.lsnr
          ONLINE ONLINE ip-172-30-15-45
                                             STABLE
ora.LOGS.dq
          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.db1.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.
- 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)
   )
 )
                                                       # line added
ENABLE GLOBAL DYNAMIC ENDPOINT LISTENER=ON
by Agent
                                                      # line added
VALID NODE CHECKING REGISTRATION LISTENER=ON
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.
```

 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 -
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 string db name db1 string db unique name db1 LA boolean global names 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 SOL> SQL> show parameter log archive config TYPE VALUE NAME _____ ___ _____ 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/undotbs1.270.

NAME

6248293

_____ _____ +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;

		ILE/temp.287.1146248371
_		01A66EE9797E0632D0F1EAC5F59/TEMPFILE/temp.288.1140
248375		1110011979710002201111001097111111111111
	LA/03C5E!	FD07C41A1FAE0632D0F1EAC9BD8/TEMPFILE/temp.290.114
248463	·	
+DATA/DB1	LA/03C5F	ODDF35CA2B6E0632D0F1EAC8B6B/TEMPFILE/temp.291.114
248463 —		-
+DATA/DB1	LA/03C5F	1C9B142A2F1E0632D0F1EACF21A/TEMPFILE/temp.292.114
248463 —		
SQL> select	c group#,	, type, member from v\$logfile order by 2, 1;
GROUP#	TYPE	MEMBER
		+LOGS/DB1_LA/ONLINELOG/group_1.259.1146248349
	ONLINE	+DATA/DB1_LA/ONLINELOG/group_1.280.1146248347
	ONLINE	+DATA/DB1_LA/ONLINELOG/group_2.281.1146248351
	ONLINE	+LOGS/DB1_LA/ONLINELOG/group_2.258.1146248353
	ONLINE	+DATA/DB1_LA/ONLINELOG/group_3.282.1146248355
	ONLINE	+LOGS/DB1_LA/ONLINELOG/group_3.260.1146248355
		+DATA/DB1_LA/ONLINELOG/group_4.283.1146248357
		+LOGS/DB1_LA/ONLINELOG/group_4.261.1146248359
		+DATA/DB1_LA/ONLINELOG/group_5.284.1146248361
		+LOGS/DB1_LA/ONLINELOG/group_5.262.1146248363
		+LOGS/DB1_LA/ONLINELOG/group_6.263.1146248365
		+DATA/DB1_LA/ONLINELOG/group_6.285.1146248365
		+LOGS/DB1_LA/ONLINELOG/group_7.264.1146248369
7	STANDBY	+DATA/DB1_LA/ONLINELOG/group_7.286.1146248367
1 / .	lected.	
14 rows se		
	t name, 🤉	open_mode from v\$database;
SQL> select	OPEN MODI	

9. 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.

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.			

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 instantaniouly.

Clone standby databse 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 relavant 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.

	oCenter®				🏚 🔤 🚱 - 👤 ad	ministrator SnapCenterAdmin 🖡 Sign Ou
	ONTAP Storage					
Dashboard	Type ONTAP SVMs Search b	by Name				t 📩
Resources	ONTAP Storage Connections					
Monitor	Name	IE IP	Cluster Name	User Name	Platform	Controller License
Reports	svm.ora		Ip-172-30-15-25.ec2.internal		F5x	Not applicable
Hosts						
Storage Systems						
E Settings						

3. Add AWS ec2-user to $\ensuremath{\mathsf{Credential}}$ in Settings.

Image: Section 2 Image: Section 2 <th>← C (▲ N</th> <th>Not secure https://snapctr:8146/Administration?ViewName=RunAs</th> <th></th> <th>A th th th de de de de de</th>	← C (▲ N	Not secure https://snapctr:8146/Administration?ViewName=RunAs		A th th th de de de de de
Control Search by Credential Hame Authentication Mode Details P Recourses Fedential Name Authentication Mode Details Monitor e2/sure Unux User/dec2/sure Monitor e2/sure Unux User/dec2/sure Monitor e2/sure User/dec2/sure User/dec2/sure Monitor <th>n NetApp Snap</th> <th>pCenter®</th> <th></th> <th>🌲 🔤 😯 - 👤 administrator SnapCenterAdmin 🖡 Sign O</th>	n NetApp Snap	pCenter®		🌲 🔤 😯 - 👤 administrator SnapCenterAdmin 🖡 Sign O
Resources Credential Atlant Authentication Mode Details Monitor ec2-user Unux User/dec2-user Nonitor ec2-user User/dec2-user User/dec2-user Nonitor ec2-user User/dec2-user User/dec2-user Nonitor ec2-user User/dec2-user User/dec2-user Nonitor ec2-user User/dec2-user User/dec2-user Stetlings ec2-user User/dec2-user User/dec2-user		Global Settings Policies Users and Access Roles Credential		
Nonitor e2user Linux Userdie:2user Monitor Anoritor Linux Userdie:2user Reports	Dashboard	Search by Credential Name		
Image: Settings	Resources	Credential Name	Authentication Mode	Details
House Systems Settings	Honitor	ec2-user	Linux	Userld:ec2-user
Storage Systems	Reports			
Settings	🛔 Hosts			
	- Storage Systems	s		
A Mers	Settings			
	Alerts			

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

\leftarrow	C 🔺 Not	t secure	https://snapctr:8146/Host						A [™] ☆	\$ \$ \$	
	NetApp Snap(Center@	D					• =	🥹 🔹 administra	tor SnapCenterAdmin	🚺 Sign Out
<		Manag	ed Hosts Disks Shares								
	Dashboard	Searc	th by Name						Add		More
	Resources		Name	14.	Туре	System	Plug-in		Version	Overall Status	
-	Monitor		ip-172-30-15-126.ec2.internal		Linux	Stand-alone	UNIX, Oracle Database		4.9	Running	
a i	Reports		ip-172-30-15-67.ec2.internal		Linux	Stand-alone	UNIX, Oracle Database		4.9	Running	
٨	Hosts										
÷,	Storage Systems										
	Settings										
▲	Alerts										

(i)

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.

	Center®			♠ ⊠ ⊕·	Ladministrator SnapCenterAdmin Sign Out
	Global Settings Policies Users and Access Oracle Database				
Dashboard	Search by Name			+	/ 5 0 5
Resources		1.		New	Modily Copy Details Delote
Monitor	Name 1		Schedule Type	Replication	Verification
Reports	Oracle full DB backup	DATA, OFFLINEMOUNT	Hourly		
Hosts					
Storage Systems					
E Settings					

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

Refresh Resources
Backup succeeded

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.

	abase 💌	db1 Topology								
Search d				-						
	latabases				· 💼	=			🔧 🕕	Refres
15.16	Name			Backup to	Object Store Remove Prote	ction Database Se	ttings Back up Now	Modify Mai	ntenance Details	Refres
	db1	Manage Copies								
		3 Backups					5	ummary Card		
		0 Clones						Backups		
		Local copies						3 Data Backups		
								0 Log Backups Clones		
								ciones		
		Primary Backup(s)								
		(search)							ne Reitore Mount 1	A Inmount
		Backup Name	Count	Туре	17	End Date	Verified	Mounted	RMAN Cataloged	sc
		db1_LA_08-31-2023_17.42.01.6804_0	1	Data	08/31/2023 5:	42:29 PM 📋	Unverified	False	Not Cataloged	
		db1_LA_08-31-2023_16.30.01.6158_0	1	Data	08/31/2023 4:	30:29 PM 🛱	Unverified	False	Not Cataloged	
		db1_LA_08-31-2023_15.59.09.6092_0	1	Data	08/31/2023 3:	59:42 PM 🛱	Unverified	False	Not Cataloged	

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

1		×
Complete Database	Clone	
Clone SID	db1dev	
Exclude PDBs	Type to find PDBs	
O PDB Clone		
	Provinite	lext
	Complete Database Clone SID Exclude PDBs	• Complete Database Clone Lone SID dbdee bottom Database Clone Dbdee

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.

Name	Select the host to create	a clone					
Locations	Clone host ip-17	2-30- <mark>1</mark> 5-126.e	c2.internal				
Credentials	⊙ Datafile locations ()						
PreOps	+SC_2090922_db1dev					•	
PostOps	+SC_2342319_db1dev						Reset
Notification	⊙ Control files ①						
Summary	+SC_2090922_db1dev/db1	dev/control/co	introl01.ctl			× *	+
	+SC_2090922_db1dev/db10	dev/control/co	ntrol02.ctl			×	Reset
	🛇 Redo logs 🚯						
	Group		Size	Unit	Number of files		
	RedoGroup 1	$\left \times \right $	200	MB	2	+	Î +
	RedoGroup 2	×	200	MB	2	+	Reset
	 RedoGroup 3 	×	200	MB	2	+	•
		U.T. D.		L	L	L.I1	

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 db	1		>
1 Name	Database Credentials fo	r the clone	
2 Locations	Credential name for sys user	None • •	
3 Credentials	ASM instance Credential name	None - + O	
4 PreOps	Database port	1521	
5 PostOps	ASM Port	1521	
6 Notification	Oracle Home Settings	0	
7 Summary	Oracle Home	/u01/app/oracle/product/19.0.0/dev	
	Oracle OS User	oracle	
	Oracle OS Group	oinstall	
		Previous	Next
		Flevious	NEXL

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

	Specify scripts to run before clone operatio	on 🚯	
Locations	Prescript full path /var/opt/snapcenter/spl/s	cripts/ Enter Prescript path	
Credentials	Arguments		
PreOps	Script timeout 60 secs		
	Database Parameter settings		
PostOps	audit_file_dest	/u01/app/oracle/admin/db1dev_LA/adump	× +
Notification	audit_trail	DB	× +
	open_cursors	300	× Reset
Summary	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.

1 Name	the second as the second as the second as second as the second second as the second second second second second	
Name	Until Cancel recovery will be performed for Physical Standby Dataguard/Active Dataguard database.	
2 Locations	Create new DBID 🚯	
3 Credentials	 Create tempfile for temporary tablespace Enter SQL queries to apply when clone is created 	
4 PreOps	shutdown immediate ; startup mount ; alter database noarchivelog ; alter database open ;	+ Reset
5 PostOps		nesce
6 Notification	 Enter scripts to run after clone operation 	
7 Summary		

13. Configure email notification if desired.

Clone from d	,		
1 Name	Provide email set	tings 🕦	
2 Locations	Email preference	Never •	
3 Credentials	From	From email	
4 PreOps	То	Email to	
	Subject	Notification	
5 PostOps	🗌 Attach job report		
6 Notification			
7 Summary			

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

1 Name	Summary	
2 Locations	Clone from backup	db1_LA_08-31-2023_17.42.01.6804_0
3 Credentials	Clone SID	db1dev
Credentials	Clone server	ip-172-30-15-126.ec2.internal
4 PreOps	Exclude PDBs	none
•	Oracle home	/u01/app/oracle/product/19.0.0/dev
5 PostOps	Oracle OS user	oracle
6 Notification	Oracle OS group	oinstall
7 Summary	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/redolog/redo01_01.log RedoGroup =1 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redolog/redo02_01.log RedoGroup =2 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redolog/redo02_02.log RedoGroup =3 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redolog/redo03_01.log RedoGroup =3 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redolog/redo03_01.log RedoGroup =3 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redolog/redo03_02.log RedoGroup =4 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redolog/redo03_02.log RedoGroup =4 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redolog/redo04_01.log RedoGroup =5 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redolog/redo04_02.log RedoGroup =5 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redolog/redo05_01.log RedoGroup =5 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redolog/redo05_02.log RedoGroup =6 TotalSize =200 Path =+SC_2090922_db1dev/db1dev/redolog/redo06_02.log

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.

~ *	Clone from backup 'db1_LA_08-31-2023_17.42.01.6804_0'	
~	ip-172-30-15-126.ec2.internal	
1	Prescripts	
~	Query Host Information	
~	Prepare for Cloning	
~	Cloning Resources	
~	FileSystem Clone	
~	Application Clone	
~	Postscripts	
4	▶ Register Clone	
~	Data Collection	
~	Send EMS Messages	
	2 Rectification and Brow	
A Table	Name: ip-172-30-15-126.ec2.internal Start Time: 08/31/2023 6:02:46 PM End Time:	00/21/2022 6-11-27 DM

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

7			abase 👻						
	Dashboard	View D	atabase	Search database	5 7				efresh Resources Add
9	Resources	17.14	Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Overall Status
	Monitor		db1	Single Instance Physical Standby (Multitenant)	ip-172-30-15-67.ec2.internal		Oracle full DB backup	08/31/2023 5:42:28 PM 🛅	Backup succeeded
ñî	Reports	-	db1dev	Single Instance Physical Standby (Multitenant)	ip-172-30-15-126.ec2.internal				Not protected
4	Hosts			Standby (Multiteriant)					
ł	Storage Systems								
=	Settings								
A	Alerts								

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 -
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=Cj0KCQiA54 KfBhCKARIsAJzSrdqwQrghn6I71jiWzSeaT9Uh1-vY-VfhJixFxnv5rWwn2S7RqZOTQ0aAh7eEALw_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

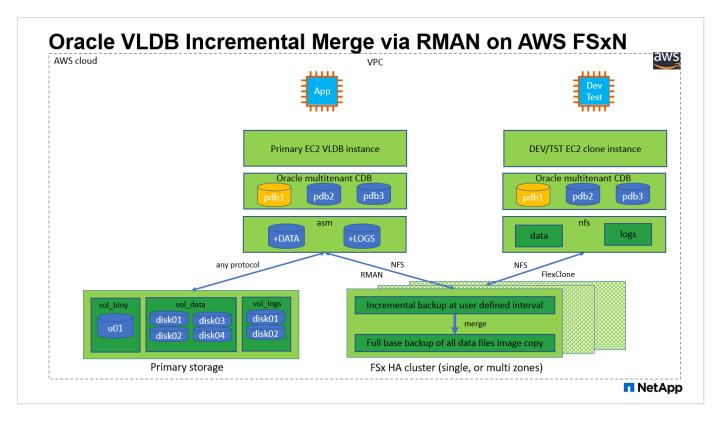
This solution is intended for the following people:

- 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



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 alway 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 actually 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.

	2 , Search	[Alt+S]	Q N. Virgir
Resource Groups & Tag Edit	tor 🔯 FSx		
Amazon FSx	×	Amazon FSx for NetApp ONTAP Amazon FSx for OpenZFS	
File systems			
Volumes			
Caches		File system details	
Backups			
ONTAP		File system The file system where this volume will be created.	
Storage virtual machines	i i	ONTAP fs-06c3c8b2a7bd56458 fsx_01	
OpenZFS		Storage virtual machine The storage virtual machine that will host this volume.	
Snapshots		svm-07915e6cff12c1e1e svm_ora	
		Volume details	
		Volume name	
		ora_01_copy	
		Maximum of 203 alphanumeric characters, plus	
		Volume size	
		102400	
		Minimum 20 MiB; Maximum 104857600 MiB	
		Volume type Select whether you're creating a Read-Write (RW) volume or a read-only Data Protection (DP) volume, which is used v	vith SnapMirror.
		• Read-Write (RW)	
		O Data Protection (DP)	
		Junction path The location within your file system where your volume will be mounted.	

 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,wsize=262144,noin
tr
```

7. Change mount point ownership to oracle:oisntall, 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

 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> {
   allocate channel c1 device type disk format
3>
'/nfsfsxn/oracopy/%U';
   allocate channel c2 device type disk format
4>
'/nfsfsxn/oracopy/%U';
   allocate channel c3 device type disk format
5>
'/nfsfsxn/oracopy/%U';
6>
   allocate channel c4 device type disk format
'/nfsfsxn/oracopy/%U';
7> delete obsolete;
   sql 'alter system archive log current';
8>
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
1 A 17-MAY-23 3009819 17-MAY-23
19
                                                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 0ilsd7at
      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$SEED
26
      6 A 17-MAY-23 2383520 12-MAY-23 NO
      Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-
SYSAUX FNO-6 001sd7c8
      Tag: ORACOPYBKUPONFSXN LEVEL 0
      Container ID: 2, PDB Name: PDB$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 011sd7bi 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 17 A 17-MAY-23 3009895 17-MAY-23 30 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 18 A 17-MAY-23 24 3009858 17-MAY-23 NO Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SYSAUX FNO-18 0mlsd7bq 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 Oulsd7de Tag: ORACOPYBKUPONFSXN LEVEL 0 Container ID: 5, PDB Name: DB1 PDB3 20 A 17-MAY-23 3009914 17-MAY-23 37 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 21 A 17-MAY-23 3009019 17-MAY-23 4 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 22 A 17-MAY-23 3009419 17-MAY-23 5 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 23 A 17-MAY-23 3009460 17-MAY-23 6 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

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 25 A 17-MAY-23 3009502 17-MAY-23 8 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 26 A 17-MAY-23 3009548 17-MAY-23 NO 9 Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SOE FNO-26 071sd6vf Tag: ORACOPYBKUPONFSXN LEVEL 0 Container ID: 3, PDB Name: DB1 PDB1 27 A 17-MAY-23 3009576 17-MAY-23 10 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 091sd711 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 0alsd720 Tag: ORACOPYBKUPONFSXN LEVEL 0 Container ID: 3, PDB Name: DB1 PDB1 30 A 17-MAY-23 3009648 17-MAY-23 13 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 Oc1sd74u 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 A 17-MAY-23 3009771 17-MAY-23 17 34 NO Name: /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-34 Of1sd788 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>
```

8. 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
_____ ____
                                               _____
  1060
1
           SYSTEM
                              YES
+DATA/DB1/DATAFILE/system.257.1136666315
           SYSAUX
   810
3
                              NO
+DATA/DB1/DATAFILE/sysaux.258.1136666361
    675
           UNDOTBS1
4
                              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 DB1 PDB2:USERS NO 16 5 +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 DB1 PDB1:SOE 22 4096 NO +DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.287.11370183 11 23 4096 DB1 PDB1:SOE NO +DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.288.11370183 59 24 DB1 PDB1:SOE 4096 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 4096 29 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 DB1 PDB1:SOE 31 4096 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 4096 DB1 PDB1:SOE NO 35 +DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/DATAFILE/soe.300.11370191 81

List of Temporary Files _____ File Size(MB) Tablespace Maxsize(MB) Tempfile Name _____ _____ 123 1 TEMP 32767 +DATA/DB1/TEMPFILE/temp.265.1136666447 123 PDB\$SEED:TEMP 32767 2 +DATA/DB1/FB864A929AEB79B9E053630F1EAC7046/TEMPFILE/temp.269.1136667 185 3 10240 DB1 PDB1:TEMP 32767 +DATA/DB1/FB867DA8C68C816EE053630F1EAC2BCF/TEMPFILE/temp.274.1136668 051 DB1 PDB2:TEMP 32767 4 123 +DATA/DB1/FB867EA89ECF81C0E053630F1EACB901/TEMPFILE/temp.279.1136668 067 5 123 DB1 PDB3:TEMP 32767 +DATA/DB1/FB867F8A4D4F821CE053630F1EAC69CC/TEMPFILE/temp.284.1136668 081

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 091sd711
-rw-r---- 1 oracle asm 4294975488 May 17 18:14 data D-DB1 I-
1730530050 TS-SOE FNO-29 0a1sd720
-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 011sd7bi
-rw-r---- 1 oracle asm 492838912 May 17 18:18 data D-DB1 I-
1730530050 TS-SYSAUX FNO-18 0mlsd7bq
-rw-r---- 1 oracle asm 849354752 May 17 18:18 data D-DB1 I-
1730530050 TS-SYSAUX FNO-3 0ilsd7at
-rw-r---- 1 oracle asm 482353152 May 17 18:18 data D-DB1 I-
1730530050 TS-SYSAUX FNO-6 0olsd7c8
-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 0ulsd7de
-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 111sd7dm
                         5251072 May 17 18:19 data D-DB1 I-
-rw-r---- 1 oracle asm
1730530050 TS-USERS FNO-16 121sd7dn
-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;
______D______
DT ______DT _______
EVENT __________
1 18-MAY-23 02.35.37.000000 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 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 001sd7c8" datafile 7 switched to datafile copy "/nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-USERS FNO-7 101sd7dl" 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 Om1sd7bq" datafile 19 switched to datafile copy "/nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-UNDOTBS1 FNO-19 0ulsd7de" 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 0alsd720"
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>
```

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 091sd711
/nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SOE FNO-29 0a1sd720
/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_011sd7bi /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_001sd7c8 /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-17_0s1sd7d4 /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-1_0h1sd7ae /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_0qlsd7cm /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-11_0nlsd7c1 /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-15_0tlsd7db /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-4_0jlsd7de /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-4_0jlsd7b4 /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-8_0vlsd7di /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-USERS_FNO-12_11lsd7dm /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-20_131sd7d1

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
                               READ WRITE NO
      3 DB1 PDB1
      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 -
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.
SOL> commit;
Commit complete.
SQL> select * from test;
```

```
ID

DT

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. Privison a new EC2 DB instance host ora_02 with same OS and version via AWS EC2 console. Configure OS kernal 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

- Configure oracle environment similiarly to primary EC2 DB instance host ora_01, such as oratab, oraInst.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 provided 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,wsize=262144,noin
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 4mlt508t
-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
                             663552 May 30 15:29 6c1tf6b8 204 1 1
-rw-r----. 1 oracle 54331
-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 6q1tf6b9 208 1 1
-rw-r----. 1 oracle 54331
                              589824 May 30 15:29 611tf6bb 213 1 1
-rw-r----. 1 oracle 54331
                              606208 May 30 15:29 6mltf6bb 214 1 1
                              368640 May 30 15:29 6oltf6bb 216 1 1
-rw-r----. 1 oracle 54331
-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 301t4ut2 -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 401t509m

```
-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
                              303104 May 30 17:33 7c1tfdiu 236 1 1
-rw-r----. 1 oracle 54331
-rw-r----. 1 oracle 54331
                              319488 May 30 17:33 7d1tfdiv 237 1 1
-rw-r----. 1 oracle 54331
                              57344 May 30 17:33 7f1tfdiv 239 1 1
-rw-r----. 1 oracle 54331
                              57344 May 30 17:33 7g1tfdiv 240 1 1
-rw-r----. 1 oracle 54331
                              57344 May 30 17:33 7h1tfdiv 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 sequency number. In this case, it is 175. Our recovery point is up to log sequency 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 ol mf 1 143 O2rotwyb .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 ol 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 03mfx17v .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 ol 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 ol mf 1 153 040x946w .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 ol mf 1 155 04tv1yvn .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 ol 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 057146my .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 ol 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 05pmg741 .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 0631tvgv .arc
-r--r-. 1 oracle 54331 164455424 May 30 16:47
o1 mf 1 172 065d94fg .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
                               9177880 bytes
Fixed Size
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 4flt506m
       Tag: ORACOPYBKUPONFSXN LEVEL 0
322
          A 30-MAY-23
                          4120175 30-MAY-23 NO
      3
      Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-
SYSAUX FNO-3 4g1t506m
       Tag: ORACOPYBKUPONFSXN LEVEL 0
      4 A 30-MAY-23 4120179 30-MAY-23 NO
317
       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
      7 A 30-MAY-23 4120207 30-MAY-23 NO
323
      Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-USERS FNO-
7 4u1t50a6
      Tag: ORACOPYBKUPONFSXN LEVEL 0
       8 A 26-MAY-23 2383520 12-MAY-23
227
                                               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 4n1t509m Tag: ORACOPYBKUPONFSXN LEVEL 0 Container ID: 3, PDB Name: DB1 PDB1 10 A 30-MAY-23 4120166 30-MAY-23 307 NO Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SYSAUX 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 411t508t 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 401t509m Tag: ORACOPYBKUPONFSXN LEVEL 0 Container ID: 4, PDB Name: DB1 PDB2 14 A 30-MAY-23 4120183 30-MAY-23 NO 318 Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SYSAUX 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 16 A 30-MAY-23 4120211 30-MAY-23 325 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-SYSAUX 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

24 A 30-MAY-23 4120162 30-MAY-23 306 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 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 30 A 30-MAY-23 4120154 30-MAY-23 309 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 31 A 30-MAY-23 4120158 30-MAY-23 NO 301 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
            34 A 30-MAY-23
    304
                                  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 4mlt508t" 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 411t508t" 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 401t509m" 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 301t4ut2" 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
```

С 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 С 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 С 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 С 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 С 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 С 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 С 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 03mfx17v .ar С 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 С 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 С 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 С 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 С 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 С 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 04tv1yvn .ar С 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 С 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 С 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 С 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 С 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 С 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 С 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 С 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 05h8lm1h .ar С 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 С 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

С 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 С 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 С 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 С 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 С 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 С 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 С 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 065d94fg .ar С 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 С 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 С 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 С 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 04tv1yv 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 05h8lm1
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 05s3001
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 0631tvg
v .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
   READ ONLY NO
     2 PDB$SEED
      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
```

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,wsize=262144,noin
tr 0 0

Update the Oracle init file from primary databse 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% _____ _ ____ _____ ____ svm ora ora 01 copy daily.2023-06-02 0010 3.59GB 28 5% daily.2023-06-03 0010 1.10GB 18 18 daily.2023-06-04 0010 608KB 08 08 daily.2023-06-05 0010 3.81GB 28 5% one-off.2023-06-05-1137 168KB 08 08 svm ora root weekly.2023-05-28 0015 1.86MB 0% 78% daily.2023-06-04 0010 152KB 08 22% weekly.2023-06-04 0015 1.24MB 08 70% daily.2023-06-05 0010 196KB 0% 27% hourly.2023-06-05 1005 156KB 08 22% hourly.2023-06-05 1105 156KB 0% 228 hourly.2023-06-05 1205 156KB 08 228 hourly.2023-06-05 1305 156KB 08 22% hourly.2023-06-05 1405 1.87MB 08 78% hourly.2023-06-05 1505 148KB 08 228 15 entries were displayed.

 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.

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,rsize=262144,wsize=262144,noin
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 8q1t105i 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 8r1t105m 283 1 1
-rw-r---- 1 oracle 54331 139264 Jun 1 20:21 8t1t105t 285 1 1
-rw-r---- 1 oracle 54331
                          3514368 Jun 1 20:21 8s1t105t 284 1 1
-rw-r---- 1 oracle 54331
                           139264 Jun 1 20:21 8ultl060 286 1 1
-rw-r---- 1 oracle 54331
                           425984 Jun 1 20:21 901t1062 288 1 1
-rw-r---- 1 oracle 54331
                           344064 Jun 1 20:21 911t1062 289 1 1
-rw-r---- 1 oracle 54331
                           245760 Jun 1 20:21 931t1063 291 1 1
-rw-r---- 1 oracle 54331
                           237568 Jun 1 20:21 941t1064 292 1 1
-rw-r---- 1 oracle 54331
                            57344 Jun 1 20:21 961t1065 294 1 1
-rw-r---- 1 oracle 54331
                            57344 Jun 1 20:21 971t1066 295 1 1
-rw-r---- 1 oracle 54331 57344 Jun 1 20:21 981t1067 296 1 1
-rw-r---- 1 oracle 54331 1040760832 Jun 1 20:23 8mltkvv2 278 1 1
-rw-r---- 1 oracle 54331 932847616 Jun 1 20:24 801tkvv2 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 8altkrhr
-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 8eltkril -rw-r---- 1 oracle 54331 246423552 Jun 5 15:21 data D-DB1 I-1730530050 TS-UNDOTBS1 FNO-19 8f1tkrj4 5251072 Jun 5 15:21 data D-DB1 I--rw-r---- 1 oracle 54331 1730530050 TS-USERS FNO-7 8h1tkrj9 5251072 Jun 5 15:21 data D-DB1 I--rw-r---- 1 oracle 54331 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 5251072 Jun 5 15:21 data D-DB1 I--rw-r---- 1 oracle 54331 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 7j1tkgk6 -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 711tkqk6 -rw-r---- 1 oracle 54331 4294975488 Jun 5 15:22 data D-DB1 I-1730530050 TS-SOE FNO-27_7pltkqrq -rw-r---- 1 oracle 54331 4294975488 Jun 5 15:22 data D-DB1 I-1730530050 TS-SOE FNO-35 811tkrap -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 7mltkqk6 -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 701tkqrj
-rw-r---- 1 oracle 54331 1241432064 Jun 5 15:30 9d1tv06n 301 1 1
-rw-r---- 1 oracle 54331 1019805696 Jun 5 15:31 9altv06m 298 1 1
-rw-r---- 1 oracle 54331 4612096 Jun 5 15:31 9e1tv0ld 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 9g1tv0lt_304_1_1
-rw-r---- 1 oracle 54331
                           319488 Jun 5 15:31 9h1tv0lt 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
                           581632 Jun 5 15:31 911tv0m5 309 1 1
-rw-r---- 1 oracle 54331
-rw-r---- 1 oracle 54331 54345728 Jun 5 15:31 9f1tv0lt 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 901tv0m6 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 -1
/nfsfsxn/archlog/DB1/archivelog/2023 06 05
total 2008864
-rw-r---- 1 oracle 54331 729088 Jun 5 14:38
o1 mf 1 190 17vwvvt9 .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
ol mf 1 192 l7vxctms .arc
-rw-r---- 1 oracle 54331 166868992 Jun 5 14:49
ol mf 1 193 l7vxjjps .arc
-rw-r---- 1 oracle 54331 166087168 Jun 5 14:52
ol mf 1 194 l7vxnxrh .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 l7vxylwp .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
ol 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, lanuch 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/db1.ctl
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 4 A 05-JUN-23 8319171 01-JUN-23 365 NO Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-UNDOTBS1 FNO-4 851tkrqf 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 8d1tkril Tag: ORACOPYBKUPONFSXN LEVEL 0 Container ID: 2, PDB Name: PDB\$SEED 6 A 01-JUN-23 2383520 12-MAY-23 NO 349 Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SYSAUX FNO-6 891tkrhr Tag: ORACOPYBKUPONFSXN_LEVEL_0 Container ID: 2, PDB Name: PDB\$SEED 7 A 05-JUN-23 8319201 01-JUN-23 372 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 10 A 05-JUN-23 8318714 01-JUN-23 NO 376 Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SYSAUX FNO-10 861tkrgo Tag: ORACOPYBKUPONFSXN LEVEL 0 Container ID: 3, PDB Name: DB1 PDB1 11 A 05-JUN-23 8318720 01-JUN-23 377 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 12 A 05-JUN-23 8318719 01-JUN-23 375 NO Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-USERS FNO-12 8iltkrj9 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 15 A 05-JUN-23 8319193 01-JUN-23 370 NO Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-UNDOTBS1 FNO-15 8eltkril Tag: ORACOPYBKUPONFSXN LEVEL 0 Container ID: 4, PDB Name: DB1 PDB2 16 A 05-JUN-23 8319206 01-JUN-23 373 NO Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-USERS FNO-16 8jltkrja Tag: ORACOPYBKUPONFSXN LEVEL 0 Container ID: 4, PDB Name: DB1 PDB2 17 A 05-JUN-23 8319188 01-JUN-23 NO 369 Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SYSTEM FNO-17 8c1tkril Taq: ORACOPYBKUPONFSXN LEVEL 0 Container ID: 5, PDB Name: DB1 PDB3 8319180 01-JUN-23 NO 18 A 05-JUN-23 367 Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SYSAUX FNO-18 881tkrhr Tag: ORACOPYBKUPONFSXN LEVEL 0 Container ID: 5, PDB Name: DB1 PDB3

19 A 05-JUN-23 8319197 01-JUN-23 NO 371 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 8kltkrjb 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 7jltkqk6 Tag: ORACOPYBKUPONFSXN LEVEL 0 Container ID: 3, PDB Name: DB1 PDB1 22 A 05-JUN-23 8318714 01-JUN-23 388 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 23 A 05-JUN-23 384 8318717 01-JUN-23 NO Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SOE FNO-23 711tkgk6 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 7mltkqk6 Tag: ORACOPYBKUPONFSXN LEVEL 0 Container ID: 3, PDB Name: DB1 PDB1 25 A 05-JUN-23 8318720 01-JUN-23 381 NO Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SOE FNO-25 7n1tkgrh Tag: ORACOPYBKUPONFSXN LEVEL 0 Container ID: 3, PDB Name: DB1 PDB1 26 A 05-JUN-23 8318714 01-JUN-23 NO 392 Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SOE FNO-26 7oltkqrj 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 7pltkqrq Tag: ORACOPYBKUPONFSXN LEVEL 0 Container ID: 3, PDB Name: DB1 PDB1 28 A 05-JUN-23 8318719 01-JUN-23 390 NO Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SOE FNO-28 7q1tkqs1 Tag: ORACOPYBKUPONFSXN LEVEL 0 Container ID: 3, PDB Name: DB1 PDB1 29 A 05-JUN-23 8318720 01-JUN-23 380 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 7sltkr3a 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 32 A 05-JUN-23 387 8318719 01-JUN-23 NO Name: /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SOE FNO-32 7u1tkr42 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 35 A 05-JUN-23 8318714 386 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 851tkrqf" datafile 5 switched to datafile copy "/nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SYSTEM FNO-5 8dltkril" 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 8altkrhr" 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 8eltkril" 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_711tkqk6"
datafile 24 switched to datafile copy "/nfsfsxn/oracopy/data D-
DB1 I-1730530050 TS-SOE FNO-24 7mltkqk6"
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 701tkqrj"
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 17vxylwp
.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 -
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_l81bhz6g_.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_7j1tkqk6 /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-22_7k1tkqk6 /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-23_711tkqk6 /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-24_7m1tkqk6 /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-25_7n1tkqrh /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-26_701tkqrj

NAME

```
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-27_7p1tkqrq
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-28_7q1tkqs1
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-29_7r1tkr32
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-30_7s1tkr3a
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-31_7t1tkr3i
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-32_7u1tkr42
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-33_7v1tkra6
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-34_801tkram
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-35_811tkrap
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-35_811tkrap
/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SOE_FNO-34_801tkram
```

NAME

/nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSAUX_FNO-18_881tkrhr

/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_8b1tkri1 /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-17_8c1tkri1 /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-1_821tkrb8 /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-5_8d1tkri1 /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-SYSTEM_FNO-9_8a1tkrhr /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_8e1tkri1 /nfsfsxn/oracopy/data_D-DB1_I-1730530050_TS-UNDOTBS1_FNO-19_8f1tkrj4

NAME

_____ _____ /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-UNDOTBS1 FNO-4 851tkrqf /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-UNDOTBS1 FNO-8 8g1tkrj7 /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-USERS FNO-12 8iltkrj9 /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
SOL>
```

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$database;
NAME OPEN_MODE LOG_MODE
_____ ____
DB1 READ WRITE
                           NOARCHIVELOG
SQL> shutdown immediate;
Database closed.
Database dismounted.
ORACLE instance shut down.
SQL> startup mount;
ORACLE instance started.
Total System Global Area 1.0737E+10 bytes
Fixed Size
                         9174800 bytes
                     1577058304 bytes
Variable Size
Database Buffers 9126805504 bytes
Redo Buffers
                        24379392 bytes
Database mounted.
SOL> exit
Disconnected from Oracle Database 19c Enterprise Edition Release
19.0.0.0 - Production
Version 19.18.0.0.0
[oracle@ip-172-30-15-124 dbs]$ 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 8dltkri - 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 8altkrh - 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 8iltkrj - 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 8eltkri - dbid changed, wrote new name
    Datafile /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-USERS FNO-
16 8jltkrj - dbid changed, wrote new name
    Datafile /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SYSTEM FNO-
17 8cltkri - 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 8kltkrj - 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 7kltkqk - dbid changed, wrote new name
    Datafile /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SOE FNO-
23 711tkqk - dbid changed, wrote new name
    Datafile /nfsfsxn/oracopy/data D-DB1 I-1730530050 TS-SOE FNO-
24 7mltkgk - 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 7oltkgr - dbid changed, wrote new name
    Datafile /nfsfsxn/oracopy/data D-DB1 I-1730530050_TS-SOE_FNO-
27 7pltkqr - 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 181bhwjg .tm -
dbid changed, wrote new name
    Datafile
/nfsfsxn/oracopy/DB1/FB864A929AEB79B9E053630F1EAC7046/datafile/o1 mf
_temp_181bhz6g_.tm - dbid changed, wrote new name
    Datafile
/nfsfsxn/oracopy/DB1/FB867DA8C68C816EE053630F1EAC2BCF/datafile/o1 mf
temp l8lbjl6t .tm - dbid changed, wrote new name
    Datafile
/nfsfsxn/oracopy/DB1/FB867EA89ECF81C0E053630F1EACB901/datafile/o1 mf
_temp_181bj135_.tm - dbid changed, wrote new name
    Datafile
/nfsfsxn/oracopy/DB1/FB867F8A4D4F821CE053630F1EAC69CC/datafile/o1 mf
```

```
_temp_181bj13g_.tm - dbid changed, wrote new name
	Control File /nfsfsxn/oracopy/dbl.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$database;
NAME OPEN_MODE LOG_MODE
_____ ____
DB1TST READ WRITE NOARCHIVELOG
SQL> show pdbs
                           OPEN MODE RESTRICTED
   CON ID CON NAME
              ----- ------
       2 PDB$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
                           OPEN MODE RESTRICTED
   CON ID CON NAME
_____ ____
      2 PDB$SEED
                                READ ONLY NO
       3 DB1 PDB1
                                READ WRITE NO
       4 DB1 PDB2
                                READ WRITE NO
                             READ WRITE NO
       5 DB1 PDB3
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=AL!4422!3!467723097970!e!!g!!aws%20ec2&ef_id=Cj0KCQiA54 KfBhCKARIsAJzSrdqwQrghn6I71jiWzSeaT9Uh1-vY-VfhJixFxnv5rWwn2S7RqZOTQ0aAh7eEALw wcB:G:s&s kwcid=AL!4422!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

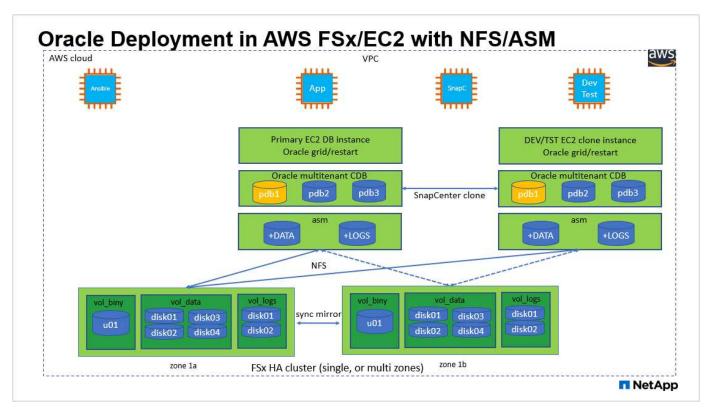
This solution is intended for the following people:

- 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	
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 alway 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 actually 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 stated in /tmp/archive on the EC2 instance.

```
[ec2-user@ip-172-30-15-58 ~]$ ls -1 /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:

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
```

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,wsize=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,wsize=65536
```

7. Change /oradata mount point ownership to the Oracle user and it's associated primary group.

sudo chown oracle:oinstall /oradata

8. Create /oralogs directory to mount Oracle logs file system

sudo mkdir /oralogs

9. Mount the log volume to /oralogs, changed to your FSx NFS lif IP address

```
sudo mount -t nfs 172.30.15.19:/ora_01_logs /oralogs -o
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsize=65536
```

10. Change /oralogs mount point ownership to the Oracle user and it's associated primary group.

sudo chown oracle:oinstall /oralogs

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,rsize=65536,wsize=65536
                                                              0
0
172.30.15.19:/ora 01 data /oradata
                                              nfs
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsize=65536
                                                              0
0
172.30.15.19:/ora 01 logs /oralogs
                                              nfs
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsize=65536
                                                              0
0
```

12. sudo to oracle user, create asm folders to store asm disk files

```
sudo su
su - oracle
mkdir /oradata/asm
mkdir /oralogs/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 infomation:

```
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/*,/oralogs/asm/*
oracle.install.asm.diskGroup.diskDiscoveryString=/oradata/asm/nfs da
ta 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.

+ Name	Dowcot	Ctata	Server
Name State details	Target	State	Server
Local Resources			
	ONLINE	ONLINE	ip-172-30-15-58
ora.DATA.dg STABLE	ONLINE	ONLINE	TD-1/2-30-13-30
ora.LISTENER.lsnr	ONLINE	ONLINE	ip-172-30-15-58
STABLE		ONLINE	10 172 30 13 30
ora.LOGS.dg	ONLINE	ONLINE	ip-172-30-15-58
STABLE	ONDINE	ONDINE	TD 112 00 10 00
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 <code>\$ORACLE HOME</code> and <code>\$ORACLE SID</code> 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.InstallEdition=EE oracle.install.db.OSDBA_GROUP=dba oracle.install.db.OSDER_GROUP=oper oracle.install.db.OSDGDBA_GROUP=dba oracle.install.db.OSKMDBA_GROUP=dba oracle.install.db.OSRACDBA_GROUP=dba oracle.install.db.OSRACDBA_GROUP=dba

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 sofware-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, lauch 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 _____ _____ Target State Server Name 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.db1.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 ONLINE ONLINE ip-172-30-15-58 1 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: /oralogs
```

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> set lin 200 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 /oralogs/asm/nfs logs disk01 MEMBER CACHED NORMAL LOGS 0001 /oralogs/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 MOUNTED 4194304 LOGS 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=AL!4422!3!467723097970!e!!g!!aws%20ec2&ef_id=Cj0KCQiA54 KfBhCKARIsAJzSrdqwQrghn6I71jiWzSeaT9Uh1-vY-VfhJixFxnv5rWwn2S7RqZOTQ0aAh7eEALw_wcB:G:s&s_kwcid=AL!4422!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

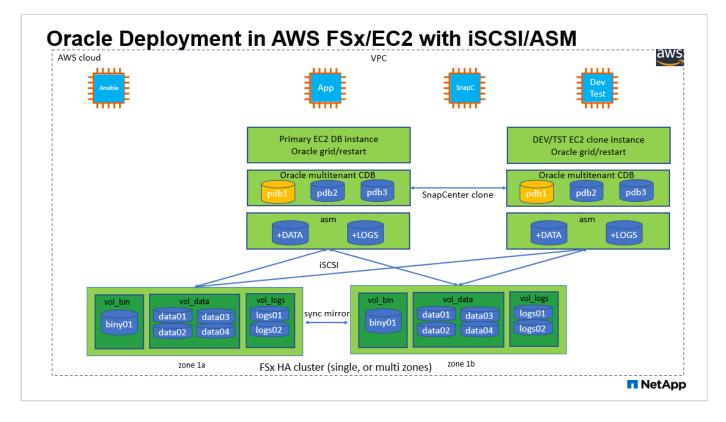
This solution is intended for the following people:

- 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

Current version offered by AWS

FSx ONTAP storage

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	napCenter Service Version			

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 alway 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 actually 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 stated in /tmp/archive on the EC2 instance.

```
[ec2-user@ip-172-30-15-58 ~]$ ls -1 /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:

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:

Vserver	Path	Igroup	LUN ID
Protocol			
svm_ora	/vol/ora_01_biny/ora_01_biny_01	ora_01	0
iscsi svm_ora	/vol/ora_01_data/ora_01_data_01	ora_01	1
iscsi svm_ora	/vol/ora_01_data/ora_01_data_02	ora_01	2
iscsi svm_ora	/vol/ora_01_data/ora_01_data_03	ora_01	3
iscsi svm_ora	/vol/ora_01_data/ora_01_data_04	ora_01	4
iscsi svm_ora	/vol/ora_01_logs/ora_01_logs_01	ora_01	5
iscsi 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 -1 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:
172.30.15.51,3260] 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 vserver(cDOT/FlashRay) lun-pathname filename adapter protocol size product /vol/ora 01 logs/ora 01 logs 02 svm ora iscsi 40g CDOT /dev/sdn host3 /vol/ora_01_logs/ora_01_logs_01 svm ora /dev/sdm host3 iscsi 40q CDOT /vol/ora 01 data/ora 01 data 03 svm ora /dev/sdk host3 iscsi 20q CDOT svm ora /vol/ora 01 data/ora 01 data 04 /dev/sdl iscsi 20q CDOT host3 svm ora /vol/ora_01_data/ora_01_data_01 iscsi 20q CDOT /dev/sdi host3 /vol/ora 01 data/ora 01 data 02 svm ora /dev/sdj host3 iscsi 20g CDOT /vol/ora 01 biny/ora 01 biny 01 svm ora 40a CDOT /dev/sdh host3 iscsi /vol/ora 01 logs/ora 01 logs 02 svm ora /dev/sdg host2 iscsi 40q CDOT /vol/ora_01_logs/ora_01_logs_01 svm ora 40g /dev/sdf host2 iscsi CDOT /vol/ora 01 data/ora 01 data 04 svm ora /dev/sde host2 iscsi 20g CDOT /vol/ora 01 data/ora 01 data 02 svm ora /dev/sdc host2 iscsi 20q 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 iscsi 20g CDOT host2 svm ora /vol/ora 01 biny/ora 01 biny 01 /dev/sda iscsi 40g CDOT host2

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 -1 /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
                           7 Mar 21 20:13
lrwxrwxrwx 1 root root
3600a09806c574235472455534e685145 -> ../dm-5
                           7 Mar 21 20:13
lrwxrwxrwx 1 root root
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:

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 {
                                 3600a09806c574235472455534e68512d
                wwid
                alias
                                 ora 01 biny 01
        }
        multipath {
                                 3600a09806c574235472455534e685141
                wwid
                alias
                                 ora 01 data 01
        }
        multipath {
                wwid
                                 3600a09806c574235472455534e685142
                alias
                                 ora 01 data 02
        }
        multipath {
                wwid
                                 3600a09806c574235472455534e685143
                alias
                                 ora 01 data 03
        }
        multipath {
                                 3600a09806c574235472455534e685144
                wwid
                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
                             7 Mar 21 20:41 ora 01 biny 01 -> ../dm-
lrwxrwxrwx 1 root root
0
                             7 Mar 21 20:41 ora 01 data 01 -> ../dm-
lrwxrwxrwx 1 root root
1
                             7 Mar 21 20:41 ora 01 data 02 -> ../dm-
lrwxrwxrwx 1 root root
2
lrwxrwxrwx 1 root root
                             7 Mar 21 20:41 ora 01 data 03 -> ../dm-
3
                             7 Mar 21 20:41 ora 01 data 04 -> ../dm-
lrwxrwxrwx 1 root root
4
lrwxrwxrwx 1 root root
                             7 Mar 21 20:41 ora 01 logs 01 -> ../dm-
5
                             7 Mar 21 20:41 ora 01 logs 02 -> ../dm-
lrwxrwxrwx 1 root root
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 it's associated primary group.

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

15. Find the UUI 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 infomation:

```
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/m
apper/ora 01 data 02,/dev/mapper/ora 01 data 03,/dev/mapper/ora 01 d
ata 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	ONT THE	ONT THE	
ora.driver.afd STABLE	ONLINE	ONLINE	ip-172-30-15-58
	ONLINE	ONLINE	ip-172-30-15-58
ora.evmd STABLE	ONTINE		TD 1/2-20-12-20
ЧЦЦЦЦЦСТО			

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 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 1048576 0 81847 81920 81847 0 N DATA/ 512 4096 1048576 81853 0 MOUNTED EXTERN N 512 0 81920 81853 N LOGS/ ASMCMD> afd state ASMCMD-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 <code>\$ORACLE_HOME</code> and <code>\$ORACLE_SID</code> 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.InstallEdition=EE oracle.install.db.OSDBA_GROUP=dba oracle.install.db.OSDER_GROUP=oper oracle.install.db.OSDGDBA_GROUP=dba oracle.install.db.OSKMDBA_GROUP=dba oracle.install.db.OSRACDBA_GROUP=dba oracle.install.db.OSRACDBA_GROUP=dba

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 sofware-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, lauch 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
             Target State Server
Name
                                                     State
details
Local Resources
                                                     STABLE
ora.DATA.dg ONLINE ONLINE
                               ip-172-30-15-58
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.db1.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> selec	ct name, open_mode from v\$database;
NAME	OPEN_MODE
DB1	READ WRITE
SQL> selec	ct name from v\$datafile;
NAME	
+DATA/DB1/ +DATA/DB1/ +DATA/DB1/ 77009 +DATA/DB1/ 77009 +DATA/DB1/ 2177009 +DATA/DB1/ 77853 +DATA/DB1/ 77853 +DATA/DB1/ 2177853	<pre>/DATAFILE/system.256.1132176177 /DATAFILE/sysaux.257.1132176221 /DATAFILE/undotbs1.258.1132176247 /86B637B62FE07A65E053F706E80A27CA/DATAFILE/system.265.11321 /86B637B62FE07A65E053F706E80A27CA/DATAFILE/sysaux.266.11321 /DATAFILE/users.259.1132176247 /86B637B62FE07A65E053F706E80A27CA/DATAFILE/undotbs1.267.113 /F7852758DCD6B800E0533A0F1EAC1DC6/DATAFILE/system.271.11321 /F7852758DCD6B800E0533A0F1EAC1DC6/DATAFILE/sysaux.272.11321 /F7852758DCD6B800E0533A0F1EAC1DC6/DATAFILE/undotbs1.270.113</pre>
7871 NAME	
77871 +DATA/DB1/ 77871 +DATA/DB1/ 2177871 +DATA/DB1/ 7889 +DATA/DB1/ 77889 +DATA/DB1/ 77889	<pre>/F785288BBCD1BA78E0533A0F1EACCD6F/DATAFILE/system.276.11321 /F785288BBCD1BA78E0533A0F1EACCD6F/DATAFILE/sysaux.277.11321 /F785288BBCD1BA78E0533A0F1EACCD6F/DATAFILE/undotbs1.275.113 /F785288BBCD1BA78E0533A0F1EACCD6F/DATAFILE/users.279.113217 /F78529A14DD8BB18E0533A0F1EACB8ED/DATAFILE/system.281.11321 /F78529A14DD8BB18E0533A0F1EACB8ED/DATAFILE/sysaux.282.11321</pre>

+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=Cj0KCQiA54 KfBhCKARIsAJzSrdqwQrghn6I71jiWzSeaT9Uh1-vY-VfhJixFxnv5rWwn2S7RqZOTQ0aAh7eEALw_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:

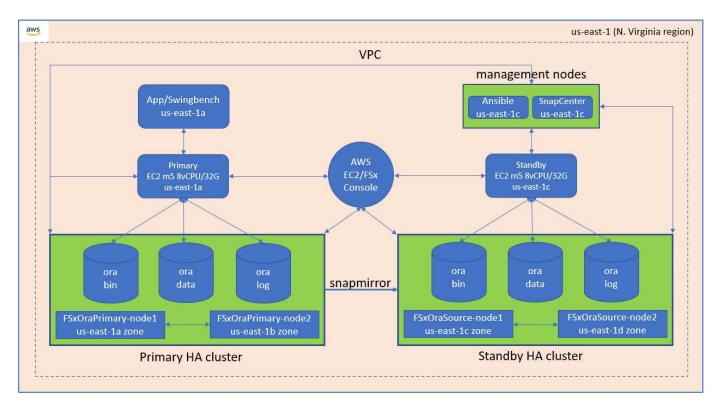
Modernize your Oracle database with hybrid cloud in AWS and FSx ONTAP, Part1 - Use case and solution architecture

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 10Gps 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

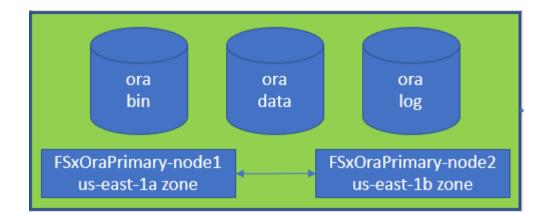
Selecting the right VM size is important for optimal performance of a relational database in a public cloud. For better performance, NetApp recommends using an EC2 M5 Series instance for Oracle deployment, which is optimized for database workloads. The same instance type is also used to power a RDS instance 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.

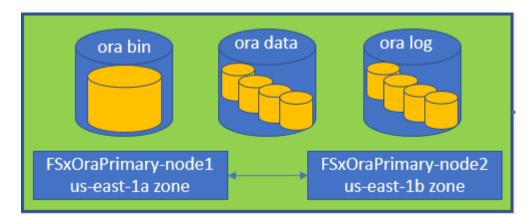
Storage layout and settings

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
Control filesData filesRedo logs	<pre>rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsize=6 5536</pre>
ORACLE_HOME ORACLE_BASE	<pre>rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsize=6 5536</pre>

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:

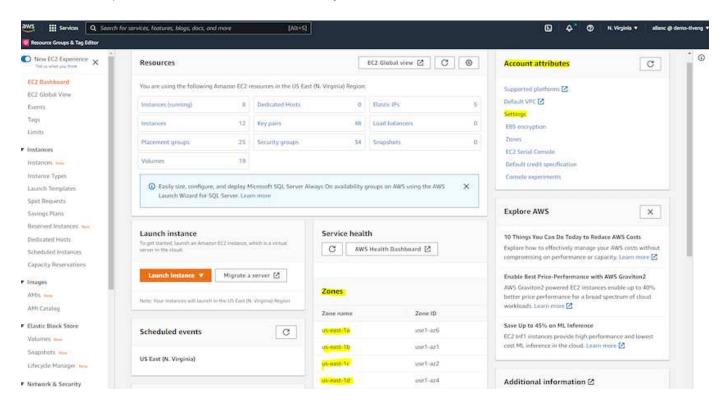
• Set up to use Amazon EC2

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.



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.

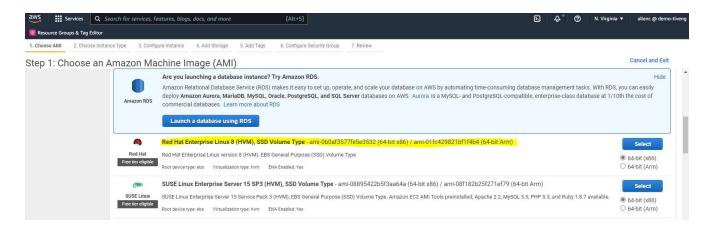
- Step 1: Launch an instance.
- Step 2: Connect to your instance.
- Step 3: Clean up your instance.

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.

Resource Groups & Tag Editor					
New EC2 Experience X	Resources		EC2 Global view 🖄 📿 💿	Account attributes	C
EC2 Dashboard	Resources			Account attributes	
EC2 Global View	You are using the following Amazon EC2 r	esources in the US East (N. Virginia) Region	:	Supported platforms	
Tags	Instances (running) 6	Dedicated Hosts 0	Elastic IPs 5	VPC Default VPC	
imits	Instances 12	Key pairs 48	Load balancers 0	none	
istances	Placement groups 25	Security groups 33	Snapshots 0	Settings EBS encryption	
nstances New	Volumes 19			Zones	
aunch Templates pot Requests avings Plans	Easily size, configure, and deploy M Launch Wizard for SQL Server. Lease	icrosoft SQL Server Always On availability n more	groups on AWS using the AWS $$ X	EC2 Serial Console Default credit specification Console experiments	
eserved Instances New edicated Hosts cheduled Instances	Launch instance	Service heal	th	Explore AWS	×
pacity Reservations	server in the cloud.		S Health Dashboard 🖸	Save up to 90% on EC2 with Spot Ins	stances
nages Mis New	Launch instance 🔻 Migrate a	server 🖸 Region US East (N. Virgi		Optimize price-performance by comb options in a single EC2 ASG. Learn mo	ining EC2 purchase

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." Search for services, features, blogs, docs, and more
 [Alt
Resource Groups & Tag Editor

. 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

m4	m4.16xlarge	64	256	EBS only	Yes	25 Gigabit	Yes
m5	m5.large	2	8	EBS only	Yes	Up to 10 Gigabit	Yes
m5	m5.xlarge	4	16	EBS only	Yes	Up to 10 Gigabit	Yes
m5	m5.2xlarge	8	32	EBS only	Yes	Up to 10 Gigabit	Yes
m5	m5.4xlarge	16	64	EBS only	Yes	Up to 10 Gigabit	Yes
m5	m5.8xlarge	32	128	EBS only	Yes	10 Gigabit	Yes
m5	m5.12xlarge	48	192	EBS only	Yes	10 Gigabit	Yes
m5	m5.16xlarge	64	256	EBS only	Yes	20 Gigabit	Yes
m5	m5.24xlarge	96	384	EBS only	Yes	25 Gigabit	Yes
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 Q Search for	r servic	es, features, blogs,	, docs, and more		[Alt+S							Ð	\$°	0	N. Virginia 🔻	allenc @	demo-tlveng
Resource Groups & Tag Editor																	
1. Choose AMI 2. Choose Instance Type	3. C	onfigure Instance	4. Add Storage	5. Add Tags	6. Configu	re Security Group	7. Review										
Step 3: Configure Instan	ce D	etails															1
No default VPC found. Select anothe	r VPC, o	or create a new de	efault VPC .														×
Configure the instance to suit your require	ements	. You can launch n	nultiple instances	from the same /	AMI, reques	t Spot instance	s to take advant	tage of the low	er pricing, ass	gn an access r	nanagement	role to the	e instance	e, and mo	ore.		
Number of instances	(\mathbf{i})	1		Launch into A	uto Scaling	Group 🕕											
Purchasing option	(j)	🗆 Request Sp	ot instances														
Network	()	vpc-04740641 No default VPC	<mark>fc537e5182</mark> C found. Create a	new default VP0		Create new	VPC										
Subnet	(1)	subnet-08c95 250 IP Address	2541f4ab282d u ses available	us-east-1a	4	Create new	subnet										
Auto-assign Public IP	(j)	Enable			4												
Hostname type	(j)	Use subnet se	etting (IP name)		4												
DNS Hostname	(j)	Enable IP n	ame IPv4 (A reco	rd) DNS requests													
		🗹 Enable reso	ource-based IPv4	(A record) DNS r	equests												
		🗌 Enable reso	ource-based IPv6	(AAAA record) D	NS request	5											
Placement group	(j)	Add instant	ce to placement g	Iroup													
Capacity Reservation	(1)	Open			4												
Domain join directory	(j)	No directory			\$	Create new	directory										
IAM role	(j)	None			4	Create new	IAM role										
	~	-												_			•
											Can	icel Pr	revious	Review	w and Launch	Next: Ad	id 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.

	ervices Q Search fo	r services, features, blog	gs, docs, and more	2	[Alt	+5]				D	\$ *	0	N. Virginia 🔻	allenc @ demo-tl
Resource Group:	ps & Tag Editor													
Choose AMI	2. Choose Instance Type	3. Configure Instance	4. Add Storage	5. Add Tags	6. Cont	igure Security Group	7. Review							
ur instance will b	d Storage I be launched with the fo of the root volume. You o n Amazon EC2.													
Volume Type (j	Device (j) Snapshot (į		Size (GiB)	i) vo	olume Type (j)		IOPS ()	Throughput (MB/s) (j)	Delete on Termination	D I	Encryptior	1	
Root	/dev/sda1	snap-03a3ad	00558b4d17c	50	G	eneral Purpose SSD (g	1p2) 🗸	150 / 3000	N/A		1	Not Encryp	oted	•
Add New Volume	ne													
Shared file	e systems 🛈								_					
ou currently don't	n't have any file systems o	on this instance. Select "A	dd file system" but	ton below to add	d a file syster	na								
dd file system	E)													

6. In Step 5, add a tag for instance identification if needed.

aws :	Services Q Search for	services, features, blog	s, docs, and more	2	[Alt+S]					D	\$°	0	N. Virginia 🔻	allenc @ demo-tlveng
🧑 Resource	Groups & Tag Editor													
1. Choose AM	Al 2. Choose Instance Type	3. Configure Instance	4. Add Storage	5. Add Tags	6. Configure Security Group	7. Review								
A tag consist A copy of a t	Add Tags ts of a case-sensitive key-valu ag can be applied to volumes, applied to all instances and vo	instances or both.												
Key (12)	8 characters maximum)		Value	(256 characters	maximum)		Instances (j)	Volumes (j)	Network Interfaces (
				This resource	currently has no tags									
					ton or click to add a Name ta ncludes permissions to creat									
Add Tag	(Up to 50 tags maximum)													
									Cancel Prev	ious	Review ar	id Launch	Next: Con	figure 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.

Resource Groups & Tag Editor Choose AMI 2. Choose Instance	Type 3, Configure Instance 4, Add Storage	5. Add Tags 6. Configure Security Group	7. Review		
2. Choose Instance	a type 3. Configure instance 4. Add Storage	5. Add rags 6. Conligure Security Group	7. Review		
ep 6: Configure Se					
			ic traffic to reach your instance. For example, if you want to set u e below. Learn more about Amazon EC2 security groups.	up a web server and allow internet traffic to rea	ch your instance, add r
	urity group: O Create a new security group	ecunty group or select norm an existing o	e below. Learning about Amazon Ecz security groups.		
Assign a sec	Select an existing security group				
	Select an existing security group				
Security Group ID	Name		Description		Actions
sg-0d746a0908b897c48	AviOccm03112021OCCM1635951256631-O	CCMSecurityGroup-B3QFHUHJRUVW	NetApp OCCM Instance External S	Security Group	Copy to new
sg-07b0625cd544aee16	AviOCCM0311OCCM1635943382952-OCCM	//SecurityGroup-1L8D4QX2SC945	NetApp OCCM Instance External S	Security Group	Copy to new
sg-0618122caef6c50e9	AviOccm1103OCCM1635944222133-OCCM	SecurityGroup-DX5PHX6CKVKC	NetApp OCCM Instance External S	Security Group	Copy to new
sg-0d63ea8c78987e660	AviOccm1209OCCM1631452667252-OCCM	ISecurityGroup-T5KVZ1Q4SH48	NetApp OCCM Instance External S	Security Group	Copy to new
sg-0aed9f8836b48c52d	AviOccmFSxOCCM1638110371156-OCCMS	SecurityGroup-N0ENZJW3TVYB	NetApp OCCM Instance External S	Security Group	Copy to new
sg-083a6ea5cba912375	connector01OCCM1631455604110-OCCMS	ecurityGroup-1790QV45PH3ZW	NetApp OCCM Instance External S	Security Group	Copy to new
sg-08148ca915189ac87	default		default VPC security group		Copy to new
sg-07f6c527620e3bb22	fsx02OCCM1633339531669-OCCMSecurity	Group-1XZYC5WM15NP7	NetApp OCCM Instance External S	Security Group	Copy to new
sg-0f359d2ba38db749f	SG-Version10-0CEc6MEs-NetAppExternalSe	ecurityGroup-N8B50KGTK58U	ONTAP Cloud firewall rules for ma	inagement and data interface	Copy to new
ound rules for sg-08148ca915	189ac87 (Selected security groups: sg-08148ca9	15189ac87)			
pe (i)	Protocol (j)	Port Range (j)	Source (j)	Description ()	
traffic	All	All	192.168.1.0/24		
traffic	All	All	sg-08148ca915189ac87 (defa	ult)	

- 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	Q Search for	services, features,	blogs, docs, and more	[Alt+S]			Σ	\$°	0	N. Virginia 🔻	allenc @ demo-	tlveng
🔯 Resource Groups & Tag	Editor											
1. Choose AMI 2. Choo	se Instance Type	3. Configure Instar	nce 4. Add Storage	5. Add Tags 6. Configure Security G	roup 7. Review							
Step 7: Review Please review your instan			to edit changes for eac	h section. Click Launch to assign a ke	y pair to your instance and complete 1	the launch process.						
 AMI Details 											Edit Al	iN
Free tier Red Hat E		sion 8 (HVM), EBS 0	olume Type - ami-0b Seneral Purpose (SSD) Ve								Edit instance typ	De
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 										E	dit security group	os
Security Group ID			Name		Description							
sg-08148ca915189a	c87		default		default VPC security	group						
All selected security	groups inbound	rules										
Туре (į)		Protocol (j)		Port Range (j)	Source (i)	Description (j)						
All traffic		All		All	192.168.1.0/24							
All traffic		All		All	sg-08148ca915189ac87 (de	fault)						
Instance Details										E	dit instance detai	Is
Storage											Edit storag	je _
										Cancel	Previous Lau	inch

d to iu to
202
101
1000
ore
~
~
his

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-db1v2.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.

aws 🔛 Services Q Search for	services, features, blogs, docs, and more [Alt+S]	▷ 🗘 ⑦ N. Virginia ▼ allenc @ demo-tlv
🧕 Resource Groups & Tag Editor		
Amazon FSx X	fsx (svm-005c6edf027866ca4)	Delete Update
File systems	Summary	
Volumes Backups • ONTAP Storage virtual machines • OpenZFS Snapshots • Windows File Server • Lustre Data repository tasks FSx on Service Quotas	SVM ID Creation time Active Directory svm-00566edf027866ca4 1 2022-01-24T18.02:24-05:00 - SVM name Ifecycle state 2 - fx 1 2 Created - UUID Subtype DEFAULT - 1a07ea1f-7d6e-11ec-97a9-7df96e22a64a EFAULT - File system 10 - - fs-0a51a3f08922224d5 1 - Resource ARN amawsr5xcus-east-1759995470648:storage-virtual-machine/fs- 0a51a3f08922224d5/xm-0056edf027866ca4 1	
	Endpoints Management DNS name Management IP address svm-00566edf027866ca4.fs-0a51a3f08922224d5.fsx.us-east-1.amazonaws.com Management IP address NFS DNS name Igs119.255.68 svm-00566edf027866ca4.fs-0a51a3f08922224d5.fsx.us-east-1.amazonaws.com Igs119.255.68 ISCSI DNS name IsCSI IP addresse iscsi svm-00566edf027866ca4.fs-0a51a3f08922224d5.fsx.us-east-1.amazonaws.com Igs119.255.68	

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.

WS Image: Services Q Search for se Resource Groups & Tag Editor	vices, features, blogs, docs, and more [Alt+5]				⊾ \$° Ø №	
Amazon FSx X	FSx > File systems					
File systems Volumes Backups	Did you know? With Amazon FSx for Windows File Server, you can reduce storage costs by 50-60% using Data De Learn how to easily enable this capability and others. Z	uplication.				×
ONTAP						-
Storage virtual machines	File systems (1)			C	Attach Actions V	Create file system
OpenZFS	Q Filter file systems					< 1 > @
Snapshots						
Windows File Server	File system File system ID ▼ File system Status name ▼ Status		Storage ⊽ type ⊽	Storage capacity ♥	Throughput capacity v	Creation time
Lustre	O rdscustomfs007 ☐ ONTAP ⊘Available	Multi-AZ	SSD	1,500 GiB	128 MB/s	2022-01-24T18:31:55- 05:00

2. Select Amazon FSx for NetApp ONTAP. Then click Next.

WS III Services Q. 3	Search for services, features, blogs, docs, and more	[Alt+5]		🛛 🔶 C
Resource Groups & Tag Editor				
F5x 3 File systems 3	Create file system			
Step 1 Select file system type	Select file system type			
Smp.2 Specify file system details	File system options			
Sup 3 Review and create	Amazon F5x for NetApp ONTAP FSX Amazon F5x for NetApp ONTAP	• Amazon FSx for OpenZFS FSXz= Amazon FSx for OpenZFS	• Aniazon PSx for Windows File Server FSX Amazon FSx for Windows File Server	Amazon FSx for Lustre
	Broadly accessible from Linux, Window Provides ONTAP's popular data manag Delivers hundreds of thousands of IDF Offers highly-available and highly-due Automatically tiers infrequently-acces	feature-rich, high-performance, and highly-reliable es, and macOS compute instances and containers (n ement capabilities like Snapshots, SnapMicro (for 5 with consistent uch-milliscend latencies, and up able multi-AZ SSD storage with support for cross-re	egion replication and built-in, fully managed backups. rage tier that can scale to petabytes in size and is cost-o	FS, SMB, and ISCSI protocols. compression / deduplication.
				Cancel

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 name - optional Info	
aws_ora_prod	
Maximum of 256 Unicode letters, whitespace, and numbers, plus + - = : /	
Deployment type Info	
O Multi-AZ	
○ Single-AZ	
SSD storage capacity Info	
1024	
Minimum 1024 GiB; Maximum 192 TiB.	
Provisioned SSD IOPS Amazon FSx provides 3 IOPS per GiB of storage capacity. You can also provision additional needed.	SSD IOPS as
 Automatic (3 IOPS per GiB of SSD storage) 	
O User-provisioned	
40000	
Maximum 80,000 IOPS	
Throughput capacity Info The sustained speed at which the file server hosting your file system can serve data. The file burst to higher speeds for periods of time.	e server can also
 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	
O No preference	
Select an IP address range	

5. In the Security & Encryption section, accept the default, and enter the fsxadmin password.

ncryption key Info WS Key Management Service (KMS) encryption key that protects y	our 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
ile system administrative password assword for this file system's "fsxadmin" user, which you can use to Don't specify a password	access the ONTAP CLI or R	EST API.
assword for this file system's "fsxadmin" user, which you can use to	o access the ONTAP CLI or R	EST API.
assword for this file system's "fsxadmin" user, which you can use to Don't specify a password	access the ONTAP CLI or R	EST API.
assword for this file system's "fsxadmin" user, which you can use to Don't specify a password Specify a password	access the ONTAP CLI or R	EST API.

6. Enter the SVM name and the vsadmin password.

Storage virtual machine name
fsxora_prod
SVM administrative password Password for this SVM's "vsadmin" user, which you can use to access the ONTAP CLI or REST API.
O Don't specify a password
Specify a password
Password
Confirm password
Active Directory
Active Directory Joining an Active Directory enables access from Windows and MacOS clients over the SMB protoco
O 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.

/olume name			
vol1			
Aaximum of 203 alphanumeric characters, plus			
unction path			
/vol1			
he location within your file system where your volume will be mounted.			
/olume size			
1024	Ĩ		
linimum 20 MiB; Maximum 104857600 MiB			
storage efficiency select whether you would like to enable ONTAP storage efficiencies on your volume: deduplication, ompression, and compaction.			
Enabled (recommended)			
Disabled			
Capacity pool tiering policy You can optionally enable automatic tiering of your data to lower-cost capacity pool storage.			
Auto	•		
Backup and maintenance - optional			
Tags - optional			

8. Review the Summary page, and click Create File System to complete FSx file system provision.

Step 1 Select file system type	Create file syste	em	
Step 2 Specify file system details	Summary Verify the following attributes before proceeding		
Step 3 Review and create	Attribute	Value	Editable after creation
	File system type	Amazon FSx for NetApp ONTAP	
	File system name	aws_ora_prod	\odot
	Deployment type	Multi-AZ	
	Storage type	SSD	
	SSD storage capacity	1,024 GiB	\odot
	Minimum SSD IOPS	40000 IOPS	\odot
	Throughput capacity	512 MB/s	\odot
	Virtual Private Cloud (VPC)	vpc-0474064fc537e5182	
	VPC Security Groups	sg-08148ca915189ac87	\odot
	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	\odot
	Automatic backup	7 dav(s)	Ø

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.

×

w

W

Enabled (recommended) Disabled		
Capacity pool tiering policy You can optionally enable automatic tiering of your dat	ta to lower-cost capacity pool	storage.
None		•

Create	vol	lume
cicate		unit

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) 		
O Disabled		
Capacity pool tiering policy You can optionally enable automatic tiering of your o	data to lower-cost capacity pool	storage.
None		
	Cancel	Confirm

X

Create	vo	lume
CICULC		curric.

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.

S				
Vo	11100	e size		
	curri	P 2150		

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) 	
O Disabled	
Capacity pool tiering policy You can optionally enable automatic ti	ering 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.

х

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:

- 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/ec2user/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-addrs
standby_cluster_name,inter_cluster_ip_address -username fsxadmin
-initial-allowed-vserver-peers *
```

 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.

FsxId00164 Peer Clust		6::> cluster Cluster	peer show Serial Number	Availabil	ity Auther	ntication
FsxId0b6a9	5149d07aa82	e 1-80-000	011	Available	ok	
FsxId00164	454fac5591e	6::> vserver	peer show			
	Peer	Peer		Pe	ering	Remote
Vserver	Vserver	State	Peer Clust	er Ap	plications	Vserver
svm_FSxOra		22. S. 197				
	svm_FSxOra	aTarget				
		peered	FsxId0b6a9	5149d07aa8	2e	
				sn	apmirror	svm_FSxOraTarget
FsxId00164	454fac5591e	6::>				

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

 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

 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.

III N	etApp		A
SnapCen	ter [®]		
Username	rdscustomval\administrator	0	
Password			
	Sign In	3 .5	

2. From the left-side menu, click Settings, and then Credential and New to add ec2-user credentials for SnapCenter plugin installation.

to his	etApp Snap(Providence State Stat				
<		Global Settings Policies I	Jsers and Access Roles Credential Software	140 T T T T T T T T T T T T T T T T T T T				
	Dashboard	Search by Credential Name		🚊 🖾 🏛				
0	Resources	Credential Name	Authentication Mode	Details				
8	Monitor	244rdscustomdb	SQL	Userid:admin				
*1	Reports	42rdscustomdb	SQL	Useridzadmin				
	(Netwice)	admin	SQL	Userid:admin				
.	Hosts	administrator	Windows	Userld:administrator				
	Storage Systems	ec2-user	Linux	UserId:ec2-user				
1.00		onpremSQL	Windows	Userid:rdscustomval\administrator				
# }	Settings	rdsdb2	Windows	Userid:administrator				
	Alerts	rdsdb244	Windows	Userld:administrator				
		rdssql	Windows	Userid:administrator				
		tst244	SQL	Userld:admin				
		tstcredfordemo	Windows	UserId:administrator				

- 3. Reset the ec2-user password and enable password SSH authentication by editing the /etc/ssh/sshd config file on the EC2 instance host.
- 4. Verify that the "Use sudo privileges" checkbox is selected. You just reset the ec2-user password in the previous step.

Credential		×
Credential Name	ec2-user	
Authentication Mode	Linux]
Username	ec2-user	0
Password	•••••]
✓ Use sudo privileges	6	
	Cancel	ОК

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 ~]\$ si	udo vi /etc/hos	sts						
[ec2-user@ip-10-0-0-151 ~]\$ cat /etc/hosts									
127.0.0.1 loc	alhost localho	st.localdomain	localhost4						
localhost4.localdomain4									
::1 loc	alhost localho	st.localdomain	localhost6						
localhost6.localdomain6									
10.0.1.233 rds	customvalsc.rd	scustomval.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.

	letApp Snap(Cente	r®					? -	1 rdscus	:omval\adminis	trator Sna	oCenterAdmin	🖡 Sign Out
<			aged Hosts Disks Shares Initiator Gr	oups iSCSI :	Session					÷		ţ	
	Dashboard									Add	Remove	Refresh	More
V	Resources		Name 12	Туре	System		Plug-in				Version	Overall St	atus
•	Monitor		RDSAMAZ-VJODQKO	Windows	Stand- alone		Microso SQL Ser		ws Server, N	licrosoft	4.5	le Host	down
м́і	Reports		rdscustommssql1.rdscustomval.com	Windows	Stand- alone		Microso SQL Ser		ws Server, N	licrosoft	4.5	🔵 Runn	ng
Α.	Hosts												
50	Storage Systems												
	Settings												
▲	Alerts												

Check Oracle Database, and, before you submit, click More Options.

			9-	L rdscustomval\administrator	SnapCenterAdmin	🖡 Sign Out	
						×	
Add Host							
Host Type	Linux		•				
Host Name	10.0.0.151						
Credentials	ec2-user		•	+ 0			
	all SnapCenter Plug-ins Package Oracle Database SAP HANA rt, Install Path, Custom Plug-Ins	for Linu	IX				
Submit Cancel]						

Check Skip Preinstall Checks. Confirm Skipping Preinstall Checks, and then click Submit After Save.

More Options		×				
Port	8145	6				
Installation Path	/opt/NetApp/snapcenter					
	Skip preinstall checks					
Custom Plug-ins ——	Choose a File	<u> </u>				
	Browse Upload					
	No plug-ins found.	*				
	Save	ancel				

You are prompted with Confirm Fingerprint, and then click Confirm and Submit.

Confirm Fingerprint								
Authenticity of the	host can	not be determined 🜖						
Host name	1E	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						
		Confirm and Submit	Close					

After successful plugin configuration, the managed host's overall status show as Running.

Mana	aged Hosts Disks	Shares	Initiator Grou	ıps iSCSI Se	ession					
Sea	arch by Name	Y					Add		Refresh	More
	Name		48.	Туре	System	Plug-in		Version	Overal	Status
	ip-10-0-0-151.ec2.int	ernal		Linux	Stand- alone	UNIX, Oracle Database		4.5	🕚 Ru	nning

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.

R NetApp Sna	pCente	rðð				0	 Letteuchtenvel-administration ScopCo	nterAtinin 🖠 Sign Out
2	One	le Dalabier						
Dushboard	.vom	Ostabase	· Selectrike	sten 🔻			-	* *
Resources		Name	Oracle Database Type	HostiCluster	Resource Group	Policies	Last flackup	Overall Status
Montar		ORC1	Single tratates	ip-10-0-0-151.ec2.memul	ond full being	Criscle full backup. Oracle log backup	03/24/2022 8:40/08 PM	Backup succeeded
A Reports								
A 100								
1. Storage System								
EE Sellings								
A Nem								

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.

6	Oracle Database			
Dashboard	View Resource Group	Search resource	grouj V	
Resources	Name	Resources	Tags	Policies
Monitor	orcl_full_bkup	1	ora_fullbkup	Oracle full backup
Reports	orcl_log_bkup	1	ora_logbkup	Oracle log backup
Hosts				
Storage Syst	tems			
E Settings				

2. Click the full backup resource name, and then click the Backup Now icon to initiate an add-hoc backup.

II Ne	tApp SnapCenter®				٠	6.	L rdscustomval\administra	tor SnapCe	enterAdmin	🛿 Sign Out
>	Oracle Database									
	Search resource groups	search					Modify Resource Group	Back up Now	Maintenance	Delete
U	Name	Resource Name	Туре	Host						
	orcl_full_bkup	ORCL	Oracle Database	lp-10-0-0-151.ec2.internal						
a i	orcl_log_bkup									
A										
} •										
÷										
▲										

3. Click Backup and then confirm the backup to start a full database backup.

Backup			×								
Create a backup for the selected resource group											
Resource Group	orcl_full_bkup										
Policy	Oracle full backup	- 0									
🗌 Verify after backt	qu										
		Cancel Bac	kup								
			12								

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.

E N	etApp Snapi	Center®				٠	= 0- 1	L rdsoutherwalladministra	tor Snapile	nter Aan	in Wilson Out	
> Ⅲ	Oracle Database Chemistri Gatas		CARCL Topology								. =	and other
0 2 4 1 1	0 000. 26 36 36 36	Manage Copies					Summary Ca 20 Backup 2 Data Backu 18 Ling Backup 0 Comes	p1				
•			(mech V) Backip Name	Court	Туре	If But Date	verted	0 T Mounted R	E		A E	
			(p.19-00-131_09-25-2022_00.34 20.4541_0) (p.10-00-131_03-25-2022_00.34 30.4541_0)		Log Data	0305/002 12:54:37 AM 8	Not Applicable Unvertified	False False	Not Catalog		1733264	İ

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.

<		Oracle Database 🔽			
D	ashboard	View Resource Group	Search resource	groui V	
J Re	esources	Name	Resources	Tags	Policies
р м	Ionitor	orcl_full_bkup	1	ora_fullbkup	Oracle full backup
Re Re	eports	orcl_log_bkup	1	ora_logbkup	Oracle log backup
ь н	osts				
St	torage Systems				
	ettings				
E Se	erringo				

2. Click the log backup resource name, and then click the Backup Now icon to initiate an add-hoc backup for archive logs.

II Ne	etApp SnapCenter®				٠	•	8-	L rdscustomval\administra	tor Snap	CenterAdmin	🖡 Sign Out
> 	Oracle Database Search resource groups	orcl_log_bkup Details							Back up Now	*	Defer
U	Name	Resource Name	Туре	Host				Modify Resource Group	Back up Now	Maintenance	Delete
۲	orcl_full_bkup	ORCL	Oracle Database	ip-10-0-0-151.ec2.internal							
- 	orcl_log_bkup										
٨											
֥											
籬											
A											

3. Click Backup and then confirm the backup to start an archive log backup.

Backup			×
Create a backup	for the selected resource	e group	
Resource Group	orcl_log_bkup		
Policy	Oracle log backup	- 0	
		Cancel	Backup
		Te	

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.

III No	tApp SnapCenter®	80				101	= 0·	Leboutone/administrator ShapCom	erAdmin	#Sgs.Out
>	Oracle Database	ORCL Topology							Ξ.	
	Scarph databases								2	<u>_</u>
0	UF 2 Name	Manage Copies								
0	040.	27 Bockupe						Summary Card		
-		n Cloves						27 Backups		
Λ.		Local copies						2 Data Dackaps		
\$+								25 Ling Backupt II Climes		
-		Primary Backup(s)								
4		(main)						<u>e 1 5 5</u>		4 I
		Backup Name	Court	туре	.17	End Date	Vertfied	Mounted RMAN Cataloged	se	N.
		10-16-06-111 (01-24-2022) OF 59-38-0238 (1	1	Log		03/25/2022 1.59/46 AME 🗖	Not Applicable	False Not Cataloge	17	10201

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.

		Orade	enninee -	1					
	Ombhowd	Vere	Detabase	Scardt data	baies Y			-	* 1
9	Resources	2*	Name	Ciracle Database Type	HMMCluster	Resource Group	Policies	Last Backup	Overall Status
-	Monitor .		ORCL	Single instance	ip-10-0-0-151.ec2.mmmail	ort1, Rull, Bikupi ort1, Yog, Skupi	Oracle full backup Oracle log backup	03/35/2022 1:10:09 PM 🛱	Balkup sixtreeder
1	Reports								
	Hists								
	Dorage Systems								
	Settings								
÷.	Colores of Colores								

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.

n Ne	tApp SnapCenter®				•	= 0· 1/	that menalistics	drator SnapCenterAd	tmin #Sig
	Oracle Database	OliCL herdegy							
	Search Autobiors							-	2
2	17 🎮 Name	Manage Copies							
2	ORL.	70 fundages					Summary	Card	
		© Corres					78 Backops		
		Local supers					5-Data Re	chupt	
							Thing that	NGR.	
							9 Oores		
		Primary Backup(s)							
		(seed)					0 1	<u>n</u> <u>n</u> <u>n</u>	
		Backup Name	Courte	Type	UF End Date	Vertfelt	Mounted	IBMAN Cataloged	SCN
		ip-10-0-0-151_03-29-2022_12-46.01.1098_1		Lig	03/25/2022 12:40/05 PM	Net Applicable	False	Not Cataloged	175425
		ap-10-0-0-151,03-25-2022_12-25/01.0000_1	1	Log	03/25/2022 12:25:09 //4 🗖	Not Applicable	. False	Not Cataloged	(1933)
		sp-10-0-0-151,03-35-2022,12.10-01,1097_1		Lóg	03/25/2022 12:10:09 PM	Not Applicative	False	Not Cataloged	17824
		(p. 10-0-0-151, 03-25-3022, 11-35.01.0500, 1	Ĭ.	Lig.	092500011135.09AM	Not Applicatio	Falter	Note Cataloged	170711
		g- to o 6 tht (03-25-2022, 11-45.61-0023, 1	1	Lig	03/25/2022 11:40:09.8M	Not Applicable	Talte	Note Callabaged	120021
		10-10-00-151,03-25-2022_11-25-01-0450_1	1	Lig	03/05/2022 11:25:09 AM	Fect Applicable	Falter	Not Cataloged	17793
		30-10-0-6-111, 13-25-2022, 11,15-01,1508,1		Lúg -	03/25/2022 11:15:17 AM	Not Applicable	False	Not Cataloged	17785
		@1080191.03252002.111501.003.0	1	Data	03/25/2022 11:15:11 AM	Unverting	Tabe	Not Cataloget	177800
		e-10-0-151,03-25-2022,11-10-01-1834,1	÷	Ling	E3/25/2022 11/10/09 AM	Not Applicable	Faise	Not Cataloged	\$72812

3. Highlight the log volume snapshot and click the Mount button to mount the volume.

Manage Copies								
78 Backups D Clottes						Summar	y Card	
						78 Backups		
Local copies						5 Data B		
						73 Log Bi	ккарь	
						0 Clanes		
Primary Backup(s)								
search Y						e t		
Backup Name	Count	Type	17	End Date	Verified	Mounted	RMAN Cataloged	SCN
Ip-10-0-0-151_03-25-2022_12.40.01,1098_1	3	Log		03/25/2022 12:40:09 PM 🗖	Not Applicable	Fabe	Not Cataloged	1784293
Ip-10-0-0-151_03-25-2022_12-25-01.0080_1	1	Log		03/25/2022 12:25:09 PM 🗖	Not Applicable	False	Not Cataloged	1783383
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 10.0.0151_03-25-2022_11.55.01.0508_1	93	Log		03/25/2022 11:55:09 AM 🛱	Not Applicable	False	Not Cataloged	1781160
ip-10-0-0151_03-25-2022_11-40.01.0323_1	1	Log		03/25/2022 11:40:09 AM 🗖	Not Applicable	False	Not Cataloged	1780268
lp-10-0-0-151_03-25-2022_11.25.01.0430_1	9.8	Log		03/25/2022 11:25:09 AM 🗖	Not Applicable	False	Not Cataloged	1779368
ip-10.0.0.151_03-25-2022_11,15.01.1500_1	1	Log		03/25/2022 11:13:17 AM	Not Applicable	False	Not Cataloged	1778540
p-10-0-0-151_03-25-2022_11.15.01.1503_0	1	Data		03/25/2022 11:15:11 AM 🛱	Unvertified	False	Not Cataloged	1778504
ip-10-0-0-151_03-25-2022_11.10.01.1834_1		Log		03/25/2022 11:10:09 AM 🗖	Not Applicable	False	Not Cataloged	1778184

4. Choose the primary EC2 instance to mount the log volume.

Mount backup	5		×
Choose the host to mount the backup	ip-10-0-0-151.ec2.internal		
Mount path :	/var/opt/snapcenter/sco/backup_mount	/ip-10-0-0-151_03-25-2022_11.15.01.1503_1/ORCL	
			Mount Cancel

5. Verify that the mount job completes successfully. Also check on the EC2 instance host to see the that log volume mounted and also the mount point path.

) Maria	Schedules	tionts Ligs											
Dathboard	(search)	by northe	v										2	
tescorces	All jobs													
Mentor	10	Status	Name					8	Start date	End date	Owner			
a Reports	4350	-	Backup of Resource Group (ord, ing thing) a	the policy the	ede log beci	kup		3/25/2022 1 400	00,0% 🗖	3/25/2022 1:K013 PM	emeut	convel-activ	visition	
	4549		Mount Service in 18-06-05-05-25-0022-11					0905-0001+380	DIN C	08/25/2022 1:36:53 PM	10103	HTOURAL IN	tronger -	
mpfs mpfs dev/nvme0n1p 98.19.255.68 98.19.255.68 dev/mapper/d	/ora_n: /ora_n:	fs_data		16G 7.7G 7.7G 9.8G 48G 48G 40G	7.0G 604K 0 5.4G 95M 3.4G 471M	48G 45G 39G	18 08 568 18 88 28	/ora_nfs_log /ora_nfs_data /rdsdbdata						
dev/nvme5n1 mpfs mpfs				25G 1.6G 1.6G	12G 0	13G 1.6G 1.6G	0%	/rdsdbbin /run/user/61001 /run/user/61005						
			583-480d-9a34-6275dab17f5b r]#		91M	48G		/var/opt/snapcenter/sco/backup_	mount/ip-10-	0-0-151_03-25-202	2_11.15		3_1/ORCL	

6. 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/
```

7. 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 Bockups						Summar	y Card				
0 Clones						80 Backups					
Local copies						5 Data 8	Jackupt				
						75 Log Bi	<i>ickups</i>				
						0 Clones					
Primary Backup(s)											
(search 🛛						• I		2	a biturt		
Backup Name	Count	Туре	lī	End Date	Verified	Mounted	RMAN Cat	aloged		5CN	
(p-10-0-0-151_03-25-2022_12-10.01.1097_1	38	Log	03/25/2022 12:	10:09 PM 🛱	Not Applicable	Fabr	Not Cr	staloger	d.	178241	ų.
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 Ca	staloge	d	178116	ø
ap-10-0-0-151_03-25-2022_11.40.01.0323_1	1	Log	03/25/2022 11:	10:09 AM	Not Applicable	False	Not Cr	staloger	d	178026	8
(p-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 Cr	staloge	d	177936	8
ip-10-0-0-151_03-25-2022_11.15.01.1503_1	1	Log	03/25/2022 11:1	15:17 AM 🛱	Not Applicable	True	Not Ci	naloger	đ	177854	ő
ip-10-0-0-151_03-25-2022_11.15.01.1503_0	Э.	Data	03/25/2022 11:	15:11 AM 🛱	Onverfiled	False	Notica	stulogei	81) - E	177850	4
lp-10-0-0-151_03-25-2022_11.10.01.1834_1	1	Log	03/25/2022 11:	🗂 MA 60:01	Not Applicable	False	Not Cr	staloge	d	177818	4

8. Check "All Datafiles" and "Change database state if needed for restore and recovery", and click Next.

Restore ORCL		×
1 Restore Scope	Restore Scope 0	
2 Recovery Scope	All Datafiles	
3 PreOps	O Tablespaces	
4 PostOps	Control files	
5 Notification	Database State Image: Change database state if needed for restore and recovery	
6 Summary	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	at

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.

Restore ORCL			×
Restore Scope	Choose Recovery Scope		
2 Recovery Scope	O All Logs		
3 PreOps	Until SCN (System Change Number) SCN 1778546	0	
4 PostOps	 Date and Time No recovery 		
5 Notification			
6 Summary	Specify external archive log files locations		
			×
		Previous	Next

10. Specify an optional prescript to run if necessary.

Restore ORCL			×
1 Restore Scope	Specify optional so	cripts to run before performing a restore job 🚯	
2 Recovery Scope	Prescript full path	/var/opt/snapcenter/spl/scripts/ Enter Prescript path	
3 PreOps	Arguments		
4 PostOps	Script timeout	60 secs	
5 Notification			
6 Summary			
		Previous	Next

11. Specify an optional afterscript to run if necessary. Check the open database after recovery.

Restore ORCL					×
Restore Scope	Specify optional sc	ripts to run after performing a re	store job 🚯		
2 Recovery Scope	Postscript full path	/var/opt/snapcenter/spl/scripts/	nter Postscript path		
3 PreOps	Arguments				
4 PostOps	Open the databas	e or container database in READ-WRIT	E mode after recovery		
5 Notification					
6 Summary					
				Previous	xt

12. Provide an SMTP server and email address if a job notification is needed.

Restore ORCL			×
Restore Scope	Provide email set	tings ()	
2 Recovery Scope	Emall preference	Never -	
3 PreOps	From	From email	1
4 PostOps	То	Email to	1
	Subject	Notification	
5 Notification	Attach Job report		
6 Summary			
		Previous	Next

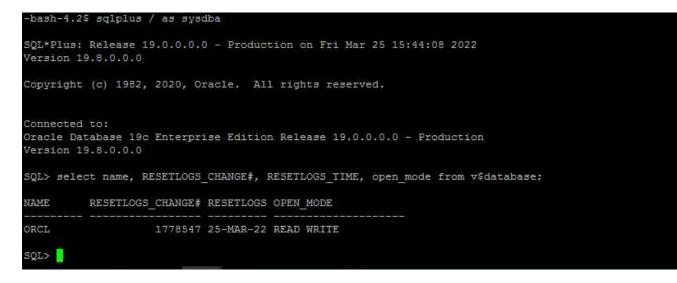
13. Restore the job summary. Click finish to launch the restore job.

Restore ORCL			×
1 Restore Scope	Summary		
2 Recovery Scope	Backup name	lp-10-0-0-151_03-25-2022_11.15.01.1503_0	
A	Backup date	03/25/2022 11:15:11 AM	
3 PreOps	Restore scope	All DataFiles	
O PostOps	Recovery scope	Until SCN 1778546	
5 Notification	Auxiliary destination		
	Options	Change database state if necessary , Open the database or container database in READ-WRITE mode after recovery	
6 Summary	Prescript full path	None	
	Prescript arguments		
	Postscript full path	None	
	Postscript arguments		
	Send email	No	
		Previous Finish	

14. Validate the restore from SnapCenter.

Job	Details	×
Rest	ore 'ip-10-0-0-151.ec2.internal\ORCL'	
~	Restore 'ip-10-0-0-151.ec2.internal\ORCL'	
1	p-10-0-0-151.ec2.internal	
~	► Prescripts	
~	► Pre Restore	
~	► Restore	
~	Post Restore	
~	Postscripts	
~	Post Restore Cleanup	
~	Data Collection	
~	Send EMS Messages	
0 Tas	sk Name: ip-10-0-0-151.ec2.internal Start Time: 03/25/2022 3:33:53 PM End Time: 03/25/2022 3:35:10 PM	•
	View Logs Cancel Job Close	
		- 13

15. Validate the restore from the EC2 instance host.



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.

🗖 Ne	tApp SnapCenter®				•	0-	1 rdscustomval\administra	tor SnapO	nterAdmin	🖡 Sign C
	Oracle Database 👻	orcl_full_bkup Details								
	Search resource groups	search					Modify Resource Group	Back up Now	Maintenance	Dele
U	Name	Resource Name	Туре	Host						
۲	orci_full_bkup	ORCL	Oracle Database	ip-10-0-0-151.ec2.internal						
2 21	orcl_log_bkup									
201 14 1-1										
₩ •										

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.

	A DATA CARDINA CARDON DATA								
	Oracle Database	CRCL Topology							
	Search Hability or Search Hability of Search Hability of Search S							E .	2
2	17 M Name	Manage Copies							
	ORC,	13 Bachan					Summar	o Caril	
1		0 Clones					T3 Backups	1 caro	
		Local opping					6 Data 6	lickups	
							B7 Log Bi	echoque.	
							0 Clones		
		Primary Backup(s)							
		(iouth V					0 1	9. 11. 1	
		(isent) (V) Rokup Name	Court	Typ#	Lif Find Date	verthed			
			Court 1	Dp#	LE Fred Dame DAVIS/2012 SISTON PM	Verhed Not Applicable			
		Backup Name	Cours 1			Not	Mounted	RMAN Cataloged	SCN 178909
2		Radiup Name @ 10.00-151_03-25.2022;17.50.01.0197_1	Court 1 1	Log	01/25/2022 5:35:09 PM 🗖	Not Applicable Not	Mounted	RMAN Cataloged Not Cataloged	SCN
2		Backup Name gr 10-0 - 151, 03-25 - 2022, 17 /5/0.01 0397, 1 gr 10-0 - 551, 03-05 - 2022, 17 /5/0.55 - 0853, 1	1	Log	03/25/2022 531:09 PM	Not Applicable Not Applicable	Mourted Fatte	RMAN Cataloged Not Cataloged Not Cataloged	SCN L78909 L78909

3. Mount the log volume snapshot taken in step 4 to the standby EC2 instance host.

Vext V	Manage Copies								
Load copes Count opes Count o	95 Backups					Summar	y Card		
Bit up Buckups Bit up Buckups Clines ines Clines Clines Clines Clines Clines Clines Clines Clines Clines Clines Clines Clines Clines Clines Clines Clines Clines Clines Clines Clines Clines Clines <th colsp<="" th=""><th></th><th></th><th></th><th></th><th></th><th></th><th></th><th></th></th>	<th></th> <th></th> <th></th> <th></th> <th></th> <th></th> <th></th> <th></th>								
Deckup Sarkup(s) Image: Count of the second of	Local copies								
verify verify									
Back up Name Court Type IF End Oate Verified Mounted RMAN Cataloged SCN ip 10 0 0 151 (03 25 2022 (15.55.01.0309.1) 1 Log 03/25/2022 6:50.03 PM Applicable Faise Not Cataloged 19/19/19/19/19/19/19/19/19/19/19/19/19/1	Primary Backup(s)								
ip-10-0-151_03-25-2022_18.55.01.0309_1 1 Log 03/25/2022_655:09 PM - Not Applicable False Not Cataloged 1892563 ip-10-0-151_03-25-2022_17.55.01.0197_1 1 Log 03/25/2022_655:09 PM - Not Applicable False Not Cataloged 1891375 ip-10-0-151_03-25-2022_17.55.01.0197_1 1 Log 03/25/2022_55:09 PM - Not Applicable False Not Cataloged 178909 ip-10-0-151_03-25-2022_17.55.01.0197_1 1 Log 03/25/2022_55:05 PM - Not Applicable False Not Cataloged 178809 ip-10-0-151_03-25-2022_17.50.50453.1 1 Log 03/25/2022_55:105 PM - Not False Not Cataloged 178809 ip-10-0-151_03-25-2022_17.50.50453.0 1 Data 03/25/2022_55:105 PM - Unvertified False Not Cataloged 178832 ip-10-0-151_03-25-2022_17.400.00738_1 1 Log 03/25/2022_55:105 PM - Not False Not Cataloged 178832 ip-10-0-0-47.ec2.internal 1 Log 03/25/2022_55:05 PM - Not Not False Not Cataloged 178832 Choose the host to mount the backup Ip-10-0	search 🛛						a 🗈 🙇		
Applicable Applicable Applicable Applicable Applicable Applicable Not Applicable Not Cataloged 1891375 ip 10 0 0 151,03 25 2022,17,55,01,0197,1 1 Log 03/25/2022 555:09 Mile Not False Not Cataloged 178909 ip 10 0 0 151,03 25 2022,17,50 55,0853,0 1 Log 03/25/2022 551:05 Mile Not False Not Cataloged 178909 ip 10 0 0 151,03 25 2022,17,50 55,0853,0 1 Log 03/25/2022 551:05 Mile Vice Applicable Not Cataloged 178882 ip 10 0 0 151,03 25 2022,17,50 55,0853,0 1 Log 03/25/2022 551:05 Mile Not False Not Cataloged 178882 ip 10 0 0 151,03 25 2022,17,60 0 9738,1 1 Log 03/25/2022 551:05 Mile Not False Not Cataloged 178882 ip 10 0 0 151,03 25 2022,17,40 0 9738,1 1 Log 03/25/2022 3-51:05 Mile Not False Not Cataloged 178882 Choose the host to mount the backup Ip-10-0-047.ec2.internal Image: state of the state of the state of the state of the state of the state of the state of t	Backup Name	Count	Type	17 End Date	Verified	Mounted	RMAN Cataloged	SCN	
Ip 10 0 0 151 03 25 2022 17.50 55 0653 1 I Log 93/25/2022 555:05 PM II Not Applicable False Not Cataloged 1789099 Ip 10 0 0 151 03 25 2022 17.50 55 0653 1 I Lug 03/25/2022 551:05 PM II Not Applicable Palse Not Cataloged 1789829 Ip 10 0 0 151 03 25 2022 17.50 55 0653 0 I Lug 03/25/2022 551:05 PM II Unverthed False Not Cataloged 1789829 Ip 10 0 0 151 03 25 2022 17.60 55 0653 0 I Data 03/25/2022 551:05 PM II Unverthed False Not Cataloged 1789829 Ip 10 0 0 151 03 25 2022 17.40 0.0 9758 1 I Log 03/25/2022 551:05 PM II Not False Not Cataloged 1789829 Ip 10 0 0 151 03 25 2022 17.40 0.0 9758 1 I Log 03/25/2022 550:05 PM II Not False Not Cataloged 1788110 Mount backups I Log 03/25/2022 550:05 PM III Not False Not Cataloged 1788110 Choose the host to mount the backup Ip-10-0-0-47.ec2.internal I IIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIII	ip-10-0-0-151_03-25-2022_18:55:01:0309_1		Log	03/25/2022 6:55:09 PM		Faise	Not Cataloged	1892563	
Ip-10-0-151_03-25-2022_17.50:55.0853_0 I Lug 03/25/2022 551:32 PM ® Avoit Applicable False Not Cataloged 1788832 Ip-10-0-151_03-25-2022_17.50:55.0853_0 1 Data 03/25/2022 551:05 PM ® Unverified False Not Cataloged 1788832 Ip-10-0-151_03-25-2022_17.60:55.0853_0 1 Log 03/25/2022 5:51:05 PM ® Unverified False Not Cataloged 1788832 Ip-10-0-151_03-25-2022_17.40:00:9758_1 1 Log 03/25/2022 5:40:08 PM ® Not False Not Cataloged 1788832 Mount backups	p-10-0-0-151_03-25-2022_18-40.00.9602_1	1	Log	03/25/2022 6:40:23 PM		False	Not Cataloged	1891375	
Applicable apploable ip-10-0-151,03-25-2022,17.5005,0853,0 1 Data d3/25/2022 5.51:05 PM III Unverified False Not Cataloged 1788832 Applicable Mount backups Choose the host to mount the backup (ip-10-0-0-47.ec2.internal	lp-10-0-0-151_03-25-2022_17.55.01.0197_1	Α.	Log	03/25/2022 5:55:09 PM		False	Not Cetaloged	1789099	
p-10-0-0151,03-25-2022_1740.00.9758_1 1 Log 03/25/2022 5:40:08 PM B Not False Not Cataloged 1788110 Mount backups Choose the host to mount the backup Ip-10-0-0-47.ec2.Internal	g-10-0-0-151_03-25-2022_17.50-55.0853_1	4	LUR	03/25/2022 551112 PM 🖨		False	Not Cataloged	1788679	
Mount backups × Choose the host to mount the backup ip-10-0-0-47.ec2.internal	p-10-0-0-151_03-25-2022_17.50.55-0853_0	1	Data	03/25/2022 5:51:05 PM	Unverified	Faise	Not Cataloged	1788832	
Choose the host to mount the backup	p-10-0-0-151_03-25-2022_17.40.00.9758_1	3	Log	03/25/2022 5:40:08 PM	Not	False	Not Cataloged	1788110	
	Mount backups								
	Choose the host to mount the backup			5-2022_17.50.55.0853_1/0	DRCL				
	Choose the host to mount the backup			5-2022_17.50.55.0853_1/0	DRCL				
	Choose the host to mount the backup			5-2022_17.50.55.0853_1/0	DRCL				
	Choose the host to mount the backup			5-2022_17.50.55.0853_1/(DRCL				
	Choose the host to mount the backup			5-2022_17.50.55.0853_1/0	DRCL				

4. Highlight the snapshot copy to be cloned for the replica, and click the Clone button to start the clone procedure.

ORCL Topology								_	
							and the second se	V	2
Manage Copies									
93 Backups						Summary	Card		
0 Clones						93 Backups			
Local copies						6 Data Ba	ckups		
						B7 Log Bac	Rups		
						0-Clones			
Primary Backup(s)									
search V						e it			1
Backup Name	Count	Type	17	End Date	Verified	Mounted	RMAN Cataloged	SCN	
p10-0-0-151_03-25-2022_17:55:01.0197_1	18	Log		03/25/2022 5:55:09 PM	Not Applicable	False	Not Cataloged	17890	199
p-10-0-0-151_03-25-2022_17:50.55.0853_1	<u>i</u>	Log		83/25/2022 5:51:12 PM 🗖	Not Applicable	False	Not Cataloged	17888	\$79
p-10-0-0-151_03-25-2022_17.50.55.0853_0	1	Data		03/25/2022 5:51:05 PM	Unverified	Talse	Not Cataloged	17088	32
p-10-0-0-151_03-25-2022_17.40.00.9758_1	3	Log		03/25/2022 5:40:08 PM	Not Applicable	False	Not Cataloged	17881	10
ip-10-0-0-151_03-25-2022_17.25.01.0539_1	3	Log		03/25/2022 5:25:08 PM 🖨	Not	False	Not Cataloged	17871	80

5. Change the replica copy name so that it is different from the primary database name. Click Next.

Clone from OR	CL			×
1 Name	Provide clone da	tabase SID		
2 Locations	Clone SID	ORCLREAD		
3 Credentials				
4 PreOps				
5 PostOps				
6 Notification				
7 Summary				
			 Previous	Next

6. Change the clone host to the standby EC2 host, accept the default naming, and click Next.

Clone from OF	RCL						×
1 Name	Select the host to	create a clone					
2 Locations	Clone host	ip-10-0-0-47.ec2.in	ternal	•			
3 Credentials	Datafile locations	0					
4 PreOps	/ora_nfs_data_ORC	LREAD				A. W	Reset
5 PostOps							
6 Notification	🕞 Control files 🚯		Í				
7 Summary	/ora_nfs_data_ORC	LREAD/ORCLREAD/col	ntrol/contro	ol01.ctl		× *	+ Reset
	🖸 Redo logs 🚯					1	iii 1
	Group		Size	Unit	Number of files		
	 RedoGroup 1 	×	128	MB	1	+	+
	/ora_nfs_date	a_ORCLREAD/ORCLRE/	AD/redolog	/redo04.log		×	Reset
	RedoGroup 2	×	128	MB	1	+	•
_							
						F	Previous Next

7. Change your Oracle home settings to match those configured for the target Oracle server host, and click Next.

Clone from OR	RCL		×
1 Name	Database Credentials for	the clone	
2 Locations	Credential name for sys user	None - + O	
3 Credentials	Database port	1521	
4 PreOps	Oracla Homa Sattings		
5 PostOps	Oracle Home Settings	/rdsdbbin/oracle	
6 Notification	Oracle OS User	rdsdb	
7 Summary	Oracle OS Group	database	
_			
		Previous Next	

8. Specify a recovery point using either time or the SCN and mounted archive log path.

Clone from OF	RCL	×
1 Name	Recover Database	
 2 Locations 3 Credentials 4 PreOps 5 PostOps 6 Notification 7 Summary 	 Until Cancel Date and Time Date-time format: MM/DD/YYYY hh:mm:ss Until SCN (System Change Number) 1788879 Specify external archive log locations (1788879) Specify external archive log locations (1788879) (1788879) (1788879)<td></td>	
	 Create new DBID ¹ Create tempfile for temporary tablespace ¹ Enter SQL queries to apply when clone is created Enter scripts to run after clone operation ¹ 	
	Previous	Next

9. Send the SMTP email settings if needed.

Clone from OR	CL		×
1 Name	Provide email sett	ings 🕕	
2 Locations	Email preference	Never -	
3 Credentials	From	From email	
4 PreOps	То	Email to	
	Subject	Notification	
5 PostOps	Attach job report		
6 Notification			
7 Summary			
		Previous	Next

10. Clone the job summary, and click Finish to launch the clone job.

	RCL
	Summary
lp-10-0-0-151_03-25-2022_17:50.55.0853_0	Clone from backup
ORCLREAD	Clone SID
ip-10-0-0-47.ec2.internal	Clone server
/rdsdbbin/oracle	Oracle home
rdsdb	Oracle OS user
database	Oracle OS group
/ora_nfs_data_ORCLREAD	Datafile mountpaths
/ora_nfs_data_ORCLREAD/ORCLREAD/control/control01.ctl	Control files
RedoGroup =1 TotalSize =128 Path =/ora_nfs_data_ORCLREAD/ORCLREAD/redolog/redol04.log RedoGroup =2 TotalSize =128 Path =/ora_nfs_data_ORCLREAD/ORCLREAD/redolog/redol03.log RedoGroup =3 TotalSize =128 Path =/ora_nfs_data_ORCLREAD/ORCLREAD/redolog/redol02.log RedoGroup =4 TotalSize =128 Path =/ora_nfs_data_ORCLREAD/ORCLREAD/redolog/redol01.log	Redo groups
Until SCN 1788879	Recovery scope
none	Prescript full path
	Prescript arguments
none	Postscript full path
	Postscript arguments
No	Send email
No	C22_08

11. Validate the replica clone by reviewing the clone job log.

Job Details	×
Clone from backup 'ip-10-0-0-151_03-25-2022_17.50.55.0853_0'	^
Clone from backup 'lp-10-0-0-151_03-25-2022_17.50.55.0853_0'	- 1
✓ ip-10-0-0-47.ec2.internal	
Prescripts	
 Query Host Information 	
 Prepare for Cloning 	
 Cloning Resources 	
✓ ► FileSystem Clone	
 Application Clone 	
 Postscripts 	
Register Clone	
V Unmount Clone	
 Data Collection 	
 Send EMS Messages 	
Task Name: ip-10-0-0-47.ec2.internal Start Time: 03/25/2022 9:08:32 PM End Time: 03/25/2022 9:12:03 PM	
View Logs Cancel Job	Close

The cloned database is registered in SnapCenter immediately.

4	Bracke Del	utere 🔸						
Dethered	View 0	atatase	Search databases	v				* *
e Resources	17 10	Name	Oracle Database Type	Host/Cluster	Resource Group	Potoes	Last Backup	Overall Status
Monitor		CRES.	tingle instance	ig-10-0-151.m2.internal	criti, Sull, pkcap criti, Jog, bikup	Oracle full backup Oracle log backup	GB/05/000-910/09/M 🛱	Backup suineede
al Augusta		ORCLADAD	tingle tratares	ip-10-0-47.ec2.internal				Not protected
A 1665								
1 Storage System								
E Setting								
A Arts								

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.
- 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

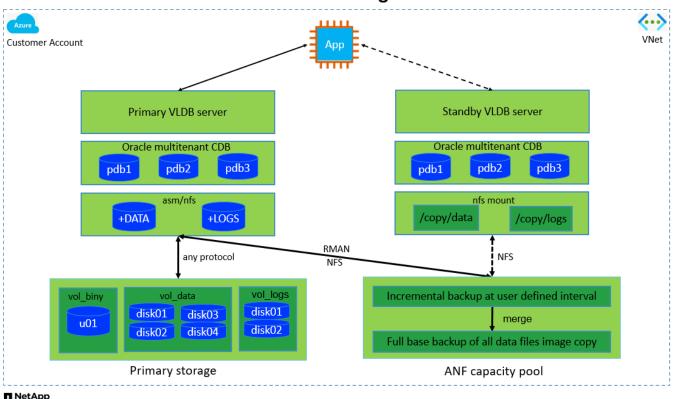
This solution is intended for the following people:

- 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

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.

≡ Microsoft Azure	${\cal P}$ Search resources, services, and docs (G+/)		5 0 © R	acao@netapp.com HYBRID CLOUD TME
Home > Azure NetApp Files > ANFAVSAcct Capacity pools >				
ANFAVSAcct Capacity pools «	Capacity pool	Acct/database) 🖉 ☆ …		>
🔎 Search 🛛 🐇 🔶 Add pool 🕐 Refresh	₽ Search «	🖉 Resize 🧷 Change QoS type 🗸 Enable Cool Access 📋	Delete	
Overview V Search pools	🚡 Overview	∧ Essentials		JSON View
Activity log Name Capacity	Activity log	Resource group ANFAVSRG	NetApp account ANFAVSAcct	
Access control (IAM)	Access control (IAM)	Subscription	Location	
 Tags Tags nimcp 2 TiB 	🔷 Tags	Hybrid Cloud TME Onprem	South Central US	
X Diagnose and solve problems	🗙 Diagnose and solve problems	Subscription ID 0efa2dfb-917c-4497-b56a-b3f4eadb8111	Service level Premium	
Settings	Settings	Size 2 TiB	QoS type Auto	
D Quota	Properties	Pool Allocated to Volume Throughput MiB/s	Pool Provisioned Throughput MiB/s	
Advisor recommendations	🔒 Locks	81.25	128 Cool Access	
Properties	Storage service		Disabled	
Locks	Volumes	Usage		
Azure NetApp Files	Monitoring	-		
Active Directory connections	Metrics	Current sum of all volume allocated sizes		
A Encryption				
Storage service	Automation			
E Capacity pools	Tasks (preview)			
Volumes				
Application volume groups	Help			
Data protection	O Support + Troubleshooting	63.5		
Snapshot policies		1.27 TIB ALLOCATED		
Sackups				
Backup vaults				
Backup Policies				

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

E Microsoft Azure	P Search resources, services, and doo	:s (G+/)		N 0 Ø	0 R		tapp.com
ome > Azure NetApp Files > ANFAVSAcct Capacity po	ools > database (ANFAVSAcct/database)						
ANFAVSAcct Capacity pools	« 🗧 database (ANFA _{Capacity pool}	VSAcct/database) Volume	S ☆ …				
P Search « + Add pool	🕐 Refresh 🔎 Search] « 🛛 🕂 Add volume 🕂 Add data replica	tion 💍 Refresh				
Overview	Overview	𝒴 Search volumes					
Activity log Name	↑↓ Capacity	Name ↑↓ Quota	ர்↓ Max. Throughput ↑்	Protocol type 1	Nount path	Ť.	Service
Access control (IAM)	Access control (IAM)	anne 1, 2000	6.25 MiB/s	NFSv3	172.30.136.68		
♦ Tags	2 TiB	ora-01-u02 500 GiB	31.25 MiB/s	NFSv3	172.30.136.68		
X Diagnose and solve problems	Diagnose and solve problems	a ora-01-u03 400 GiB	25 MiB/s	NFSv3	172.30.136.68		
Settings	Settings	ora-02-u01 100 GiB	6.25 MiB/s	NFSv3	172.30.136.68		
D Quota	Properties	ora-02-u02 100 GiB	6.25 MiB/s	NFSv3	172.30.136.68		
Advisor recommendations	Locks	ora-02-u03 100 GiB	6.25 MiB/s	NFSv3	172.30.136.68		
Properties	Storage service	4 6					Þ
Locks	Volumes						
Azure NetApp Files							
Active Directory connections	Monitoring						
Encryption	mi Metrics						
26.7 4.6.	Automation						
Storage service							
Capacity pools	Export template						
Volumes	Help						
Application volume groups	③ Support + Troubleshooting						
Data protection							
Snapshot policies							
ackups							

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

Basics Protocol Tags Revi	ew + create	
This page will help you create an Azur within your virtual network. Learn mo	e NetApp Files volume in your subscription and enable you to access the volume from reabout Azure NetApp Files $arsigma^2$	om
Volume details		
Volume name *	ora-01-u02-copy	~
Available quota (GiB) 🕕	748	
	74	8 Gie
Quota (GiB) * 🕕	500	~
	500	0 Gil
Available throughput (MiB/s) 🕠	46.75	
Иах. Throughput (MiB/s) 🕕	31.25	
nable Cool Access ①		
Coolness Period 🛈	31	
Cool Access Retrieval Policy 🕕	Default	\sim
/irtual network * 🕕	ANFAVSVal (172.30.136.64/26,172.30.137.128/25,172.30.152.0/27) Create new virtual network	\sim
Delegated subnet * 🕡	ANF_Sub (172.30.136.64/26) Create new subnet	\sim
Network features ①	Basic Standard	
Availability Zone 🕕	None	\sim
incryption key source 🛈		\sim
show advanced section		

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

configure access to your volume. coress otocol type onfiguration e path e path ora-01-u02-copy ersions NFSv3 ora-01-u02-copy ersions ora-01-u02-copy ora-01-u02-copy ersions ora-01-u02-copy ora-01-u02-copy ersions ora-01-u02-copy ora-01-u02-	enfigure access to your volume.	reate a volume				
cress otocol type NFS SMB Dual-protocol onfiguration e path ora-01-u02-copy ersions * NFSV3 Enabled Disabled otop Enabled Disabled Disabled Disabled Disabled O770 cure VMware Solution DataStore Origont policy mode to top ↓ Move to bottom Delete Index Allowed clients Access Root Access Chown Mode I 172.30.137.128/25,1√ Read & Write ✓ On ✓ Restricted ✓ ✓	configure access to your volume. coss otocol type NFS SMB Dual-protocol arbitrons Image: stress s	Basics Protocol Tags I	Review + create			
ccess otacol type ● NFS SMB Dual-protocol anfiguration e path * ② ora-01-u02-copy ersions * NF5V3 erberos Enabled ③ Disabled App Enabled ④ Disabled App Enabled ④ Disabled mix Permissions ③ 0770 cure VMware Solution DataStore ③ onfigure the volume's export policy. This can be edited later. Learn more of * Move up ↓ Move down ↑ Move to top ↓ Move to bottom ⑧ Delete I ndex Allowed clients Access Root Access Chown Mode I 1 172:30.137.128/25.1√ Read & Write ✓ On ✓ Restricted Review + create < Previous Next : Tags >	cccess otocol type ● NFS ● SMB ● Dual-protocol ora-01-u02-copy ersions * NF5v3 ✓ erberos ● Enabled ● Disabled App ● Enabled ● Disabled hix Permissions ● 0770 cure VMware Solution DataStore ● □ cport policy onfigure the volume's export policy. This can be edited later. Learn more 0° ↑ Move up ↓ Move down ↑ Move to top ↓ Move to bottom ● Delete I 1 172:30.137.128/25,1√ Read & Write ✓ On ✓ Restricted 1 172:30.137.128/25,1√ Read & Write ✓ On ✓ Restricted Review + create Previous Next : Tags >					
onfiguration le path * ○ orra-01-u02-copy ersions * NFSv3 erberos Enabled ● Disabled OAP Enabled ● Disabled nix Permissions ○ 0770 cure VMware Solution DataStore ○ export policy onfigure the volume's export policy. This can be edited later. Learn more of ↑ Move up ↓ Move down ↑ Move to top ↓ Move to bottom ● Delete Index Allowed clients Access Root Access Chown Mode I 1 12.230.137.128/25.1√ Read & Write ↓ On ↓ Restricted ✓	onfiguration le path * ○ ora-01-u02-copy ersions * NFSv3 erberos Enabled ● Disabled OAP Enabled ● Disabled nix Permissions ○ 0770 cure VMware Solution DataStore ○ cure VMware Solution DataStore ○ cord point cure VMware Solution DataStore ○ cure VMware Solution DataStore ○ cord point cure VMware Solution DataStore ○	Access				
e path * ① ora-01-u02-copy vrsions * NFSv3 orberos Enabled ● Disabled opap Enabled ● Disabled opap Enabled ● Disabled opap 0770 cure VMware Solution DataStore ① 0770 cure VMware Solution DataStore ① 0 opapt policy opapt policy. This can be edited later. Learn more © ↑ Move up ↓ Move down ↑ Move to top ↓ Move to bottom ● Delete Index Allowed clients Access Root Access Chown Mode 1 172.30.137.128/25,1√ Read & Write ↓ On ↓ Restricted	le path * ①	rotocol type	● NFS 〇	SMB O Dual-protocol	I	
e path * ① ora-01-u02-copy vrsions * NFSv3 orberos Enabled ● Disabled opap Enabled ● Disabled opap Enabled ● Disabled opap 0770 cure VMware Solution DataStore ① 0770 cure VMware Solution DataStore ① 0 opapt policy opapt policy. This can be edited later. Learn more © ↑ Move up ↓ Move down ↑ Move to top ↓ Move to bottom ● Delete Index Allowed clients Access Root Access Chown Mode 1 172.30.137.128/25,1√ Read & Write ↓ On ↓ Restricted	le path * ①	'aufiguestion				
ersions * NFSv3 ✓ erberos Enabled Disabled DAP Enabled Disabled Disabl	ersions * NFSv3		ora-01-u02-0	000		
erberos	erberos					
AAP Construction DataStore Construction Data	DAP DAP Dap <tr< td=""><td></td><td></td><td>Disabled</td><td></td><td></td></tr<>			Disabled		
nix Permissions ③ 0770 ture VMware Solution DataStore ③ □ sport policy onfigure the volume's export policy. This can be edited later. Learn more © ² ↑ Move up ↓ Move down ↑ Move to top ↓ Move to bottom Delete □ Index Allowed clients Access Root Access Chown Mode □ 1 172.30.137.128/25,1√ Read & Write ✓ On ✓ Restricted □ ↓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓	nix Permissions ③ 0770 ture VMware Solution DataStore ③ □ sport policy onfigure the volume's export policy. This can be edited later. Learn more r ² ↑ Move up ↓ Move down ↑ Move to top ↓ Move to bottom Delete □ Index Allowed clients Access Root Access Chown Mode □ 1 172.30.137.128/25,1√ Read & Write ✓ On ✓ Restricted □ ↓ ↑ ♥ ♥ ♥ ♥ ♥ ♥ ♥ ♥ ♥ ♥ ♥ ♥ ♥ ♥ ♥ ♥ ♥ ♥		<u> </u>	-		
ture VMware Solution DataStore ○	cure VMware Solution DataStore ○ sport policy onfigure the volume's export policy. This can be edited later. Learn more C ² ↑ Move up ↓ Move down ↑ Move to top ↓ Move to bottom Delete Index Allowed clients Access Root Access Chown Mode 1 1172.30.137.128/25,1~ Read & Write ~ On ~ Restricted ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓	DAP	Enabled			
sport policy Infigure the volume's export policy. This can be edited later. Learn more C ² Move up ↓ Move down ↑ Move to top ↓ Move to bottom Î Delete Index Allowed clients Access Root Access Chown Mode I 1 172.30.137.128/25,1√ Read & Write ✓ On ✓ Restricted I 1 172.30.137.128/25,1√ Read & Write ✓ On ✓ Restricted I 4 Previous Next : Tags >	sport policy Infigure the volume's export policy. This can be edited later. Learn more C ² Move up ↓ Move down ↑ Move to top ↓ Move to bottom Î Delete Index Allowed clients Access Root Access Chown Mode I 172.30.137.128/25,1√ Read & Write ∨ On ∨ Restricted Previous Next : Tags >	ix Permissions 🕕	0770			
Infigure the volume's export policy. This can be edited later. Learn more c ² Image: Move down The Move to top Image: Move to bottom Image: Delete Image: Image	nfigure the volume's export policy. This can be edited later. Learn more C ² Move up ↓ Move down ↑ Move to top ↓ Move to bottom Delete Index Allowed clients Access Root Access Chown Mode 1 172.30.137.128/25,1√ Read & Write ✓ On ✓ Restricted ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓	ure VMware Solution DataStore	0			
Review + create < Previous Next : Tags >	Review + create < Previous Next : Tags >	Index	Allowed clients	Access	Root Access	
		Index	Allowed clients	Access Read & Write	Root Access	Restricted
		Index	Allowed clients	Access Read & Write	Root Access	Restricted
a volume tag if desired.	a volume tag if desired.	Index	Allowed clients	Access Read & Write	Root Access	Restricted
a volume tag if desired.	a volume tag if desired.	Index	Allowed clients	Access Read & Write ✓	Root Access	Restricted
		Index I I Review + create	Allowed clients 172.30.137.128/25,1	Access Read & Write ✓	Root Access	Restricted
		Index 1 Review + create	Allowed clients 172.30.137.128/25,1	Access Read & Write ✓	Root Access	Restricted
		Index 1 Review + create	Allowed clients 172.30.137.128/25,1	Access Read & Write ✓	Root Access	Restricted
		Index I I I Review + create	Allowed clients 172.30.137.128/25,1	Access Read & Write ✓	Root Access	Restricted
		Index I I I Review + create	Allowed clients 172.30.137.128/25,1	Access Read & Write ✓	Root Access	Restricted
		Index I I I Review + create	Allowed clients 172.30.137.128/25,1	Access Read & Write ✓	Root Access	Restricted
		Index 1 Review + create	Allowed clients 172.30.137.128/25,1	Access Read & Write ✓	Root Access	Restricted

Create a volume	
Basics Protocol Tags Review + create	
Tags are name/value pairs that enable you to categorize reso multiple resources and resource groups. Learn more about t	purces and view consolidated billing by applying the same tag to ${\sf ags}\ {\sf C}$
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. Review and create the 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

sudo mkdir /nfsanf

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

8. Change mount point ownership to oracle:oisntall, 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';
   allocate channel c3 device type disk format
5>
'/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:
             Key Completion Time Filename/Handle
Type
_____
                        -----
                  1
                        18-MAR-24
Backup Set
 Backup Piece 1
                        18-MAR-24
/u03/orareco/NTAP1/autobackup/2024 03 18/o1 mf s 1163958359 04h19dg
r .bkp
Backup Set
                   2
                         18-MAR-24
                         18-MAR-24
 Backup Piece
                  2
/u03/orareco/NTAP1/autobackup/2024 03 18/o1 mf s 1163961675 0711m21
g .bkp
Backup Set
                   3
                        18-MAR-24
Backup Piece
                  3
                        18-MAR-24
/u03/orareco/NTAP1/autobackup/2024 03 18/o1 mf s 1163962888 08p6y71
x .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 18-MAR-24 Backup Set 7 7 Backup Piece 18-MAR-24 /u03/orareco/NTAP1/autobackup/2024 03 18/o1 mf s 1163966806 0dbyx34 4 .bkp 18-MAR-24 Backup Set 8 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 10 18-MAR-24 Backup Piece /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 12 18-MAR-24 Backup Piece /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 Obd3tqg3 .bkp RECID=5 STAMP=1163964705 deleted backup piece backup piece handle=/u03/orareco/NTAP1/autobackup/2024 03 18/o1 mf s 1163965895 Ochx6mzt .bkp RECID=6 STAMP=1163965906 deleted backup piece backup piece handle=/u03/orareco/NTAP1/autobackup/2024 03 18/o1 mf s 1163966806 Odbyx344 .bkp RECID=7 STAMP=1163966814 deleted backup piece backup piece handle=/u03/orareco/NTAP1/autobackup/2024 03 18/o1 mf s 1163968012 Ofgvg805 .bkp RECID=8 STAMP=1163968018 deleted backup piece backup piece handle=/u03/orareco/NTAP1/autobackup/2024 03 18/o1 mf s 1163968919 0q9x5t1v .bkp RECID=9 STAMP=1163968926 deleted backup piece backup piece handle=/u03/orareco/NTAP1/autobackup/2024 03 18/o1 mf s 1163969821 Oh4rfdzj .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 0j804wk8 .bkp RECID=11 STAMP=1163971032 Deleted 3 objects deleted backup piece backup piece handle=/u03/orareco/NTAP1/autobackup/2024 03 18/o1 mf s 1163971931 0k3pnn2o .bkp RECID=12 STAMP=1163971938 Deleted 3 objects

```
deleted backup piece
backup piece
handle=/u03/orareco/NTAP1/autobackup/2024 03 18/o1 mf s 1163972835
0kyg92t1 .bkp RECID=13 STAMP=1163972837
Deleted 4 objects
sql statement: alter system archive log current
Starting recover at 20-MAR-24
no copy of datafile 1 found to recover
no copy of datafile 3 found to recover
no copy of datafile 4 found to recover
no copy of datafile 31 found to recover
no copy of datafile 32 found to recover
Finished recover at 20-MAR-24
Starting backup at 20-MAR-24
no parent backup or copy of datafile 1 found
no parent backup or copy of datafile 3 found
no parent backup or copy of datafile 4 found
no parent backup or copy of datafile 19 found
no parent backup or copy of datafile 20 found
channel c1: starting datafile copy
input datafile file number=00021
name=/u02/oradata/NTAP1/NTAP1 pdb1/soe 01.dbf
channel c2: starting datafile copy
input datafile file number=00022
name=/u02/oradata/NTAP1/NTAP1 pdb1/soe 02.dbf
channel c3: starting datafile copy
input datafile file number=00023
name=/u02/oradata/NTAP1/NTAP1 pdb1/soe 03.dbf
channel c4: starting datafile copy
input datafile file number=00024
name=/u02/oradata/NTAP1/NTAP1 pdb1/soe 04.dbf
output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-
SOE FNO-22 0g2m6brl tag=ORACOPYBKUPONANF LEVEL 0 RECID=4
STAMP=1164132108
channel c2: datafile copy complete, elapsed time: 01:06:39
channel c2: starting datafile copy
input datafile file number=00025
name=/u02/oradata/NTAP1/NTAP1 pdb1/soe 05.dbf
```

output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-24 0i2m6brl tag=ORACOPYBKUPONANF LEVEL 0 RECID=5 STAMP=1164132121 channel c4: datafile copy complete, elapsed time: 01:06:45 channel c4: starting datafile copy input datafile file number=00026 name=/u02/oradata/NTAP1/NTAP1 pdb1/soe 06.dbf output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-23 0h2m6brl tag=ORACOPYBKUPONANF LEVEL 0 RECID=6 STAMP=1164132198 channel c3: datafile copy complete, elapsed time: 01:08:05 channel c3: starting datafile copy input datafile file number=00027 name=/u02/oradata/NTAP1/NTAP1 pdb1/soe 07.dbf output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-21 0f2m6brl tag=ORACOPYBKUPONANF LEVEL 0 RECID=7 STAMP=1164132248 channel c1: datafile copy complete, elapsed time: 01:08:57 channel c1: starting datafile copy input datafile file number=00028 name=/u02/oradata/NTAP1/NTAP1 pdb1/soe 08.dbf output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-25 0j2m6fol tag=ORACOPYBKUPONANF LEVEL 0 RECID=9 STAMP=1164136123 channel c2: datafile copy complete, elapsed time: 01:06:46 channel c2: starting datafile copy input datafile file number=00029 name=/u02/oradata/NTAP1/NTAP1 pdb1/soe 09.dbf output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-26 0k2m6fot tag=ORACOPYBKUPONANF LEVEL 0 RECID=8 STAMP=1164136113 channel c4: datafile copy complete, elapsed time: 01:06:36 channel c4: starting datafile copy input datafile file number=00030 name=/u02/oradata/NTAP1/NTAP1 pdb1/soe 10.dbf output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-27 012m6frc 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 On2m6jlr 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 0u2m6ngs tag=ORACOPYBKUPONANF LEVEL 0 RECID=21 STAMP=1164140471 channel c2: datafile copy complete, elapsed time: 00:01:33 channel c2: starting datafile copy input datafile file number=00013 name=/u02/oradata/NTAP1/NTAP1 pdb2/system01.dbf output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SYSAUX FNO-18 112m6nrt tag=ORACOPYBKUPONANF LEVEL 0 RECID=22 STAMP=1164140476 channel c3: datafile copy complete, elapsed time: 00:00:57 channel c3: starting datafile copy input datafile file number=00017 name=/u02/oradata/NTAP1/NTAP1 pdb3/system01.dbf output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SYSAUX FNO-6 122m6nti tag=ORACOPYBKUPONANF LEVEL 0 RECID=23 STAMP=1164140488 channel c4: datafile copy complete, elapsed time: 00:00:25 channel c4: starting datafile copy input datafile file number=00005 name=/u02/oradata/NTAP1/pdbseed/system01.dbf output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SYSTEM FNO-13 142m6ntp tag=ORACOPYBKUPONANF LEVEL 0 RECID=24 STAMP=1164140532 channel c2: datafile copy complete, elapsed time: 00:01:06 channel c2: starting datafile copy

input datafile file number=00008 name=/u02/oradata/NTAP1/pdbseed/undotbs01.dbf output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SYSTEM FNO-17 152m6nts tag=ORACOPYBKUPONANF LEVEL 0 RECID=25 STAMP=1164140539 channel c3: datafile copy complete, elapsed time: 00:01:03 channel c3: starting datafile copy input datafile file number=00015 name=/u02/oradata/NTAP1/NTAP1 pdb2/undotbs01.dbf output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SYSTEM FNO-9 132m6ntm tag=ORACOPYBKUPONANF LEVEL 0 RECID=26 STAMP=1164140541 channel c1: datafile copy complete, elapsed time: 00:01:13 channel c1: starting datafile copy input datafile file number=00019 name=/u02/oradata/NTAP1/NTAP1 pdb3/undotbs01.dbf output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SYSTEM FNO-5 162m6nuc tag=ORACOPYBKUPONANF LEVEL 0 RECID=27 STAMP=1164140541 channel c4: datafile copy complete, elapsed time: 00:00:41 channel c4: starting datafile copy input datafile file number=00007 name=/u02/oradata/NTAP1/users01.dbf output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-UNDOTBS1 FNO-8 172m6nvr tag=ORACOPYBKUPONANF LEVEL 0 RECID=28 STAMP=1164140552 channel c2: datafile copy complete, elapsed time: 00:00:16 channel c2: starting datafile copy input datafile file number=00012 name=/u02/oradata/NTAP1/NTAP1 pdb1/users01.dbf output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-UNDOTBS1 FNO-15 182m6nvs tag=ORACOPYBKUPONANF LEVEL 0 RECID=30 STAMP=1164140561 channel c3: datafile copy complete, elapsed time: 00:00:24 channel c3: starting datafile copy input datafile file number=00016 name=/u02/oradata/NTAP1/NTAP1 pdb2/users01.dbf output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-USERS FNO-7 1a2m6o01 tag=ORACOPYBKUPONANF LEVEL 0 RECID=29 STAMP=1164140560 channel c4: datafile copy complete, elapsed time: 00:00:16 channel c4: starting datafile copy input datafile file number=00020 name=/u02/oradata/NTAP1/NTAP1 pdb3/users01.dbf output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-UNDOTBS1 FNO-19 192m6nvv tag=ORACOPYBKUPONANF LEVEL 0 RECID=31 STAMP=1164140564

```
channel c1: datafile copy complete, elapsed time: 00:00:21
output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-
USERS FNO-12 1b2m6o0e tag=ORACOPYBKUPONANF LEVEL 0 RECID=32
STAMP=1164140564
channel c2: datafile copy complete, elapsed time: 00:00:02
output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-
USERS FNO-16 1c2m6o0k tag=ORACOPYBKUPONANF LEVEL 0 RECID=34
STAMP=1164140565
channel c3: datafile copy complete, elapsed time: 00:00:01
output file name=/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-
USERS FNO-20 1d2m6o0k tag=ORACOPYBKUPONANF LEVEL 0 RECID=33
STAMP=1164140565
channel c4: datafile copy complete, elapsed time: 00:00:01
Finished backup at 20-MAR-24
Starting Control File and SPFILE Autobackup at 20-MAR-24
piece
handle=/nfsanf/archlog/NTAP1/autobackup/2024 03 20/o1 mf s 116414056
5 5g56ypks .bkp comment=NONE
Finished Control File and SPFILE Autobackup at 20-MAR-24
released channel: c1
released channel: c2
released channel: c3
released channel: c4
RMAN> **end-of-file**
RMAN>
```

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

RMAN> list copy of database tag 'OraCopyBKUPonANF_level_0'; List of Datafile Copies _____ Key File S Completion Time Ckp SCN Ckp Time Sparse 1 A 20-MAR-24 4161498 14 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 A 20-MAR-24 4161589 20-MAR-24 4 NO Name: /nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-UNDOTBS1 FNO-4 0t2m6nml Tag: ORACOPYBKUPONANF LEVEL 0 5 A 20-MAR-24 2379694 18-MAR-24 NO 27 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 A 20-MAR-24 2379694 6 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 7 A 20-MAR-24 4161872 20-MAR-24 NO 29 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 A 20-MAR-24 4161835 20-MAR-24 NO 9 Name: /nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SYSTEM FNO-9 132m6ntm Tag: ORACOPYBKUPONANF LEVEL 0 Container ID: 3, PDB Name: NTAP1 PDB1 10 A 20-MAR-24 19 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 11 A 20-MAR-24 4161780 21 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 12 A 20-MAR-24 4161880 20-MAR-24 32 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 13 A 20-MAR-24 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 19 A 20-MAR-24 4161869 20-MAR-24 31 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 20 A 20-MAR-24 4161887 20-MAR-24 33 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 Of2m6brl Tag: ORACOPYBKUPONANF LEVEL 0 Container ID: 3, PDB Name: NTAP1 PDB1 22 A 20-MAR-24 4152518 20-MAR-24 NO 4 Name: /nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-22 Og2m6brl 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 Oh2m6brl Tag: ORACOPYBKUPONANF LEVEL 0 Container ID: 3, PDB Name: NTAP1 PDB1 24 A 20-MAR-24 4152529 20-MAR-24 5 NO Name: /nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-24 Oi2m6brl 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 26 A 20-MAR-24 4156130 20-MAR-24 NO 8 Name: /nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-26 Ok2m6fot Tag: ORACOPYBKUPONANF LEVEL 0 Container ID: 3, PDB Name: NTAP1 PDB1

27 A 20-MAR-24 4156159 20-MAR-24 10 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 28 A 20-MAR-24 4156183 20-MAR-24 11 NO Name: /nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-28 Om2m6fsu Tag: ORACOPYBKUPONANF LEVEL 0 Container ID: 3, PDB Name: NTAP1 PDB1 29 A 20-MAR-24 12 4158795 20-MAR-24 NO Name: /nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-29 On2m6jlr Tag: ORACOPYBKUPONANF LEVEL 0 Container ID: 3, PDB Name: NTAP1 PDB1 A 20-MAR-24 4158803 13 30 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 Oq2m6jsi Tag: ORACOPYBKUPONANF LEVEL 0 Container ID: 3, PDB Name: NTAP1 PDB1

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

1060 1 SYSTEM YES /u02/oradata/NTAP1/system01.dbf 1000 SYSAUX 3 NO /u02/oradata/NTAP1/sysaux01.dbf 695 UNDOTBS1 4 YES /u02/oradata/NTAP1/undotbs01.dbf 400 PDB\$SEED:SYSTEM 5 NO /u02/oradata/NTAP1/pdbseed/system01.dbf 440 PDB\$SEED:SYSAUX 6 NO /u02/oradata/NTAP1/pdbseed/sysaux01.dbf 5 USERS 7 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 NTAP1 PDB1:SYSAUX 10 520 NO /u02/oradata/NTAP1/NTAP1 pdb1/sysaux01.dbf NTAP1 PDB1:UNDOTBS1 YES 11 580 /u02/oradata/NTAP1/NTAP1 pdb1/undotbs01.dbf 12 5 NTAP1 PDB1:USERS NO /u02/oradata/NTAP1/NTAP1 pdb1/users01.dbf 410 NTAP1 PDB2:SYSTEM 13 YES /u02/oradata/NTAP1/NTAP1 pdb2/system01.dbf 500 14 NTAP1 PDB2:SYSAUX NO /u02/oradata/NTAP1/NTAP1 pdb2/sysaux01.dbf 235 NTAP1 PDB2:UNDOTBS1 YES 15 /u02/oradata/NTAP1/NTAP1 pdb2/undotbs01.dbf NTAP1 PDB2:USERS 16 5 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 NTAP1 PDB3:UNDOTBS1 YES 19 235 /u02/oradata/NTAP1/NTAP1 pdb3/undotbs01.dbf NTAP1 PDB3:USERS 20 5 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 NTAP1 PDB1:SOE 23 31744 NO /u02/oradata/NTAP1/NTAP1 pdb1/soe 03.dbf 24 31744 NTAP1 PDB1:SOE NO

____ ____

```
/u02/oradata/NTAP1/NTAP1 pdb1/soe 04.dbf
          NTAP1 PDB1:SOE NO
25
   31744
/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
   31744 NTAP1 PDB1:TEMP
3
                            32767
/u02/oradata/NTAP1/NTAP1 pdb1/temp01.dbf
          NTAP1 PDB2:TEMP 32767
4
   123
/u02/oradata/NTAP1/NTAP1 pdb2/temp01.dbf
   123 NTAP1 PDB3:TEMP 32767
5
/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 -1 /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_0f2m6br1
-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 012m6frc
-rw-r---- 1 oracle oinstall 33286004736 Mar 20 19:12 data D-
NTAP1 I-2441823937 TS-SOE FNO-28 Om2m6fsu
-rw-r---- 1 oracle oinstall 33286004736 Mar 20 20:14 data D-
NTAP1 I-2441823937 TS-SOE FNO-29 On2m6jlr
-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_1b2m600e -rw-r----- 1 oracle oinstall 5251072 Mar 20 20:22 data_D-NTAP1_I-2441823937_TS-USERS_FNO-16_1c2m600k -rw-r----- 1 oracle oinstall 5251072 Mar 20 20:22 data_D-NTAP1_I-2441823937_TS-USERS_FNO-16_1c2m600k -rw-r----- 1 oracle oinstall 5251072 Mar 20 20:22 data_D-NTAP1_I-2441823937_TS-USERS_FNO-20_1d2m600k -rw-r----- 1 oracle oinstall 5251072 Mar 20 20:22 data_D-NTAP1_I-2441823937_TS-USERS_FNO-7_1a2m6001 [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 -
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

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 0u2m6ngs" 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 012m6frc" 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 On2m6jlr" 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 0flsd788
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>
```

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 012m6frc
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-28 0m2m6fsu
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-29 0n2m6jlr
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-30 0o2m6jlr
/nfsanf/oracopy/data D-NTAP1_I-2441823937_TS-SOE_FNO-31_0p2m6jrb
NAME
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-32 0q2m6jsi
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SYSAUX FNO-10 0v2m6nqs
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SYSAUX FNO-14 102m6nr3
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SYSAUX FNO-18 112m6nrt
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SYSAUX FNO-3 0s2m6nl1
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SYSAUX FNO-6 122m6nti
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SYSTEM FNO-13 142m6ntp
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SYSTEM FNO-17 152m6nts
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SYSTEM FNO-1 0r2m6nhk
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SYSTEM FNO-5 162m6nuc
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SYSTEM FNO-9 132m6ntm
```

NAME

```
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-UNDOTBS1 FNO-
11 Ou2m6nqs
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-UNDOTBS1 FNO-
15 182m6nvs
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-UNDOTBS1 FNO-
19 192m6nvv
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-UNDOTBS1 FNO-4 0t2m6nml
/nfsanf/oracopy/data D-NTAP1 I-2441823937_TS-UNDOTBS1_FNO-8_172m6nvr
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-USERS FNO-12 1b2m6o0e
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-USERS FNO-16 1c2m6o0k
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-USERS FNO-20 1d2m6o0k
/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-USERS FNO-7 1a2m6o01
/u02/oradata/NTAP1/NTAP1 pdb1/temp01.dbf
/u02/oradata/NTAP1/NTAP1 pdb1/temp02.dbf
NAME
_____
/u02/oradata/NTAP1/NTAP1 pdb2/temp01.dbf
/u02/oradata/NTAP1/NTAP1 pdb3/temp01.dbf
/u02/oradata/NTAP1/control01.ctl
/u02/oradata/NTAP1/pdbseed/temp012024-03-18 16-07-32-463-PM.dbf
/u02/oradata/NTAP1/temp01.dbf
/u03/orareco/NTAP1/control02.ctl
/u03/orareco/NTAP1/onlinelog/redo01.log
/u03/orareco/NTAP1/onlinelog/redo02.log
/u03/orareco/NTAP1/onlinelog/redo03.log
42 rows selected.
```

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

SQL> alter session set container=ntap1_pdb1;
Session altered.
SQL> select * from test;
ID
DT
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 kernal 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,rsize=262144,wsize=262144,noin
tr
```

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

-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 212mbcdn 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 202mbcdv 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 li2m9cap

-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 1o2m9qd4 -rw-r----. 1 oracle oinstall 33286004736 Mar 22 15:31 data D-NTAP1 I-2441823937 TS-SOE FNO-31 1r2m9kfk -rw-r----. 1 oracle oinstall 33286004736 Mar 22 15:31 data D-NTAP1 I-2441823937 TS-SOE FNO-29 1p2m9ju6 -rw-r----. 1 oracle oinstall 33286004736 Mar 22 15:31 data D-NTAP1 I-2441823937 TS-SOE FNO-32 1s2m9kgg -rw-r----. 1 oracle oinstall 33286004736 Mar 22 15:31 data D-NTAP1 I-2441823937 TS-SOE FNO-25 112m9g3u -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 392mbg1i 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 3k2mbq3a 116 1 1 -rw-r----. 1 oracle oinstall 57344 Mar 22 15:38 312mbg3f 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
ol 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
ol 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
ol mf 1 8 9yctxjgy .arc
-r--r-. 1 oracle oinstall 185494528 Mar 22 15:24
ol mf 1 9 9yfrj0bl .arc
-r--r--. 1 oracle oinstall 134470144 Mar 22 15:29
ol 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 - 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
                               9174800 bytes
Fixed Size
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
  ---- ---- - -----
                                 ---- -----
                                            _____ ____
      1 A 22-MAR-24
82
                           4598427 22-MAR-24
                                                  NO
      Name: /nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-
SYSTEM FNO-1 1t2m9nij
       Tag: ORACOPYBKUPONANF LEVEL 0
       3 A 22-MAR-24 4598423 22-MAR-24 NO
83
       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
       5 A 21-MAR-24 2379694 18-MAR-24
58
                                                    NO
       Name: /nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-
SYSTEM FNO-5 282m90em
       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
           A 21-MAR-24 2379694 18-MAR-24 NO
       8
       Name: /nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-
UNDOTBS1 FNO-8 292m90em
       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 10 A 22-MAR-24 68 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 11 A 22-MAR-24 4598304 66 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 17 A 22-MAR-24 4598450 91 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 20 A 22-MAR-24 93 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 21 A 22-MAR-24 4598318 22-MAR-24 81 NO Name: /nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-21 lh2m9cap Tag: ORACOPYBKUPONANF LEVEL 0 Container ID: 3, PDB Name: NTAP1 PDB1 22 A 22-MAR-24 4598304 22-MAR-24 NO 72 Name: /nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-22 li2m9cap Tag: ORACOPYBKUPONANF LEVEL 0 Container ID: 3, PDB Name: NTAP1 PDB1 4598308 22-MAR-24 23 A 22-MAR-24 73 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

24 A 22-MAR-24 4598313 22-MAR-24 80 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 112m9q3u 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 27 A 22-MAR-24 70 4598308 22-MAR-24 NO Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-27 ln2m9gcg Tag: ORACOPYBKUPONANF LEVEL 0 Container ID: 3, PDB Name: NTAP1 PDB1 4598313 22-MAR-24 NO 75 28 A 22-MAR-24 Name: /nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-28 1o2m9qd4 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 30 A 22-MAR-24 67 4598304 22-MAR-24 NO Name: /nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-30 1g2m9k7a Tag: ORACOPYBKUPONANF LEVEL 0 Container ID: 3, PDB Name: NTAP1 PDB1 76 31 A 22-MAR-24 4598308 22-MAR-24 NO Name: /nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-31 1r2m9kfk Tag: ORACOPYBKUPONANF LEVEL 0

```
Container ID: 3, PDB Name: NTAP1_PDB1

78 32 A 22-MAR-24 4598313 22-MAR-24 NO

Name: /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-

32_1s2m9kgg

Tag: ORACOPYBKUPONANF_LEVEL_0

Container ID: 3, PDB Name: NTAP1_PDB1
```

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

```
RMAN> switch database to copy;
Starting implicit crosscheck backup at 22-MAR-24
allocated channel: ORA DISK 1
channel ORA DISK 1: SID=12 device type=DISK
Crosschecked 33 objects
Finished implicit crosscheck backup at 22-MAR-24
Starting implicit crosscheck copy at 22-MAR-24
using channel ORA DISK 1
Crosschecked 31 objects
Finished implicit crosscheck copy at 22-MAR-24
searching for all files in the recovery area
cataloging files...
cataloging done
List of Cataloged Files
_____
File Name:
/nfsanf/archlog/NTAP1/autobackup/2024 03 20/o1 mf s 1164140565 5q56
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 282m90em"
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 112m9g3u" datafile 26 switched to datafile copy "/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-26 1m2m9g9j" datafile 27 switched to datafile copy "/nfsanf/oracopy/data D-NTAP1 I-2441823937_TS-SOE_FNO-27_1n2m9gcg" datafile 28 switched to datafile copy "/nfsanf/oracopy/data D-NTAP1 I-2441823937 TS-SOE FNO-28 1o2m9gd4" datafile 29 switched to datafile copy "/nfsanf/oracopy/data D-NTAP1_I-2441823937 TS-SOE FNO-29 1p2m9ju6"

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

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

```
RMAN> run {
2> set until sequence=11;
3> recover database;
4> }
executing command: SET until clause
Starting recover at 22-MAR-24
using channel ORA DISK 1
starting media recovery
archived log for thread 1 with sequence 4 is already on disk as file
/nfsanf/archlog/NTAP1/archivelog/2024 03 22/o1 mf 1 4 9y6gn5co .arc
archived log for thread 1 with sequence 5 is already on disk as file
/nfsanf/archlog/NTAP1/archivelog/2024 03 22/o1 mf 1 5 9y7p68s6 .arc
archived log for thread 1 with sequence 6 is already on disk as file
/nfsanf/archlog/NTAP1/archivelog/2024 03 22/01 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 22/01 mf_1 10 9yomybbc_.ar
С
archived log file
name=/nfsanf/archlog/NTAP1/archivelog/2024 03 22/o1 mf 1 4 9y6gn5co
.arc thread=1 sequence=4
archived log file
name=/nfsanf/archlog/NTAP1/archivelog/2024 03 22/o1 mf 1 5 9y7p68s6
.arc thread=1 sequence=5
archived log file
name=/nfsanf/archlog/NTAP1/archivelog/2024 03 22/o1 mf 1 6 9y8ygtss
```

.arc thread=1 sequence=6 archived log file name=/nfsanf/archlog/NTAP1/archivelog/2024 03 22/o1 mf 1 7 9ybjdp55 .arc thread=1 sequence=7 archived log file name=/nfsanf/archlog/NTAP1/archivelog/2024 03 22/o1 mf 1 8 9yctxjgy .arc thread=1 sequence=8 archived log file name=/nfsanf/archlog/NTAP1/archivelog/2024 03 22/o1 mf 1 9 9yfrj0b1 .arc thread=1 sequence=9 archived log file name=/nfsanf/archlog/NTAP1/archivelog/2024 03 22/o1 mf 1 10 9yomybb c .arc thread=1 sequence=10 media recovery complete, elapsed time: 00:01:17 Finished recover at 22-MAR-24 RMAN> exit

Recovery Manager complete.



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

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

SQL> select 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
                                       OPEN MODE RESTRICTED
   CON ID CON NAME
                  ----- -----
        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_272m9oe1 /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-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_112m9g3u /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-26_1m2m9g9j /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-27_1n2m9gcg /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-28_1o2m9gd4 /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-29_1p2m9ju6 /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-30_1q2m9k7a /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-31_1r2m9kfk /nfsanf/oracopy/data_D-NTAP1_I-2441823937_TS-SOE_FNO-32_1s2m9kgg

```
31 rows selected.
SQL> select member from v$logfile;
MEMBER
_____
/nfsanf/oracopy/redo03.log
/nfsanf/oracopy/redo02.log
/nfsanf/oracopy/redo01.log
SQL> select name from v$controlfile;
NAME
_____
_____
/nfsanf/oracopy/NTAP1.ctl
SQL> alter session set container=ntap1_pdb1;
Session altered.
SQL> select * from test;
     ID
_____
DT
_____
_____
EVENT
_____
_____
      1
21-MAR-24 03.15.03.000000 PM
test oracle incremental merge switch to copy
      2
22-MAR-24 02.22.06.000000 PM
test recovery on a new Azure VM host with image copy on \ensuremath{\mathsf{ANF}}
```

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

SQL> select name from v\$tempfile;

```
_____
/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>
```

- 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,rsize=262144,wsize=262144,noin
tr 0 0
```

- Update the Oracle init file from primary databse 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-premisses and Azure cloud

Audience

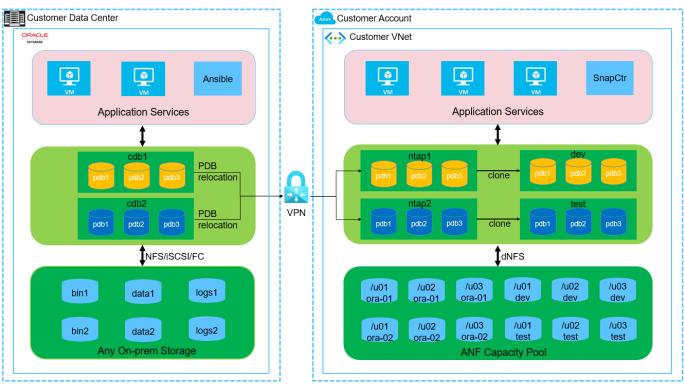
This solution is intended for the following people:

- 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



NetApp

Hardware and software components

	Hardware	
Azure NetApp Files	Current offering in Azure by Microsoft	One capacity pool with Premium service level
Azure VM for DB server	Standard_B4ms - 4 vCPUs, 16GiB	Two Linux virtual machine instances for concurrent deployment
Azure VM for SnapCenter	Standard_B4ms - 4 vCPUs, 16GiB	One Windows virtual machine instance
	Software	
RedHat Linux	RHEL Linux 8.6 (LVM) - x64 Gen2	Deployed RedHat subscription for testing
Windows Server	2022 DataCenter; Azure Edition Hotpatch - x64 Gen2	Hosting SnapCenter server
Oracle Database	Version 19.18	Applied RU patch p34765931_190000_Linux-x86- 64.zip
Oracle OPatch	Version 12.2.0.1.36	Latest patch p6880880_190000_Linux-x86- 64.zip
SnapCenter Server	Version 5.0	Workgroup deployment
Open JDK	Version java-11-openjdk	SnapCenter plugin requirement on DB VMs
NFS	Version 3.0	Oracle dNFS enabled
Ansible	core 2.16.2	Python 3.6.8

Oracle database configuration in the lab environment

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

Key factors for deployment consideration

- Oracle database storage layout. In this automated Oracle deployment, we provision three database volumes for each database to host Oracle binary, data, and logs by default. The volumes are mounted on Oracle DB server as /u01 binary, /u02 data, /u03 logs via NFS. Dual control files are configured on /u02 and /u03 mount points for redundancy.
- **Multiple DB servers deployment.** The automation solution can deploy an Oracle container database to multiple DB servers in a single Ansible playbook run. Regardless of the number of DB servers, the playbook execution remains the same. You can deploy multiple container databases to a single VM

instance by repeating the deployment with different database instance IDs (Oracle SID). But ensure there is sufficient memory on the host to support deployed databases.

- **dNFS configuration.** By using dNFS (available since Oracle 11g), an Oracle database running on an Azure Virtual Machine can drive significantly more I/O than the native NFS client. Automated Oracle deployment configures dNFS on NFSv3 by default.
- Allocate large size volume to speed up deployment. ANF file system IO throughput is regulated based on the size of volume. For initial deployment, allocate large size volumes can speed up the deployment. The volumes subsequently can be downsized dynamically without application impact.
- **Database backup.** NetApp provides a SnapCenter software suite for database backup, restore, and cloning with a user-friendly UI interface. NetApp recommends implementing such a management tool to achieve fast (under a minute) snapshot backup, quick (minutes) database restore, and database clone.

Solution deployment

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

Prerequisites for deployment

Deployment requires the following prerequisites.

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



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

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

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

- 3. From the Azure cloud portal, provision a Windows server to run the NetApp SnapCenter UI tool with the latest version. Refer to the following link for details: Install the SnapCenter Server
- 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-premisses 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
azureuser 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 ANF, linux and oracle
######
*****
##
### ANF env specific config variables ###
# Prerequisite to create three volumes in NetApp storage pool from
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.
# NFS lif ip address to access database volumes in ANF storage pool
(retrievable from cloud dashboard)
nfs lif: 172.30.136.68
****
### Linux env specific config variables ###
redhat sub username: XXXXXXXX
redhat sub password: XXXXXXXX
*******
### 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: XXXXXXX
```

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 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
                        0 0
xfs
       defaults
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
                                                         defaults
                                                 xfs
0 0
/dev/mapper/rootvg-usrlv /usr
                                                         defaults
                                                 xfs
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,wsize=65536 0 0
172.30.136.68:/ora-01-u02 /u02 nfs
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsize=65536 0 0
172.30.136.68:/ora-01-u03 /u03 nfs
rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsize=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	08	/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	48	/home
	/dev/sda15	495M	5.9M	489M	28	/boot/efi
	/dev/mapper/rootvg-tmplv	12G	8.4G	3.7G	70%	/tmp
	tmpfs	1.6G	0	1.6G	08	/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	18	/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 - 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
                         TNSLSNR for Linux: Version 19.0.0.0.0 -
Version
Production
                         01-FEB-2024 16:13:49
Start Date
                         0 days 0 hr. 11 min. 49 sec
Uptime
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.hr2z2nbmhnqutdsxqscjtuxizd.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 -
Production
Version 19.18.0.0.0
SQL> select name, open mode, log mode from v$database;
NAME OPEN MODE
                      LOG MODE
----- ------
NTAP1 READ WRITE
                          ARCHIVELOG
SQL> show pdbs
                                      OPEN MODE RESTRICTED
  CON ID CON NAME
                _____ ____
        2 PDB$SEED
                                      READ ONLY NO
        3 NTAP1 PDB1
                                      READ WRITE NO
                                     READ WRITE NO
        4 NTAP1 PDB2
```

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/undotbs01.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_pdb3/system01.dbf
/u02/oradata/NTAP1/NTAP1_pdb3/sysaux01.dbf
/u02/oradata/NTAP1/NTAP1_pdb3/undotbs01.dbf
/u02/oradata/NTAP1/NTAP1_pdb3/undotbs01.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/onlinelog/redo03.log

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

C Not secure https://172.30.137.187:5500/em/login?reason=need_login&retur	mon-yenysner			A 🗘 🗘
		CLE ENTERPRIS		
	Username system Password	<u></u>		
	Container Name	761		
	Log in			
			ORACLE	
Copyright 2013, 2020, Oracle and	d/or its affiliates. All rights reserved.			
← C S Not secure Https://172.30.137.187.5500/em/shell				A* ☆ 田 倉
V Voc secure mitpa//1/2.50.157.107.5500/em/sites				~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~
CRACLE Enterprise Manager Database Express				
CRACLE Enterprise Manager Database Express INTAP1 (19.18.0.0.0) Performance * Storage * Database Home				
■NTAP1(19.18.0.0.0) Performance v Storage v Database Home Time Zone				
Storage TAP1 (19.18.0.0,0) Performance Storage				1 min Auto-Refresh 🔹
■NTAP1(19.18.0.0.0) Performance v Storage v Database Home Time Zone	Performance			1 min Auto-Refresh 🛛 🔻 🗌
Storage Storage Storage Storage Storage Storage Storage Storage Up Time 34 minutes. 43 seconds	Performance Activity Services Containers			1 min Auto-Refresh 🛛 🔻 🗌
Status Up Time 34 minutes, 43 seconds Type Single Instance (NTAP1) CCB (3 PDB(o)				
Storage St	Activity Services Containers			Cther Cther Ketvork
Status Up Time 34 minutes. 43 seconds Type Single Instance (NTAP1) CD8 3 PD8(0) Version 19.18.00.0 Enterprise Edition Platform Name Linux 86 64-bit Treed 1	Activity Services Containers			Cother Network System 10 User 10
Status Variation Status Up Time 34 minutes. 43 seconds Type Single Instance (NTAP1) CD8 (3 PD8(s)) Version 13 18.0.00 Enterprise Edition Platform Name Linux x86 64-bit Thread 1 Archiver Stated	Activity Services Containers 40	040500 PM 041020 PM 041540 PM 0421	00 PM 042620 PM 043140 PM 043700 PM	Cther Netvork System (/O User (/O C CD
Status Up Time 34 minutes. 43 seconds Type Single Instance (NTAP1) CD8 3 PD8(0) Version 19.18.00.0 Enterprise Edition Platform Name Linux 86 64-bit Treed 1	Activity Services Containers	040500 PM 041020 PM 041540 PM 0421	00 PM 042620 PM 043140 PM 043700 PM	Cther Netvork System (/O User (/O C CD
■ NtAP1 (1918.0.00) Performance ▼ Statabase Home Time Zone Browser (GMT-00.00) ▼ Status Up Time 34 minutes. 43 seconds Up Time 34 minutes. 43 seconds Type 5 ingle instance (NTAP1) CD8 3 PDB(tr) Version 19.18.0.00 Enterprise Edition Platform March 1 Archiver Started Last Backup Time N/A Incident(s) € 1	Activity Services Containers 40	040500 PM 041020 PM 041540 PM 0421;	00 PM 042620 PM 043140 PM 043700 PM	Cther Netvork System (/O User (/O C CD
■ NAP1 (1918.0.0.0) Performance ▼ Storage ▼ ■ Database Home ■ Browser (GMT-00.00) ▼ Status Up Time 34 minutes, 43 reconds Up Time 54 minutes, 43 reconds Up Time 54 minutes, 43 reconds Up Time 94 minutes, 43 reconds Up Time 54 reconds Version 19.18.0.00 Enterprise Edition Platom Name Linux56 Febitit Timed 1 Archiver Stared Last Backup Time N/A Incident(s) ● 1 Resources ■	Activity Services Containers 40			Cther Netvork System (/O User (/O C CD
■ NAP1 (1918.0.00) Performance ▼ Storage ▼ Database Home Time Zone Browser (GMT-00.00) ▼ Status Up Time 34 minutes. 43 seconds Up Time 34 minutes. 43 seconds Type Single instance (NTAP1) CD8 (3 POB(d)) Version 19.18.0.0.0 Enterprise faition Platform Kart Platform Kart Last Backup Time N/A Last Backup Time N/A Incident(s) ● 1 Time 1% 60 1%	Activity Services Containers 40	14 GB	3.7 G8	Cther Netvork System (/O User (/O C CD
■ NAP1 (1918.0.0.0) Performance ▼ Storage ▼ Database Home Time Zone Browser (GMT-00.00) ▼ Status Up Time 34 minutes. 43 reconds Up Time 54 minutes. 43 reconds Up Time 54 minutes. 43 reconds Up Time 1918.0.00 Enterprise Edition Platom Name Linux 56 de-bit Thread 1 Archiver Started Last Backup Time N/A Incident(s) ● 1 The Status \$0 1 Status 0 0	Activity Services Containers 40		total_sga total_sga total_sga	Other Network System I/O User I/O Cold2200 PM (0447360 PM
■ NAP1 (19:180.00) Performance ▼ Storage ▼ Database Home Time Zone Browser (GMT-00:00) ▼ Status Up Time 34 minutes, 43 seconds Up Time 34 minutes, 43 seconds Up Time 34 minutes, 43 records Up Time 34 minutes, 43 records Up Time 34 minutes, 43 records Version 19:18:0.00 Enterprise Edition Platform Name Linux 86 6-bit Time 41 Archiver Started Last Backup Time N/A Incident(s) ● 1 Tis 50 0% ● Detervition	Activity Services Containers 40	14.68	1.7.68 total_sga total_gga target_gga shared pool 1.9.68	Other Network System I/O User I/O CFU Od4220 PM (044740 PM
NAP1 (19180.00) Performance ▼ Storage ▼ Database Home Time Zone Browser (GMT-0000) ▼ Status Up Time 34 minutes. 43 seconds Type Single Instance (VTAP1) CD8 (§ PDB(s)) Version 1918.00.00 Enterprise Edition Platform Name Inax 265 64-bit Thread 1 Archiver Started Last Backup Time N/A Incident(s) ● 1	Activity Services Containers	14 (56 112 (58 8.4 (56	3.7 GB total_ga 2.8 GB tatagt_pa 1.9 GB staret pool staret pool	Other Network System Od44220 PM 0647A0 PM
■ NAP1 (1918.0.0.0) Performance * Storage * Database Home Time Zone Browser (GMT-00:00) * Status Up Time 34 minutes. 43 seconds Type Single Instance (VTAP1) CDB (§ PDB(s)) Version 1918.0.0.0 Enterprise Edition Platform Mane Linux 365 4-bit Thread 1 Archiver Started Last Backup Time N/A Incident(s) 1% 0% 0% 0% 0% 0% 0% 0% 0% 0% 0%	Activity Services Containers 40	14.66 112.08 8.4.08 5.6.68 2.8.66 08	1.7.68 total_sga total_gga target_gga shared pool 1.9.68	Other National States of the states
■ NAP1 (1918.0.0.0) Performance * Storage * Database Home Time Zone Browser (GMT-00:00) * Status Up Time 34 minutes. 43 seconds Type Single Instance (VTAP1) CDB (§ PDB(s)) Version 1918.0.0.0 Enterprise Edition Platform Mane Linux 365 4-bit Thread 1 Archiver Started Last Backup Time N/A Incident(s) 1% 0% 0% 0% 0% 0% 0% 0% 0% 0% 0%	Activity Services Containers 40	14 08 112 68 8.4 66 5.6 68 2.8 68	1.7 GB total_sga total_gga target_gga sharef pol large pool 953.7 MB	Other Other Network System VO User VO User VO User VO Other Other Network System VO Other Network System VO Other Network System VO Other Notwork System VO Other Notwork System VO Other Notwork System VO Other Notwork System VO Other Notwork System VO Other Notwork System VO Other Notwork System VO Other Notwork System VO Other Notwork System VO Other Notwork Notwork System VO Other Notwork System VO Other Notwork System VO Other Notwork Notwork Notwork System VO Other Notwork
NAP1 (19180.00) Performance ▼ Storage ▼ Database Home Time Zone Browser (GMT-0000) ▼ Status Up Time 34 minutes.43 seconds Type Single Instance (NTAP1) COB (§ PDB(s)) Werion 1918.00.00 Enterprise Edition Platform Mane (INXAP) CDB (§ PDB(s)) Werion 1918.00.00 Enterprise Edition Platform Mane (INXAP) Thread 1 Archiver Started Last Backup Time N/A Incident(s) 1% 0%	Activity Services Containers 40	14.66 112.08 8.4.08 5.6.68 2.8.66 08	1.7 GB total_sga total_gga target_gga sharef pol large pool 953.7 MB	Other National States of the states
■ NAP1 (1918.0.0.0) Performance ▼ Storage ▼ Database Home Time 20ne Browser (GMT-00.00) ▼ Status Up Time 34 minutes. 43 seconds Up Time 34 minutes. 43 seconds Type Single Instance (NTAP1) CDG (3 POB(0)) Version 19.18.0.0.0 Enterprise Edition Platom back the NA Last Backup Time NA Incident(s) ● 1 Resources 1% 0% 0% 0% Heat CPU SQL Monitor - Last Hour (20 max)	Activity Services Containers 40	14 08 112 08 84 08 36 08 2.8 08 0 8 Memory	3.7 GB total_pgs 2.8 GB total_pgs 2.8 GB target_pgs 1.9 GB there pool large pool 953.7 MB Shared IO P 0.8	Other National States of the states
NAP1 (19180.00) Performance ▼ Storage ▼ Database Home Time Zone Browser (GMT-0000) ▼ Status Up Time 34 minutes.43 seconds Type Single Instance (NTAP1) COB (§ PDB(s)) Werion 1918.00.00 Enterprise Edition Platform Mane (INXAP) CDB (§ PDB(s)) Werion 1918.00.00 Enterprise Edition Platform Mane (INXAP) Thread 1 Archiver Started Last Backup Time N/A Incident(s) 1% 0%	Activity Services Containers 40	14 08 112 08 84 08 36 08 2.8 08 0 8 Memory	1.7 GB total_sga total_gga target_gga sharef pol large pool 953.7 MB	Other Other Other System I/O User I/O Odd220 PM 064740 PM
NAP1 (19180.00) Performance ▼ Storage ▼ Database Home Time Zone Browser (GMT-0000) ▼ Status Up Time 34 minutes.43 seconds Type Single Instance (NTAP1) COB (§ PDB(s)) Version 1918.00.00 Enterprise Edition Platform Mare (INXPF) COB (§ PDB(s)) Version 1918.00.00 Enterprise Edition Platform Mare (INXPF) COB (§ PDB(s)) Version 1918.00.00 Enterprise Edition Platform Mare (INXPF) Last Backup Time N/A Incident(s) 1% 0% 0% 0% 0% 0% 0% 10 0% 10 0% <td>Activity Services Containers 40 </td> <td>14 08 112 08 84 08 36 08 2.8 08 0 8 Memory</td> <td>3.7 GB total_pgs 2.8 GB total_pgs 2.8 GB target_pgs 1.9 GB there pool large pool 953.7 MB Shared IO P 0.8</td> <td>Other National States of the states</td>	Activity Services Containers 40	14 08 112 08 84 08 36 08 2.8 08 0 8 Memory	3.7 GB total_pgs 2.8 GB total_pgs 2.8 GB target_pgs 1.9 GB there pool large pool 953.7 MB Shared IO P 0.8	Other National States of the states

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-aservice 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

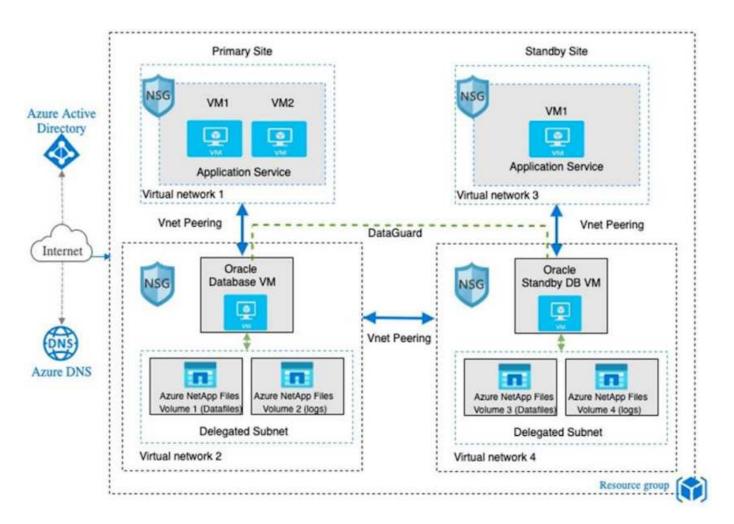
This documentation describes in detail how to deploy, configure, and protect an Oracle database with an Azure virtual machine and Azure NetApp Files storage service that delivers performance and durability similar to an on-premises system. For best-practices guidance, see TR-4780 Oracle Databases on 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.

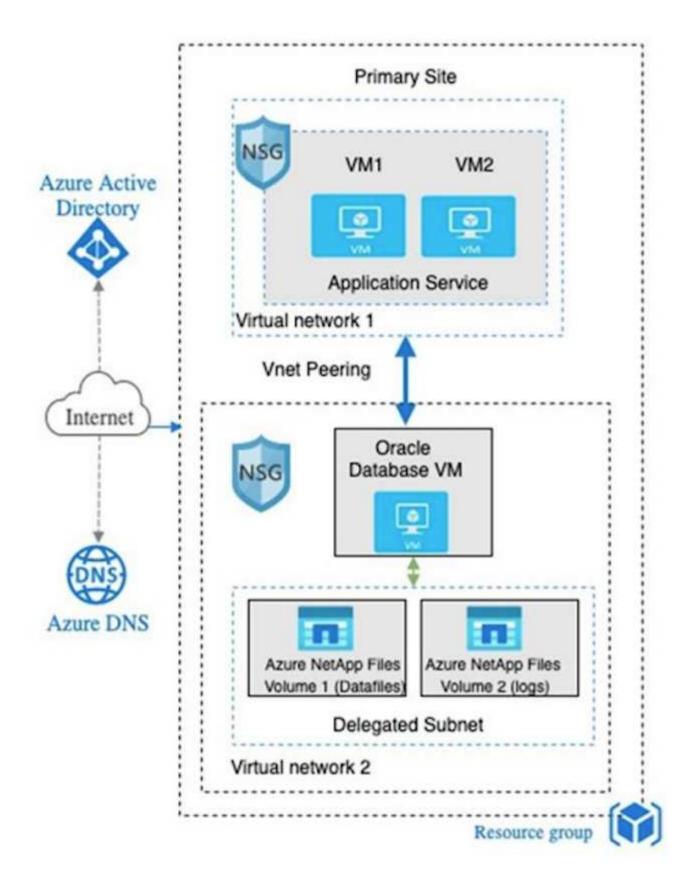
Within the environment, the Oracle compute instance is deployed via an Azure services VM console. There are multiple Azure instance types available from the console. NetApp recommends deploying a database-oriented Azure VM instance that meets your expected workload.

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.

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.

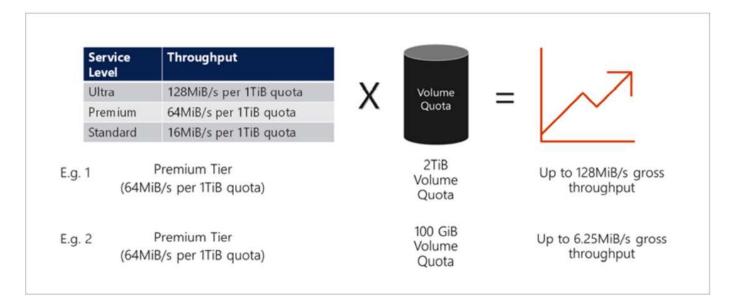
For example, a 10TiB-provisioned single-capacity pool with the premium service level provides an overall available throughput for all volumes in this capacity pool of 10x 64MBps, so 640MBps with 40,000 (16K) IOPs or 80,000 (8K) IOPs.

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.

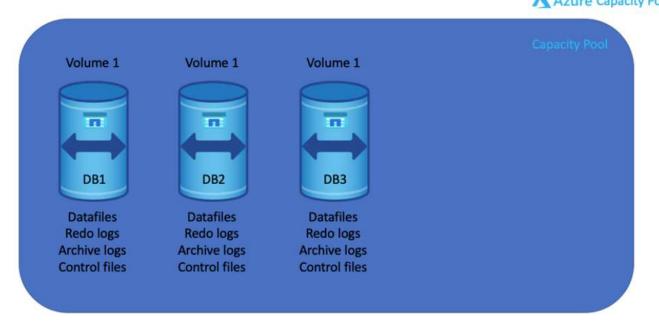
Storage layout and settings

NetApp recommends the following storage layout:

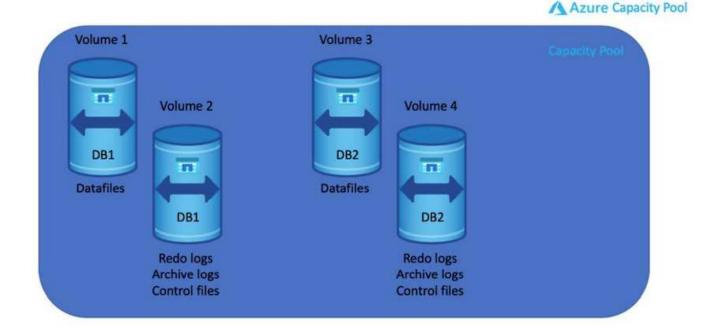
• For small databases, using single volume layout for all Oracle files.



Microsoft Azure



• 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.

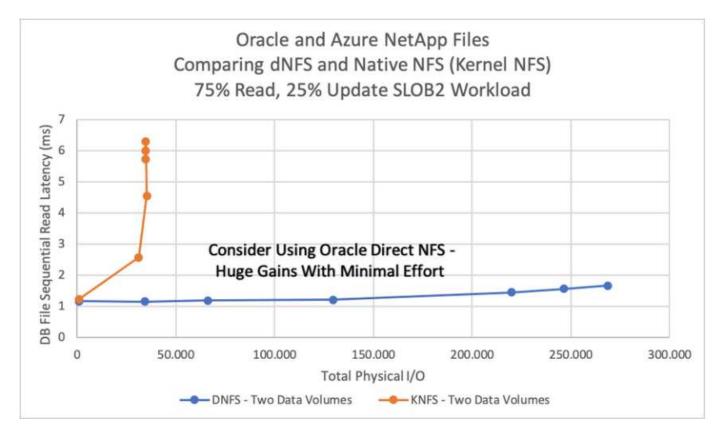


NFS configuration

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
Control filesData filesRedo logs	<pre>rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsize=6 5536</pre>
ORACLE_HOME ORACLE_BASE	<pre>rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsize=6 5536</pre>

(i)

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

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

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

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

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

Create and consume Azure resources

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

1. Create a virtual network or VNet

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

See Create a virtual network using the Azure portal to create a VNet.

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

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

See Set up Azure NetApp Files and create an NFS volume to set up an ANF capacity pool.

3. Provision Azure VM for Oracle

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

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

Microsoft Azure			rvices, and docs (G+/)				l	E 6 4 © 0	R acao@netap HYBRID CL	рр.com 🧧
lome >										
/irtual machines 🛷 … ybrid Cloud TME										×
$+$ Create \lor \rightleftharpoons Switch to classic 🕚 Reserv	vations 🗸 🛞 Manage view 🗸	🖒 Refresh 🞍 Export to CSV	/ 😤 Open query 🕴 🖉	Assign tags ▷ Start	🤆 Restart 🔲 Stop 📋 D	Delete 📧 Services 🗸 🤌	Maintenance \checkmark			
Filter for any field Subscription equa	is all Type equals all	Resource group equals all $ imes$	Location equals all $ imes$	⁺ _▼ Add filter						
								No grouping	∽ I≡ List view	\sim
Name 🗘	Type ↑↓	Subscription \uparrow_{\downarrow}	Resource group \uparrow_{\downarrow}	Location \uparrow_{\downarrow}	Status ↑↓	Operating system \uparrow_{\downarrow}	Size ↑↓	Public IP address \uparrow_{\downarrow}	Disks ↑↓	
📮 💶 acao-ora01	Virtual machine	Hybrid Cloud TME Onprem	TMEtstres	South Central US	Stopped (deallocated)	Linux	Standard_B4ms	13.65.63.157	1	
ANFAVFVal2JH	Virtual machine	Hybrid Cloud TME Onprem	ANFAVSVAL2	West Europe	Running	Windows	Standard_DS2_v2	20.229.80.88	1	-
ANFAVSfio01	Virtual machine	Hybrid Cloud TME Onprem	anfavsrg	South Central US	Stopped (deallocated)	Linux	Standard_D32ds_v4		1	
ANFAVSfioAZ1	Virtual machine	Hybrid Cloud TME Onprem	anfavsrg	South Central US	Running	Linux	Standard_E32as_v4	40.124.74.246	1	
ANFAVSfioAZ2	Virtual machine	Hybrid Cloud TME Onprem	anfavsrg	South Central US	Stopped (deallocated)	Linux	Standard_E32as_v4	40.124.178.111	1	
ANFAVSfioAZ3	Virtual machine	Hybrid Cloud TME Onprem	anfavsrg	South Central US	Stopped (deallocated)	Linux	Standard_E32as_v4	40.124.194.32	1	
ANFAVSvalDC	Virtual machine	Hybrid Cloud TME Onprem	anfavsrg	South Central US	Stopped (deallocated)	Windows	Standard_B4ms	-	1	
ANFAVSvalJH	Virtual machine	Hybrid Cloud TME Onprem	anfavsrg	South Central US	Running	Windows	Standard_B2ms	70.37.66.218	1	-
ANFAVSvalJH2	Virtual machine	Hybrid Cloud TME Onprem	anfavsrg	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	
ANFCVODRDemo	Virtual machine	Hybrid Cloud TME Onprem	anfcvodrdemo-rg	West Europe	Running	Linux	Standard_E4s_v3		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.

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 C^a

Project details

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

Subscription * 🕡	Hybrid Cloud TME Onprem	\sim
Resource group * 🛈	ANFAVSRG	\sim
	Create new	
Instance details		
Virtual machine name * 🕕	acao-ora01	~
Region * 🛈	(US) South Central US	\sim
Availability options	No infrastructure redundancy required	\sim
Security type ①	Standard	\sim
Image * 🛈	🌨 Red Hat Enterprise Linux 8.0 (LVM) - Gen2	\sim
	See all images Configure VM generation	
Run with Azure Spot discount ①		
Size * 🛈	Standard_D8s_v3 - 8 vcpus, 32 GiB memory (\$273.02/month)	\sim
	See all sizes	
Administrator account		
Authentication type 🛈	SSH public key	
	Password	
Review + create < Prev	vious Next : Disks >	

■ Microsoft Azure		$\mathcal P$ Search resources, services, and docs
Home > Virtual machines >		
Create a virtual mach	nine	
Size * 🕕	Standard_D8s_v3 - 8 vcpus, 32 GiB memory (\$27 See all sizes	(3.02/month)
Administrator account Authentication type ①	 SSH public key Password 	
Username * 🛈	azureuser	\checkmark
Password *	•••••	\checkmark
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 * ①	None Allow selected ports	
Select inbound ports *	SSH (22)	\sim
	A 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.

s (G+/

■ Microsoft Azure	$\mathcal P$ Search resources, services, and docs (G+/)
Home > Virtual machines >	
Create a virtual machine	
Basics Disks Networking Management Advanced Tags Review + c	reate
Azure VMs have one operating system disk and a temporary disk for short-term storage. You The size of the VM determines the type of storage you can use and the number of data disks	
Disk options	
OS disk type * () Premium SSD (locally-redundant storage)	\checkmark
Delete with VM 🕦 🔽	
Enable encryption at host 🛈	
Fincryption at host is not registered for the selected subscription. Learn more about enabling	this feature 🖉
Encryption type * (Default) Encryption at-rest with a platform-mar	naged key 🗸 🗸
Enable Ultra Disk compatibility ①	
Data disks for acao-ora01	
You can add and configure additional data disks for your virtual machine or attach existing di temporary disk.	sks. This VM also comes with a
LUN Name Size (GiB) Disk type Host cachin	ng 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.

Create a virtual machine

Network interface

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

Virtual network * 🕕	ANFAVSVal	\sim
	Create new	
Subnet * 🕡	VM_Sub (172.30.137.128/25)	\sim
	Manage subnet configuration	
Public IP 🕕	(new) acao-ora01-ip	\sim
	Create new	
NIC network security group 🛈	O None	
	Basic	
	O Advanced	
Public inbound ports * 🕕	O None	
	Allow selected ports	
Select inbound ports *	SSH (22)	\sim
	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 ①		
Enable accelerated networking ①		
Load balancing		
You can place this virtual machine in the ba	ackend pool of an existing Azure load balancing solution. Learn more 🗷	
Place this virtual machine behind an existing load balancing solution?		
Review + create < Previ	ous Next : Management >	

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

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
	O Disable
Enable OS guest diagnostics ①	
Identity	
Enable system assigned managed identity ()	
Azure AD	
Login with Azure AD ①	
	RBAC role assignment of Virtual Machine Administrator Login or Virtual Machine User Login is required when using Azure AD login. <u>Learn more</u> [™]
	based authentication. You will need to use an SSH client that supports OpenSSH id Shell from the Azure Portal. Learn more $arepsilon^{*}$
Auto-shutdown	
Enable auto-shutdown ①	
Backup	
Review + create < Prev	ious 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.

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 (i)

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

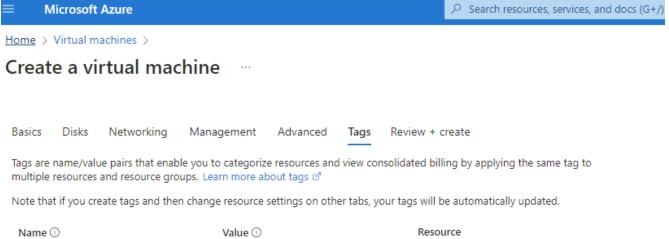
〕 Your image must have a code to support consumption of custom data. If your image suppo	orts cloud-init, custom-data will be
processed by cloud-init. Learn more about custom data for VMs and	

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 C⁴

Enable user data	
Review + create	< Previous Next : Tags >

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



database	:	oracle	12 selected V	Ì
	:		12 selected V	

Review + create	< Previous	Next : Review + create >	
-----------------	------------	--------------------------	--

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

4. Provision ANF database volumes for Oracle

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

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

≡ Microsoft Azure	₽ Sea	rch resources, service	s, and docs (G+/)						
	Azure services								
	+ •	«· >	•		2	*	٩	×	\rightarrow
	Create a Azure NetAp resource Files	p Virtual networks	Virtual machines	Storage accounts	Users	Subscriptions	Azure Active Directory	Quickstart Center	More services
	Resources								
	Recent Favorite								
	Name			Туре				Last Viewed	
	ANFAVSAcct			NetApp	account			a few seconds	ago
	↔ ANFAVSVal			Virtual n	etwork			3 hours ago	
	롲 acao-ora01			Virtual n	nachine			5 days ago	
	📍 Hybrid Cloud TME Onprem			Subscrip	tion			2 weeks ago	
	WEANFAVSacct			NetApp	account			2 weeks ago	
	📕 ANFAVSAcct/CapPool/acao	-ora01-u03		Volume				2 weeks ago	
	📕 ANFAVSAcct/CapPool/acao	-ora01-u02		Volume				2 weeks ago	
	ANFAVSAcct/CapPool/acao	-ora01-u01		Volume				2 weeks ago	
	acao-ora01_OsDisk_1_673b	ad70ccce4709afc8127	78e2bc97cb	Disk				2 weeks ago	
	🚮 acao-ora0166			Network	Interface			3 weeks ago	
	[iii] TMEtstres			Resource	group			3 weeks ago	
	See all								

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

E Microsoft Azure			,₽ se	earch resources, services	, and docs (G+/)				$\mathbf{\Sigma}$	Ģ	Q (
Home > Azure NetApp Files > Azure NetApp Files Hybrid Cloud TME + Create @ Manage view > ····	«	ANFAVSAcct NetApp account Search (Ctrl+/)	\$	💼 Delete							
Filter for any field Name †: ANFAVSAcct WEANFAVSacct		A Trac		Essentials Resource group (move Location Subscription (move) Subscription ID Tags (edit)) : ANEAVSRG : South Central US : Hybrid Cloud TME Onprem : Oefa2dfb-917c-4497-b56a-b3f4ead : product_line : Field use - various	28111	Provisionin	g state : Succeeded			
		III Properties △ Locks Azure NetApp Files ◆ Active Directory connections Storage service ☜ Capacity pools ☜ Volumes Data protection ☑ Snapshot policies Storage service add-ons Image service add-ons			Connect to Active Directory Connect your NetApp to Active Learn more B View AD connections	Azure NetApp Fi migrate a	nterprise files storage, powered les makes it easy for enterprise line-of-business ind run complex, file-based applications with no of Capacity pools Purchase pools of capacity with a service le which you provision volumes. Learn more of View capacity pools	(LOB) and storage professi code change. Learn more p Volumes vel in Container for	ø active file d snapshe		
< Page 1 v of 1 >		Automation Automation Automatic (preview) Support template Support + troubleshooting R New Support Request									

Microsoft Azure		P Search resources, services,	and docs (G+/)				🗆 🛛 🖓 🖉 🎯 (ତ ନ [ା] acac	o@netapp.co
ome > Azure NetApp Files > ANFAVSAcct								-	
verine NetApp Files «	ANFAVSAcct Volu	imes							
🕂 Create 🕘 Manage view 🗸 \cdots		+ Add volume +	Add data replication	🕐 Refresh					
ilter for any field	Overview	₽ Search volumes							
ame 1	Activity log	Name	î↓ Quota	↑↓ Throughput	↑↓ Protocol type	↑↓ Mount path ↑↓ Service level	↑↓ Network features ↑↓	Capacity pool	¢ψ
	Access control (IAM)	anf2+z1-stdds01	200 GiB	25.6 MiB/s	NFSv3	172.30.136.70;/anf2-z1-std Ultra	Standard	cappool	
ANTAVSACCI	🗳 Tags	anf2-z1-stdds02	200 GiB	25.6 MiB/s	NFSv3	172.30.136.70:/anf2-z1-std Ultra	Standard	cappool	
WEANFAVSacct	Settings	anf2-z1-stdds03	100 GiB	12.8 MiB/s	NFSv3	172.30.136.70:/anf2-z1-std Ultra	Standard	cappool	
		anf2-z1-stdds04	100 GiB	12.8 MiB/s	NFSv3	172.30.136.70:/anf2-z1-std Ultra	Standard	cappool	
	Quota	anf2-z1-stdds05	100 GiB	12.8 MiB/s	NFSv3	172.30.136.70:/anf2-z1-std Ultra	Standard	cappool	
	Properties	anf2-z1-stdds06	100 GiB	12.8 MiB/s	NFSv3	172.30.136.70:/anf2-z1-std Ultra	Standard	cappool	
	A Locks	anf2-z1-stdds07	100 GiB	12.8 MiB/s	NFSv3	172.30.136.70:/anf2-z1-std Ultra	Standard	cappool	
	Azure NetApp Files	anf2-z1-stdds08	100 GiB	12.8 MiB/s	NFSv3	172.30.136.70:/anf2-z1-std Ultra	Standard	cappool	
	Active Directory connections	anf-z1-stdds01	6 TiB	786.432 MiB/s	NFSv3	172.30.136.70:/anf-z1-stdc Ultra	Standard	cappool	
	Storage service	anf-z1-stdds02	200 GiB	25.6 MiB/s	NFSv3	172.30.136.70:/anf-z1-stdc Ultra	Standard	cappool	

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

Microsoft Azure		م	Search resources, services, and docs (G+/)	
Home > Azure NetApp Files > ANFAVSA	Acct Volumes >			
RetApp account	nes	«	Create a volume	
	+ Add volume ····			
Overview	${\cal P}$ Search volumes	•	Basics Protocol Tags Review	w + create
Activity log	Name ↑↓ Quota	- 1	This page will help you create an Azure within your virtual network. Learn more	NetApp Files volume in your subscription and enable you to access the volume from about Azure NetApp Files rd
Access control (IAM)	록 anf2-z1-stdds01 200 GiB		Volume details	
🗳 Tags	록 anf2-z1-stdds02 200 GiB		Volume name *	acao-ora01_u01 🗸
Settings	록 anf2-z1-stdds03 100 GiB		Capacity pool * 🕡	CapPool V
 []] Quota	anf2-z1-stdds04 100 GiB			
Properties	anf2-z1-stdds05 100 GiB		Available quota (GiB) 🕕	572 572 GiB
A Locks	anf2-z1-stdds06 100 GiB		Quota (GiB) * 🕕	100 ~
—	anf2-z1-stdds07 100 GiB		Quota (Gib) ····································	100 GiB
Azure NetApp Files	록 anf2-z1-stdds08 100 GiB		Available throughput (MiB/s) 🕕	73.22
Active Directory connections	록 anf-z1-stdds01 6 TiB		Throughout (1450 (c)	12.5
Storage service	anf-z1-stdds02 200 GiB		Throughput (MiB/s) 🛈	12.3
🚡 Capacity pools	anf-z1-stdds03 1 TiB		Enable Cool Access 🛈	
Volumes	록 anf-z1-stdds04 200 GiB		Coolness Period 🛈	31
Data protection	록 anf-z1-stdds06 200 GiB	- 1	Virtual network * 🕕	ANFAVSVal (172.30.136.64/26,172.30.137.128/25,172.30.152.0/27)
Snapshot policies	anf-z1-stdds07 200 GiB			Create new virtual network
Shapshot policies	anf-z1-stdds08 200 GiB	- 1	Delegated subnet * 🕕	ANF_Sub (172.30.136.64/26)
Storage service add-ons	anf-zq-stdds05 1 TiB	- 1		Create new subnet
NetApp add-ons	vol1 1 TiB	- 1	Network features	🔘 Basic 💿 Standard
Automation	vol3basic 100 GiB	- 1	Availability Zone 🕕	None V
🖧 Tasks (preview)	volnfsbasic 100 GiB	- 1	Show advanced section	
😫 Export template	volnfsstd 100 GiB	- 1		—
Support + troubleshooting	volnfsstdnew 100 GiB	- 1		
Rew Support Request	zone1basic 6 TiB			
//	400 CiD	- 1	Review + create < P	revious Next : Protocol >

4. 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		𝒫 Search resources, services, and docs (G+/)
Home > Azure NetApp Files > ANFAVSA	cct Volumes >	
RetApp account	es «	« Create a volume …
	+ Add volume ····	
🔟 Overview	$\mathcal P$ Search volumes	Basics Protocol Tags Review + create
Activity log	Name ↑↓ Quota	Configure access to your volume.
Access control (IAM)	anf2-z1-stdds01 200 GiB	Access
Tags	anf2-z1-stdds02 200 GiB	Protocol type
Settings	anf2-z1-stdds03 100 GiB	Configuration
D Quota	anf2-z1-stdds04 100 GiB	File path * () acao-ora01_u01
Properties	anf2-z1-stdds05 100 GiB	
A Locks	anf2-z1-stdds06 100 GiB	Versions *
Azure NetApp Files	anf2-z1-stdds07 100 GiB	Kerberos O Enabled O Disabled
	anf2-z1-stdds08 100 GiB	LDAP O Enabled O Disabled
Active Directory connections	anf-z1-stdds01 6 TiB	Azure VMware Solution DataStore 💿 🗌
Storage service	anf-z1-stdds02 200 GiB	
🐁 Capacity pools	anf-z1-stdds03 1 TiB	Export policy
Volumes	anf-z1-stdds04 200 GiB	Configure the volume's export policy. This can be edited later. Learn more 🖻
Data protection	anf-z1-stdds06 200 GiB	\uparrow Move up \downarrow Move down $ar{\uparrow}$ Move to top \downarrow Move to bottom $ar{\blacksquare}$ Delete
Snapshot policies	anf-z1-stdds07 200 GiB	Index Allowed clients Access Root Access
Storage service add-ons	anf-z1-stdds08 200 GiB	1
NetApp add-ons	anf-zq-stdds05 1 TiB	□ 2 172.30.137.142 ✓ Read & Write ✓ On ✓ ····
	Vol1 1 TiB	
Automation	vol3basic 100 GiB	
Tasks (preview)	volnfsbasic 100 GiB	
Export template	volnfsstd 100 GiB	
Support + troubleshooting	volnfsstdnew 100 GiB	
📯 New Support Request	zone1basic 6 TiB	Review + create < Previous Next : Tags >
	zone2basic 100 GiB	▼

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

Home > Azure NetApp Files > ANFAVS	Acct Volumes >		
RetApp account	nes	« Create a volume …	
✓ Search (Ctrl+/) «	+ Add volume ····		
🔟 Overview	${\cal P}$ Search volumes	Basics Protocol Tags Review + create	
Activity log	Name ↑↓ Quota	Tags are name/value pairs that enable you to categorize resources and view consolidated billing by applying the san multiple resources and resource groups. Learn more about tags 🕫	me tag to
Access control (IAM)	anf2-z1-stdds01 200 GiB	Note that if you create tags and then change resource settings on other tabs, your tags will be automatically update	ed.
🗳 Tags	anf2-z1-stdds02 200 GiB		
Settings	anf2-z1-stdds03 100 GiB	Name O Value O	_ ~
 Quota	anf2-z1-stdds04 🛛 100 GiB	database : oracle	Î.
Properties	anf2-z1-stdds05 100 GiB		
Locks	anf2-z1-stdds06 100 GiB		
	anf2-z1-stdds07 100 GiB		
Azure NetApp Files	anf2-z1-stdds08 100 GiB		
Active Directory connections	anf-z1-stdds01 6 TiB		
Storage service	anf-z1-stdds02 200 GiB		
🐁 Capacity pools	록 anf-z1-stdds03 1 TiB		
Volumes	anf-z1-stdds04 200 GiB		
Data protection	anf-z1-stdds06 200 GiB		
Snapshot policies	anf-z1-stdds07 200 GiB		
	anf-z1-stdds08 200 GiB		
Storage service add-ons	anf-zq-stdds05 1 TiB		
NetApp add-ons	록 vol1 1 TiB		
Automation	Nol3basic 100 GiB		
Tasks (preview)	Not volnfsbasic 100 GiB		
😫 Export template	Not volnfsstd 100 GiB	_	
Support + troubleshooting	volnfsstdnew 100 GiB		
Request	록 zone1basic 6 TiB	Review + create < Previous Next : Review + create >	
	zone2basic 100 GiB		

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

Microsoft Azure		𝒫 Search resources, services, and docs (G+/)
Home > Azure NetApp Files > ANFAV	SAcct Volumes >	
ANFAVSAcct Volui	mes	« Create a volume …
	+ Add volume	✓ Validation passed
🖬 Overview	$\mathcal P$ Search volumes	
Activity log	Name ↑↓ Quota	Basics Protocol Tags Review + create
Access control (IAM)	록 anf2-z1-stdds01 200 GiB	
🗳 Tags	anf2-z1-stdds02 200 GiB	Basics
Settings	록 anf2-z1-stdds03 🛛 100 GiB	Subscription Hybrid Cloud TME Onprem
Quota	anf2-z1-stdds04 100 GiB	Resource group ANFAVSRG Region South Central US
—	anf2-z1-stdds05 100 GiB	Region South Central US Volume name acao-ora01-u01
Properties	anf2-z1-stdds06 100 GiB	Capacity pool CapPool
🔒 Locks	anf2-z1-stdds07 100 GiB	Service level Ultra
Azure NetApp Files	anf2-z1-stdds08 100 GiB	Quota 100 GiB Encryption key source Microsoft.NetApp
Active Directory connections	anf-z1-stdds01 6 TiB	Availability Zone None
Storage service	록 anf-z1-stdds02 200 GiB	
🚡 Capacity pools	anf-z1-stdds03 1 TiB	Networking
Volumes	anf-z1-stdds04 200 GiB	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)
Data protection	anf-z1-stdds06 200 GiB	Network features Standard
-	anf-z1-stdds07 200 GiB	-
Snapshot policies	anf-z1-stdds08 200 GiB	Protocol
Storage service add-ons	anf-zq-stdds05 1 TiB	Protocol NFSv3
NetApp add-ons	vol1 1 TiB	File path acao-ora01-u01
Automation	vol3basic 100 GiB	Tags
🔓 Tasks (preview)	volnfsbasic 100 GiB	database oracle
😫 Export template	volnfsstd 100 GiB	-
Support + troubleshooting	volnfsstdnew 100 GiB	-
Rew Support Request	zone1basic 6 TiB	
rivew support request	zone2basic 100 GiB	Create < Previous Next > Download a template for automation

Install and configure Oracle on Azure VM with ANF

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

Set up an Ansible controller

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

Obtain Oracle deployment automation toolkit

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

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

Execute the toolkit with your configuration

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

the Azure console rather than the CLI.



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

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

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

Install AzAcSnap tool

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

./azacsnap installer v5.0.run -I

Configure Oracle connectivity

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

1. Set up AzAcSnap database user

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

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

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

2. Create the user.

CREATE USER azacsnap IDENTIFIED BY password;

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

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

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

ALTER PROFILE default LIMIT PASSWORD LIFE TIME unlimited;

5. Validate azacsnap connectivity for the database.

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

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

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

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

cat /etc/oratab

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

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

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

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

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

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

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

```
echo "# Connection string
AZACSNAP=\"(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=172.30.137.142)(POR
T=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 <code>azacsnap.json</code> configuration file, add all data volumes, including the binary volume, to <code>dataVolume</code> and all log volumes to <code>otherVolume</code>. 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.

Microsoft Azure		${\cal P}$. Search resources, services, and docs (G+/)				acao@netapp.com нувяю столо тме
ome > Azure NetApp Files > ANFA	VSAcct Volumes > acao-ora01-u01 (ANFAVSAc					
ANFAVSAcct Volu	imes «	Colume acao-ora01-u01 (AN	NFAVSAcct/CapPool	l/acao-ora01-u01) Snapshots		
© Search (Ctrl+/) «	+ Add volume ····		+ Add snapshot 💍 Refres	h		
Overview		Overview	Search snapshots			
Activity log	Name ↑↓ Quota	Activity log	Name	↑↓ Location	↑↓ Created	¢ψ.
Access control (IAM)	acao-ora01-u01 100 GiB	Access control (IAM)	(D) acao-ora01-data_2022-	09-09T165255-8258852Z South Central US	09/09/2022, 12:53:22 PM	
Tags	acao-ora01-u02 100 GiB	Tags	() acao-ora01-data_2022-	09-12T160536-9809839Z South Central US	09/12/2022; 12:05:55 PM	
ettings	acao-ora01-u03 100 GiB	Settings				
Quota	anf2-z1-stdds01 200 GiB	III Properties				
Properties	🖷 anf2-z1-stdds02 200 GiB	🔒 Locks				
Locks	anf2-z1-stdds03 100 GiB	Storage service				
zure NetApp Files	anf2-z1-stdds04 100 GiB	Mount instructions				
Active Directory connections	anf2-z1-stdds05 100 GiB	Export policy				
orage service	anf2-z1-stdds06 100 GiB	(C) Snapshots				
Capacity pools	anf2+z1+stdds07 100 GiB	Replication				
Volumes	anf2-z1-stdds08 100 GiB	Monitoring				
	anf-z1-stdds01 6 TiB					

Microsoft Azure		𝒫 Search resources, services, and docs (G+/	1			acao@netapp.com HYBRID CLOUD TME
ANFAVSAcct Volu		() acao-ora01-u03		ol/acao-ora01-u03) Snapshots		×
P Search (Ctrl+/) «	+ Add volume ····	Overview	≪ + Add snapshot ◯ Refre	sh		
Activity log	Name ↑↓ Quota	Activity log	Name	↑↓ Location	↑↓ Created	¢ψ
Access control (IAM)	acao-ora01-u01 100 GiB	Access control (IAM)	(b) acao-ora01-data_2022	-09-12T160628+8754798Z South Central US	09/12/2022, 12:06:31 PM	
Tags	acao-ora01-u02 100 GiB	Tags	(L) acao-ora01-log_2022-	09-12T164501-7965925Z South Central US	09/12/2022, 12:45:04 PM	
ettings	acao-ora01-u03 100 GiB	Settings	(U) acao-ora01-log_2022-	09-12T171501-8032661Z South Central US	09/12/2022, 01:15:04 PM	
Quota	anf2-z1-stdds01 200 GiB	Properties	(L) acao-ora01-log_2022-	09-12T173001-4787919Z South Central US	09/12/2022, 01:30:04 PM	
Properties	anf2-z1-stdds02 200 GiB	🔒 Locks	(L) acao-ora01-log_2022-	09-12T174501-5010614Z South Central US	09/12/2022, 01:45:04 PM	
Locks	anf2-z1-stdds03 100 GiB	Storage service	(U) acao-ora01-log_2022-	09-12T180053-5029874Z South Central US	09/12/2022, 02:00:55 PM	
zure NetApp Files	anf2-z1-stdds04 100 GiB	Mount instructions	(E) acao-ora01-log_2022-	09-12T181502-3857027Z South Central US	09/12/2022, 02:15:05 PM	
Active Directory connections	anf2-z1-stdds05 100 GiB	Export policy	(b) acao-ora01-log_2022-	09-12T183002-4407395Z South Central US	09/12/2022, 02:30:07 PM	
orage service	anf2-z1-stdds06 100 GiB	(O) Snapshots				
Capacity pools	anf2-z1-stdds07 100 GiB	Replication				
Volumes	anf2-z1-stdds08 100 GiB	Monitoring				
ata protection	anf-z1-stdds01 6 TiB	ai Metrics				

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;
ΙD
_____
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 - Production
Version 19.8.0.0.0
```

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

Microsoft Azure		𝒫 Search resources, services, and docs (G+/)				ြာ acao@netapp.com нувкю ссоир тм	
Home > Azure NetApp Files > ANFA	WSAcct Volumes > acao-ora01-u03 (ANFAVSA	cct/CapPool/acao-ora01-u03)					
ANFAVSAcct Volu	umes	acao-ora01-u03	(ANFAVSAcct/CapPoo	l/acao-ora01-u03) Snapshots			
	+ Add volume ····	Search (Ctrl+/)	« 🕂 Add snapshot 💍 Refre	sh			
Overview	♪	Cverview					
Activity log	Name ↑↓ Quota	Activity log	Name	↑↓ Location	↑↓ Created	↑↓	
Access control (IAM)	acao-ora01-u01 100 GiB	Access control (IAM)	(I) acao-ora01-data_2022-	09-12T160628-8754798Z South Central US	09/12/2022, 12:06:31 PM		
Tags	acao-ora01-u02 100 GiB	Tags	(b) acao-ora01-log_2022-0	9-13T120122-8173645Z South Central US	09/13/2022, 08:01:25 AM		
Settings	acao-ora01-u03 100 GiB	Settings	(b) acao-ora01-log_2022-0	9-13T121501-6760417Z South Central US	09/13/2022, 08:15:04 AM		
Quota	anf2-z1-stdds01 200 GiB	Properties	(b) acao-ora01-log_2022-0	9-13T123001-8000548Z South Central US	09/13/2022, 08:30:05 AM		
Properties	anf2+z1+stdds02 200 GiB	🔒 Locks	(U) acao-ora01-log_2022-0	9-13T124501-7472967Z South Central US	09/13/2022, 08:45:04 AM		
Locks	anf2-z1-stdds03 100 GiB	Storage service	(L) acao-ora01-log_2022-0	9-13T131501-6270448Z South Central US	09/13/2022, 09:15:04 AM		
Azure NetApp Files	anf2-z1-stdds04 100 GiB	O Mount instructions	(U) acao-ora01-log_2022-0	9-13T133001-7773619Z South Central US	09/13/2022, 09:30:04 AM		
Active Directory connections	anf2-z1-stdds05 100 GiB	Export policy	(E) acao-ora01-log_2022-0	9-13T134502-0449919Z South Central US	09/13/2022, 09:45:04 AM	store to new volume	
torage service	anf2-z1-stdds06 100 GiB	(D) Snapshots				vert volume	
Capacity pools	anf2-z1-stdds07 100 GiB	Replication			De De		
Volumes	anf2-z1-stdds08 100 GiB	Monitoring				iete	
Application volume groups	anf-z1-stdds01 6 TiB	iá Metrics					
	anf-z1-stdds02 200 GiB						
Data protection	anf-z1-stdds03 1 TiB	Automation					

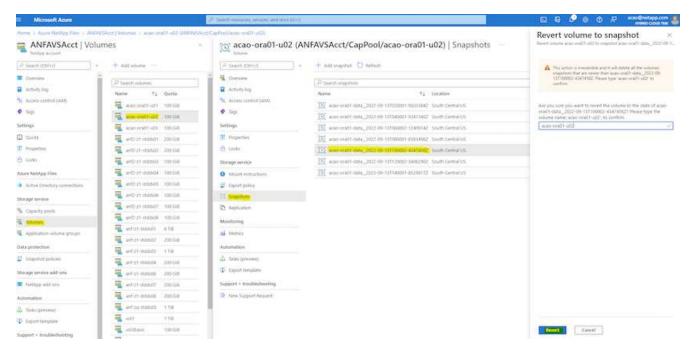
4. Confirm revert volume and click **Revert** to complete the volume reversion to the latest available backup.

Microsoft Azure		(Pittaard resistor, seven, and dos ()-	t de la companya de la		
ANFAVSAcct Vol	wskod j Warnen (Facas craft sids (ANFA) umes		ANFAVSAcct/CapPool/acao-ora01-u03) Snapshots	Revert volume to snapshot >> Anot stars social? di la regent social ling_322 = 0	
/ Tearth (Chin) + + Add volume		JP Search (Ctrix /)	+ automptor () setest.	A This action is incremitive and it will strate all the outprist	
Dummer Ametricage	P Seath volumes	a Guernese	J ² Guich stagebots	examples that are rever than also well in a _3022.00 121134022-06400162 Please type local cost11402 to continue	
Access correct DAMO	Nate 14 October accessed with 100 Ge accessed 1401 100 Ge accessed 1401 100 Ge	Access control JAAMS	Noive * Economic [0] scale (stab) data_2502 dilli UT100008 XT147082 Scale Careford IIII [0] scale (stab) data_2502 dilli UT100028 XT147082 Scale Careford IIII [0] scale (stab) data_2502 dilli UT100028 XT147082 Scale Careford IIII	Are plot size pool wath thi revert the volume to the role of acasi- exact lines_call do 11113400 044001027 Nesse type the volume name, acan ocal 1 v007, to confirm.	
Settingi Di Costa		Settings III Proparties	332 acar-read img_2022-00-11112001-62801172 South Control LS 332 acar-read img_2022-00-11112001-00006482 South Control LS		
T Properties	R +0-11 +5+62 - 300 G4	🖰 toda	acar over mg_pacy or ring over an over a sub- contact of acar over mg_pacy or ring over a sub- sub- sub- contact over a sub- sub- sub- sub- sub- sub- sub- sub-		
A soos	-D-11-H36531 109-04	Morage service	30 acar-insol exp_2022-00-11101301-62104462 South Central US		
Kaura Netikop Film	Bill oth strated a 100 GB	0 Multistudies	III acar-crapt-log_2022-04-1111/2001 /7716152 (com/carmelists		
 Active Directory conhections 	-100 GB	B bpoteko	33 and and any and an intrace country sound and the		
itorage service	100 GB	12 Septem			
Casely point		C Replication			
Application infume groups	artici sociali il 18	al United			
U trupilot pulkies	art (1.60022) 118	de talas (preview)			
torage service add erro	🗮 art 21 statuts - 200 GR	Toport turoplate			
Nitiap add oni	art st ussui? - 200.06	Support - traubieshooting			
lutonution	📲 erfist essaita - 200.048	R New Support Anguint			
(oport tempide	R and periodicity 1.176	1			
Support + traulileshooting	📲 woldbasic 🛛 100 Gal			Cancel	

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

Microsoft Azure		P Search resources, services, and docs (G	s+/)			асао@netapp.coi нтвяр стопр ти
iome > Azure NetApp Files > ANFA	VSAcct Volumes > acao-ora01-u02 (ANFAVSA					
ANFAVSAcct Volu	Imes	C acao-ora01-u02	(ANFAVSAcct/CapPool	/acao-ora01-u02) Snapshots 🦷		
P Search (Ctrl+/) «	+ Add volume	P Search (Ctrl+/)	🛛 « 🛛 🕂 Add snapshot 🕐 Refresh			
Overview	Search volumes	Overview	Search snapshots			
Activity log	Name 14 Quota	Activity log	Name	↑↓ Location	↑↓ Created	¢↓
Access control (IAM)	acao-ora01-u01 100 GiB	Access control (IAM)	acao-ora01-data_2022-0	9-13T020001-9503384Z South Central US	09/12/2022, 10:00:19 PM	
Tags	acao-ora01-u02 100 GiB	Tags	(L) acao-ora01-data_2022-0	9-13T040001-9341340Z South Central US	09/13/2022, 12:00:18 AM	
ettings	📕 acao-ora01-u03 100 GiB	Settings	(L) acao-ora01-data_2022-0	-13T060002-1240914Z South Central US	09/13/2022, 02:00:19 AM	
Quota	anf2-z1-stdds01 200 GiB	Properties	(E) acao-ora01-data_2022-0	-13T080001-8383498Z South Central US	09/13/2022, 04:00:18 AM	
Properties	anf2-z1-stdds02 200 GiB	🔒 Locks	(C) acao-ora01-data_2022-0	-13T100002-4347456Z South Central US	09/13/2022, 06:00:18 / 🕚 Restore	e to new volume
Locks	anf2-z1-stdds03 100 GiB	Storage service	acao-ora01-data_2022-0	9-13T120002-3406290Z South Central US	09/13/2022, 08:00:18 / 🔗 Revent	
zure NetApp Files	anf2-z1-stdds04 100 GiB	1 Mount instructions	(U) acao-ora01-data_2022-0	9-13T140001-8529817Z South Central US	09/13/2022, 10:00:17 /	
Active Directory connections	anf2-z1-stdds05 100 GiB	Export policy				
orage service	anf2-z1-stdds06 100 GiB	(G) Snapshots				
Capacity pools	anf2-z1-stdds07 100 GiB	Replication				
Volumes	anf2-z1-stdds08 100 GiB	Monitoring				
Application volume groups	anf-z1-stdds01 6 TiB	Metrics				
ita protection	anf-z1-stdds02 200 GiB	Automation				

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



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

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

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

```
[oracle@acao-ora01 ~]$ sqlplus / as sysdba
SQL*Plus: Release 19.0.0.0.0 - Production on Tue Sep 13 15:10:17 2022
Version 19.8.0.0.0
Copyright (c) 1982, 2019, Oracle. All rights reserved.
Connected to an idle instance.
SQL> startup mount;
ORACLE instance started.
Total System Global Area 6442448984 bytes
Fixed Size
                            8910936 bytes
Variable Size
                         1090519040 bytes
Database Buffers
                         5335154688 bytes
Redo Buffers
                            7864320 bytes
Database mounted.
```

SQL> recover database using backup controlfile until cancel; ORA-00279: change 3188523 generated at 09/13/2022 10:00:09 needed for thread 1 ORA-00289: suggestion : /u03/orareco/ORATST/archivelog/2022_09_13/o1_mf_1_43__22rnjq9q_.arc ORA-00280: change 3188523 for thread 1 is in sequence #43 Specify log: {<RET>=suggested | filename | AUTO | CANCEL} ORA-00279: change 3188862 generated at 09/13/2022 10:01:20 needed for thread 1 ORA-00289: suggestion : /u03/orareco/ORATST/archivelog/2022 09 13/o1 mf 1 44 29f2lgb5 .arc ORA-00280: change 3188862 for thread 1 is in sequence #44 ORA-00278: log file '/u03/orareco/ORATST/archivelog/2022 09 13/o1 mf 1 43 22rnjq9q .arc' no longer needed for this recovery Specify log: {<RET>=suggested | filename | AUTO | CANCEL} ORA-00279: change 3193117 generated at 09/13/2022 12:00:08 needed for thread 1 ORA-00289: suggestion : /u03/orareco/ORATST/archivelog/2022 09 13/o1 mf 1 45 29h6qqyw .arc ORA-00280: change 3193117 for thread 1 is in sequence #45 ORA-00278: log file '/u03/orareco/ORATST/archivelog/2022 09 13/o1 mf 1 44 29f2lgb5 .arc' no longer needed for this recovery Specify log: {<RET>=suggested | filename | AUTO | CANCEL} ORA-00279: change 3193440 generated at 09/13/2022 12:01:20 needed for thread 1 ORA-00289: suggestion : /u03/orareco/ORATST/archivelog/2022 09 13/o1 mf 1 46 %u .arc ORA-00280: change 3193440 for thread 1 is in sequence #46 ORA-00278: log file '/u03/orareco/ORATST/archivelog/2022 09 13/o1 mf 1 45 29h6qqyw .arc' no longer needed for this recovery Specify log: {<RET>=suggested | filename | AUTO | CANCEL} cancel Media recovery cancelled.

```
SQL> alter database open resetlogs;
Database altered.
SQL> select * from testsnapshot;
 ID
_____
EVENT
_____
DT
_____
       1
insert a data marker to validate snapshot restore
12-SEP-22 07.07.35.000000 PM
SQL> select systimestamp from dual;
SYSTIMESTAMP
 _____
13-SEP-22 03.28.52.646977 PM +00:00
```

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

Database migration from on-premises to Azure cloud

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

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

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

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

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

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

or PDB.

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

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

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

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

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

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

7. Plug in the single instance non-CDB.

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



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

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

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

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

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

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

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

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

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

Migrate on-premises Oracle databases to Azure with PDB relocation

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

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

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

git clone https://github.com/NetApp-Automation/na ora aws migration.git

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

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

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

```
ansible-playbook -i hosts ora_pdb_relocate.yml -u admin -k -K -t ora pdb relo onprem
```



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

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

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



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

Additional Oracle database migration options

Please see the Microsoft documentation for additional migration options: Oracle database migration decision process.

On-Premises/Hybrid Cloud

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

App2 App1 Oracle DB1 with Restart Oracle DB2 with Restart Snap center Ansible ora_02 ora_01 control public network: mtu 1500 iscsi a: mtu 9000 iscsi b: mtu 9000 a0a a0a e0e e0e e0f e0g e0g e0h e0f e3a e3a e3b data04 data03 data03 NetApp ASA controller-1 NetApp ASA controller-2

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

Oracle database configuration in the lab environment

Server	Database	DB Storage
ora_01	NTAP1(NTAP1_PDB1,NTAP1_PD B2,NTAP1_PDB3)	iSCSI luns on ASA A400
ora_02	NTAP2(NTAP2_PDB1,NTAP2_PD B2,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 multipath 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.

- 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"

 (\mathbf{i})

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
######
*****
# Enter the supported ONTAP platform: on-prem, aws-fsx.
ontap platform: on-prem
# Enter ONTAP cluster management user credentials
username: "xxxxxxx"
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 nodel:
 - {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: xxxxxxx
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: xxxxxxx
```

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
                     7.8G 0 7.8G 0% /sys/fs/cgroup
tmpfs
                     44G 38G 6.8G 85% /
/dev/mapper/rhel-root
                    1014M 258M 757M 26% /boot
/dev/sda1
                     1.6G 12K 1.6G 1% /run/user/42
tmpfs
                     1.6G 4.0K 1.6G 1% /run/user/1000
tmpfs
/dev/mapper/ora 01 biny 01p1 40G 21G 20G 52% /u01
[oracle@ora 01 ~]$ asm
[oracle@ora 01 ~]$ crsctl stat res -t
_____
_____
         Target State Server
Name
                                           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 Ol
                                           STABLE
_____
Cluster Resources
   _____
```

```
_____
ora.cssd
    1 ONLINE ONLINE
                        ora 01
                                           STABLE
ora.diskmon
    1
         OFFLINE OFFLINE
                                           STABLE
ora.driver.afd
    1
       ONLINE ONLINE ora 01
                                           STABLE
ora.evmd
      ONLINE ONLINE ora_01
    1
                                           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
 0 N DATA/ MOUNTED EXTERN N 512 512 4096 4194304 78880 81920 78880 0 0 N LOGS/ ASMCMD> lsdsk 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.

←	\rightarrow	C		Not secure	https://10.61.180.21:5500/em/login
---	---------------	---	--	------------	------------------------------------

	Username Password Container Name Log in		
		ORACLE	
RACLE Enterprise Manager Database Express		Ŕ	
C Not secure Https://10.61.180.21:5500/em/shell RACLE Enterprise Manager Database Express TAP1 (19.18.0.0) Performance Storage tabase Home eZone wwser (GMT-05:00)			Image: The second se
TAP1 (19.18.0.00) Performance * Storage * tabase Home > Zone	Performance Activity Services Containers 40	1 min Ac	system uto-Refresh V Refresh

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.GETHTTPSPORT() _____ 0 SQL> exec DBMS XDB CONFIG.SETHTTPSPORT(5501); PL/SQL procedure successfully completed. SQL> select dbms xdb config.gethttpsport() from dual; DBMS_XDB_CONFIG.GETHTTPSPORT() -----5501 login to NTAP1 PDB1 from port 5501.

Performance Activity Services 12 0.0 0.4 0.0 0.51500 PM 01:57:20 PM 01:51:00 PM 01:57:20 PM 01:51:00 PM 01:57:20 PM 01:51:00 PM 02:03:40 PM 02:10:00 PM 02:10:00 PM 02:10:00 PM 02:10:00 PM 02:10:00 PM 02:22:40 PM 02:25:20 PM 02:41:40 PM 02:41:40 PM 02:48:00 PM	Browser (GMT-05:00) 🛛 🔻							1 min Auto-R	Refresh 🔻	Refre
Activity Services 12 0.8 0.0 0.15720 PM 020340 PM 022240 PM 022240 PM 022500 PM 024140 PM 024800 PM 024800 PM										
1.2 0.8 0.4 0.0 01:57:00 PM 02:03:00 PM 02:16:00 PM 02:16:20 PM 02:22:40 PM 02:29:00 PM 02:35:20 PM 02:41:40 PM 02:45:00 PM	Status		Performance							
04 04 00 01:57:00 PM 02:03:00 PM 02:10:00 PM 02:16:20 PM 02:22:40 PM 02:29:00 PM 02:35:20 PM 02:41:40 PM 02:45:00 PM		1 hours, 23 minutes, 14 seconds	Activity Se	ervices						
04 04 00 01:57:00 PM 02:03:00 PM 02:10:00 PM 02:16:20 PM 02:22:40 PM 02:29:00 PM 02:35:20 PM 02:41:40 PM 02:45:00 PM		Single Instance (NTAP1)								
0.4 0.4 0.0 01:51:00 PM 02:03:40 PM 02:10:00 PM 02:16:20 PM 02:22:40 PM 02:29:00 PM 02:35:20 PM 02:41:40 PM 02:46:00 PM		19.18.0.0.0 Enterprise Edition								
00		Linux x86 64-bit	0.8	0.8						
01:51:00 PM 01:57:20 PM 02:03:40 PM 02:10:00 PM 02:16:20 PM 02:22:40 PM 02:29:00 PM 02:35:20 PM 02:41:40 PM 02:48:00 PM	Thread		0.4							
	Archiver Stopped									
Nov 8, 2023 GMI-10500	Last Dealure Times		0.0						CPU	
	Last Backup Time	N/A	01:51:00 PM 01		02:16:20 PM 02:22:40 PM	1 02:29:00 PM 02	:35:20 PM 02:41:40 P		CPU	
1717.1.0	Last Backup Time Incident(s)	N/A	01:51:00 PM 01		02:16:20 PM 02:22:40 PM	4 02:29:00 PM 02	:35:20 PM 02:41:40 P		CPD	
	Incident(s) Resources	N/A 0	01:51:00 PM 01 Nov 8, 2023 GMT-	7-05:00	02:16:20 PM 02:22:40 PM	И 02:29:00 PM 0;			CP0	
	Incident(s) Resources 0 ms/s	N/A 0	015100 PM 01 Nev 8, 2023 GMT-	г-05:00 171.7 МВ.	02:16:20 PM 02:22:40 PM	и 02:29:00 РМ 0;	2 GB			
143.1 MB 1.7 GB USER	Incident(s) Resources	N/A 0	01:51:00 PM 01 Nov 8, 2023 GMT-	171.7 M8	02:16:20 PM 02:22:40 PM	и 02:29:00 РМ 0;	2 G8 1.7 GB			
143.1 MB 1.7 GB USER 114.4 MB 1.4 GB USER	Incident(s) Resources 0 ms/s 0 ms/s	N/A 0	015100 PM 01 Nov 8, 2023 GMT-	171.7 M8 143.1 M8 144.1 M8 114.4 M8	02:16:20 PM 02:22:40 Ph		2 GB 1.7 GB 1.4 GB			UNDO
143.1 MB 1.7 GB USER 144. MB 14.6 GB USER 1144. MB 11.6 UNDO 85.8 MB total_spa CPU 572. MC total_spa total_spa total_spa total_spa	Incident(s)	N/A 0	0.035 0.020 0.025	171.7 M8 143.1 M8 1144.4 M8 CPU 85.8 M8	02:16:20 PM 02:22:40 PM	total_sga	2 GB 1.7 GB 1.4 GB 1.1 GB 858.3 MB			UNDO TEMPORARY
143.1 M8 1.7 G8 USER 144.4 M8 1.4 G8 UNOO 114.4 M8 1.1 G8 UNOO 11.1 G8 TEMPORARY 57.2 M8 558.3 M8 S72.2 M8 S752.4 M8 S7572.4 M8	Incident(s)	N/A 0 Background Foreground	0.035 0.030 0.035 0.030 0.025 0.030 0.015 0.010	171.7 M6 143.1 M8 114.4 M8 ■ CPU \$7.2 M8	02:16:20 PM 02:22:40 PM	total_sga	2 GB 1.7 GB 1.4 GB 1.1 GB 8553 MB 572.2 MB			UNDO TEMPORARY SYSTEM
143.1 M8 1.7 GB USER 114.4 M8 1.4 GB UNDO 85.8 MB total_sign 85.8 MB TMPORARY 26.6 MB total_pga 572.2 MB SYSTEM	Incident(s)	N/A 0 Background Foreground	0.035 0.020 0.015 0.020 0.015 0.010 0.015	171.7 M8 143.1 M8 114.4 M8 CPU 28.6 M8 28.6 M8	02:16:20 PM 02:22:40 PM	total_sga	2 (68 1.7 (68 1.4 (68 1.1 (68 8553.3 MB 572.2 MB 286.1 MB			UNDO TEMPORARY SYSTEM SYSAUX
171.7 MR 2 150		N/A	01:51:00 PM 01		02:16:20 PM 02:22:40 PM	4 02:29:00 PM 02	:35:20 PM 02:41:40 P		CPU	
143.1 MB 1.7 GB	Incident(s) Resources 0 ms/s 0 ms/s	N/A 0	015100 PM 01 Nov 8, 2023 GMT-	171.7 M8	02:16:20 PM 02:22:40 PM	и 02:29:00 РМ 0;	2 G8 1.7 GB			JSER
143.1 MB 1.7 GB USER 114.4 MB 14.6 B USER S S S MS 11.6 B USER	Incident(s) Resources 0 ms/s 0 ms/s 0 ms/s	N/A 0	0,035 0,039	171.7 M8 143.1 M8 144.1 M8 114.4 M8	02:16:20 PM 02:22:40 PM		2 GB 1.7 GB 1.4 GB 1.1 GB			UNDO
143.1 MB 1.7 GB USER 114.4 MB 1.4 GB UNDO 116.8 SSB.MB TEMPORARY 116.8 TEMPORARY	Incident(s)	N/A 0 Background	0.035 0.020	171.7 M8 143.1 M8 1144.4 M8 CPU 85.8 M8	02:16:20 PM 02:22:40 PM	total_sga	2 GB 1.7 GB 1.4 GB 1.1 GB			UNDO TEMPORARY
143.1 MB 1.7 GB USER 114.4 MB 1.4 GB UNDO 114.4 MB 1.1 GB UNDO € CPU 85.8 MB total_sga 55.8 MB TEMPORARY 57.2 MB total_ga 57.2 MB SVSTM SVSTM	Incident(s)	N/A 0	0.035 0.020 0.025	171.7 M6 143.1 M8 114.4 M8 ■ CPU \$7.2 M8	02:16:20 PM 02:22:40 PM	total_sga	2 GB 1.7 GB 1.4 GB 1.1 GB 858.3 MB			UNDO TEMPORARY SYSTEM
143.1 MB 1.7 GB USER 144.4 MB 14 GB UNDO 114.4 MB 11.1 GB UNDO 1.1 GB TOMODO 1.1 GB TOMODO 1.1 GB TOMODO 1.1 GB TOMODO 1.1 GB SYSTEM 1.1 SYSTEM 1.1 SYSTEM 1.1 SY	Incident(s)	N/A 0 Background Foreground	0.035 0.030 0.035 0.030 0.025 0.030 0.015 0.010	171.7 M6 143.1 M8 114.4 M8 ■ CPU \$7.2 M8	02:16:20 PM 02:22:40 PM	total_sga	2 GB 1.7 GB 1.4 GB 1.1 GB 8553 MB 572.2 MB			UNDO TEMPORARY SYSTEM SYSAUX
143.1 MB 1.7 GB USER 114.4 MB 1.4 GB UNCO 114.4 MB 1.1 GB UNCO E CPU 85.8 MB total_sga 858.3 MB TMPORARY 28.6 MB total_pga 572.2 MB SYSTEM SYSTEM	Incident(s) Resources Oms/s Oms/s Oms/s Oms/s Oms/s Oms/s	N/A 0 Background Foreground	0.035 0.020 0.015 0.020 0.015 0.010 0.015	171.7 M8 143.1 M8 114.4 M8 CPU 28.6 M8 28.6 M8	02:16:20 PM 02:22:40 PM	total_sga	2 (68 1.7 (68 1.4 (68 1.1 (68 8553.3 MB 572.2 MB 286.1 MB			UNDO TEMPORARY SYSTEM SYSAUX
143.1 MB 1.7 GB USER 114.4 MB 1.4 GB UNDO 114.4 MB 1.1 GB UNDO 6 CPU 85.8 MB Itotal.sga 858.3 MB Itotmore 57.2 MB 572.2 MB Systam Systam Systam 28.6 MB 0.8 0.6 0.8 Itotal.sga	Incident(s)	N/A 0 Background Foreground	0.035 0.030 0.035 0.030 0.025 0.030 0.015 0.010	171.7 M8 143.1 M8 114.4 M8 CPU 28.6 M8 28.6 M8		total_sga	2 (68 1.7 (68 1.4 (68 1.1 (68 8553.3 MB 572.2 MB 286.1 MB	NU 02:48:00 PM		UNDO TEMPORARY SYSTEM SYSAUX

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 is 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.

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

Requirements

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 \rightarrow Inventories \rightarrow 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 \rightarrow 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 \rightarrow 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 \rightarrow 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
```

{{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"}
    - {vol_name: "{{groups.oracle[0]}}_u03", aggr_name: "aggr01_node01",
```

- 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.

```
#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: "eOM"}
# 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"
****
```

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 \rightarrow Templates \rightarrow 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 \rightarrow 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 \rightarrow 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 \rightarrow 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 \rightarrow 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 \rightarrow Jobs to monitor the job output and progress.
- I. After the linux_config role has completed, run the process again for oracle_config.
- m. Go to Resources \rightarrow 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 \rightarrow 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$database;
NAME LOG MODE
_____ ___
CDB2 ARCHIVELOG
SQL> show pdbs
  CON ID CON NAME
                               OPEN MODE RESTRICTED
_____ _ ____
                                   READ ONLY NO
       2 PDB$SEED
       3 CDB2 PDB1
                                   READ WRITE NO
                                  READ WRITE NO
READ WRITE NO
       4 CDB2 PDB2
       5 CDB2 PDB3
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
                                                    NFSVERSION
                          DIRNAME
_____
_____
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 hte Oracle listener configuration with the following command. Change to the appropriate listener port and database service name.

```
[oracle@localhost ~]$ sqlplus
system@//localhost:1523/cdb2_pdbl.cie.netapp.com
SQL*Plus: Release 19.0.0.0 - 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

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"}

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 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",

```
- {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"

#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"
```

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
```

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$database;
NAME LOG MODE
_____ ___
CDB2 ARCHIVELOG
SQL> show pdbs
  CON ID CON NAME
                               OPEN MODE RESTRICTED
_____ _____
                                   READ ONLY NO
       2 PDB$SEED
       3 CDB2 PDB1
                                   READ WRITE NO
                                  READ WRITE NO
READ WRITE NO
       4 CDB2 PDB2
       5 CDB2 PDB3
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
                                                    NFSVERSION
                         DIRNAME
_____
_____
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 hte Oracle listener configuration with the following command. Change to the appropriate listener port and database service name.

```
[oracle@localhost ~]$ sqlplus
system@//localhost:1523/cdb2_pdbl.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 - xmltodict - 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 - xmltodict - 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.

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 \rightarrow Inventories \rightarrow 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 \rightarrow 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 \rightarrow 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 \rightarrow 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 \rightarrow 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 \rightarrow 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 \rightarrow 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 \rightarrow 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 \rightarrow 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"
    - "e0q"
  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"
   - "e0q"
 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 secuirty group that allows access to on-prem
resources.
security group: "sg-123123123"
# You 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
****
```

#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"

```
#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 \rightarrow Templates \rightarrow 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 \rightarrow 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 \rightarrow 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 \rightarrow 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 \rightarrow 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 \rightarrow 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 \rightarrow 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.



Recovering Oracle Database

- 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 paraellel 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/oraInst.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

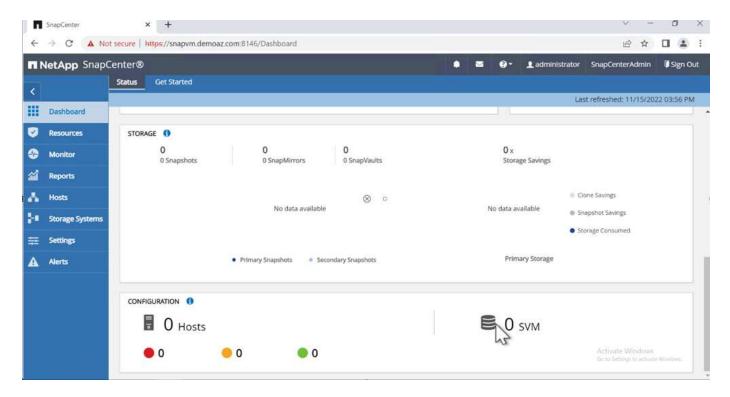
SnapCenter Configuration

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 (+).

÷	→ C ▲ No	t secure https://sna	pvm.demoa	z.com:8146/Administ	ration/Viewl	Name=RunAs				_	_	⊮ ☆	ш а
1	letApp Snap(Center®						•	i 0-	1 administrator	SnapCer	nterAdmin	🕻 Sign Out
<		Global Settings	Policies	Users and Access	Roles	Credential	Software						
	Dashboard	Search by Creden	tial Name								+ New	<i>.</i>	
2	Resources	Credential Name				Authenticati	on Mode				Deta	ails	
0	Monitor	There is no match f	or your searc	h or data is not availa	ible.								
M	Reports												
1	Hosts					\otimes	0						
	Storage Systems												
æ													
Â	Alerts												

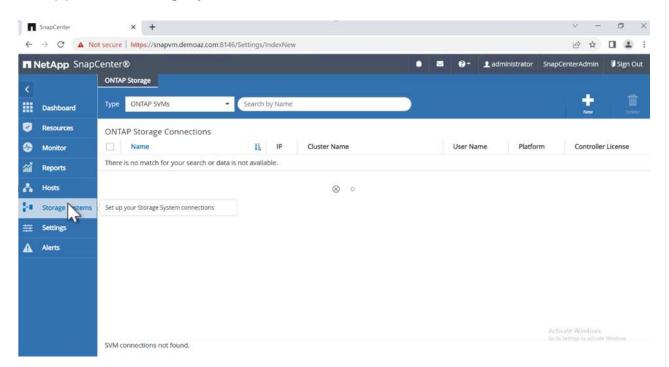
The new user must have administrator rights on the SQL Server host.

Credential Name	Demoaz	
uthentication Mode	Windows 🔹	
Username	demoaz\clusteradmin	0
Password		
	\otimes \circ	

Configure storage

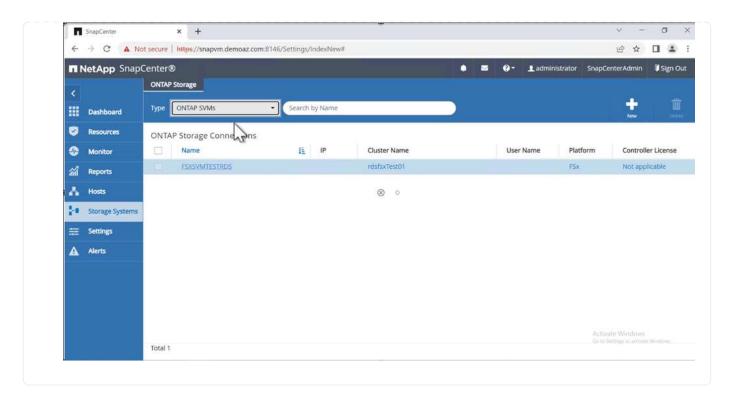
To configure storage in SnapCenter, complete the following steps:

- 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.



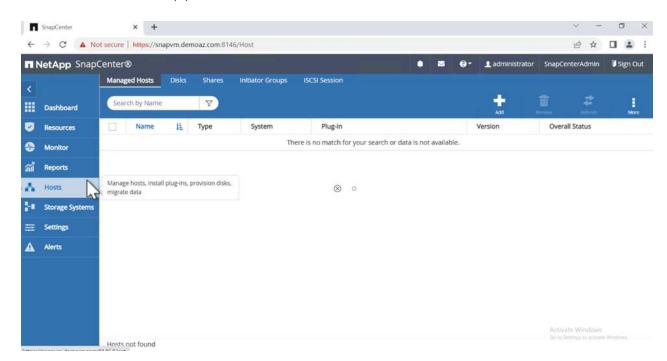
3. Provide the FSx for ONTAP management endpoint.

IN	etApp SnapCenter®		•	9-	1 administrator	SnapCenterAdmin	🖡 Sign	Out
	ONTAP Storage	Add Storage System						
		Add Storage System 0						
2	ONTAP Storage Connections	Storage System 172,30.0,98						
9	🗌 Name Ià	Username fsxadmin						
a	There is no match for your search or data is not available.	Password						
6		Event Management System (EMS) & Auto@pport Settings						
•		Send AutoSupport notification to storage system						
=		Log SnapCenter Server events to syslog						
		More Options : Platform, Protocol, Preferred IP etc						
		Submit Cancel Reset						
						Activate Windows So to Settings to activate		
	SVM connections not found.							



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.

Managed Hosts Search by Name Name Intere is no match for your search or data is not available. Host Type Host Name FSXCLUSTER.Demoaz.com Credentials Demoaz Select Plug-ins to Install SnapCenter Plug-ingPackage 4.7 for Windows Microsoft Windows Microsoft SQL Server SAP HANA More Options : Port, gMSA, Install Path, Custom Plug-Ins Submit	IN	etApp SnapCenter®			•		9-	1 administrator	SnapCenterAdmin	Sign (Ou
Search by Name Name There is no match for your search or data is not available. Host Type Windows Host Name FSXCLUSTER.Demoaz.com Credentials Demoaz Select Plug-ins to Install SnapCenter Plug-ing:Package 4.7 for Windows Image: Plug-ins to Install SnapCenter Plug-ing:Package 4.7 for Windows Image: Plug-ins to Install SnapCenter Plug-ing:Package 4.7 for Windows Image: Plug-ins to Install SnapCenter Plug-ing:Package 4.7 for Windows Image: Plug-ins to Install SnapCenter Plug-ing:Package 4.7 for Windows Image: Plug-ins to Install SnapCenter Plug-ing:Package 4.7 for Windows Image: Plug-ins to Install SnapCenter Plug-ing:Package 4.7 for Windows Image: Plug-ins to Install SnapCenter Plug-ing:Package 5erver Sap Hana More Options : Port, gMSA, Install Path, Custom Plug-ins	>	Managed Hosts									
 There is no match for your search or data is not available. Host Name FSXCLUSTER.Demoaz.com Demoaz Select Plug-ins to Install SnapCenter Plug-ing-Package 4.7 for Windows Microsoft Windows Microsoft Windows Microsoft SQL Server SAP HANA More Options : Port, gMSA, Install Path, Custom Plug-Ins 		Search by Name	Add Host								
 Mice is not available. Credentials Credentials Demoaz Select Plug-ins to Install SnapCenter Plug-ing-Package 4.7 for Windows Microsoft Windows Microsoft SQL Server SAP HANA More Options : Port, gMSA, Install Path, Custom Plug-Ins 	2	Name 1	Host Type	Windows	٠						
Select Plug-ins to Install SnapCenter Plug-ing/Package 4.7 for Windows Microsoft Windows Microsoft SQL Server SAP HANA More Options : Port, gMSA, Install Path, Custom Plug-Ins	P		Host Name	FSXCLUSTER.Demoaz.com							
Select Plug-ins to Install SnapCenter Plug-ing-Package 4.7 for Windows Microsoft Windows Microsoft SQL Server SAP HANA More Options : Port, gMSA, Install Path, Custom Plug-Ins		not available.	Credentials	Demoaz	٠	+					
	-		More Options ; Po	Microsoft Windows Vicrosoft SQL Server Vicrosoft Exchange Server SAP HANA							

The SQL Server packages are installed.

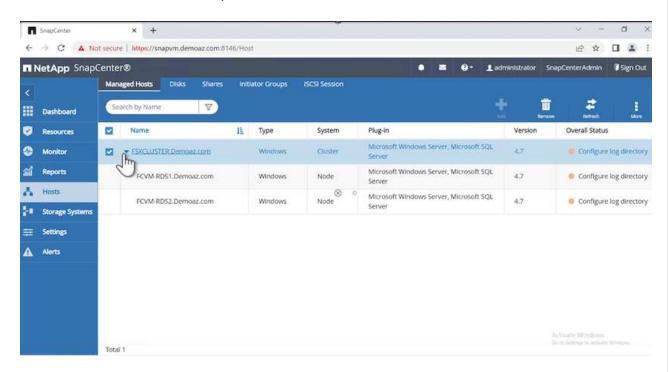
		t secure https://snapvn		141111111111111111111111111111111111111					~			25	100
Net	App SnapC							• =	9-	1 administr	rator Sna	apCenterAdmin	Sign Out
٢		Managed Hosts D	isks Shares	initi	ator Groups	ISCSI Session							
🛛 Da	shboard	Search by Name	V							+		2	More
🕗 Re	sources	Name		臣	Туре	System	Plug-in			Ver	rsion	Overall Status	r .
De Mo	onitor	ESXCLUSTER.D	emoaz.com		Windows	Cluster						Installing	plug-in
🖬 Re	ports												
Но	osts					\otimes	0						
Sto	orage Systems												
	orage Systems ttings												
⊒ Sel													
≣ Se	ttings												
E Sel	ttings												
⊒ Se	ttings												

1. After the installation is complete, go to the **Resource** tab to verify whether all FSx for ONTAP iSCSI volumes are present.

	letApp Snap(File Sys				• =	🛛 🔹 👤 administra	tor SnapCenter/	Admin 🔋 Sign Out	
< 	Dashboard	View	Path	• search	V			Referent Ref	aran New Resource Cou	
0	Resources	ja	Name 11	Host	Storage Layout	Resource Groups	Policies	Last Backup	Overall Status	
•	Monitor		D:\	F5XCLUSTER.Demo STER.Demoaz.com	FSXSVMTESTRDS:/ FCIDATA/FCIDATA				Not protected	
a i	Reports		E/\	FSXCLUSTER.Demo STER.Demoaz.com	FSXSVMTESTRDS/ I/FCILOG/FCILOG				Not protected	
А 34	Hosts		ΕA	FSXCLUSTER.Demo STER.Demoaz.com	FSXSVMTESTRDS:/ ACKUP/FCIBACKUP				Not protected	
	Storage Systems Settings		G/\	FSXCLUSTER.Demo STER.Demoaz.com	FSXSVMTESTRDS/				Not protected	
4	Alerts		H3	FSXCLUSTER.Demo STER.Demoaz.com	FSXSVMTESTRDS:/				Not protected	
			К3	FSXCLUSTER.Demo STER.Demoaz.com	FSXSVMTESTRDS:/				Not protected	
			վեղ							
		Total 6	0							

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.

SnapCenter	× +			20				v -	٥	×
← → C ▲ N	lot secure https://snapvm.der	noaz.com:8146/Host						18 ☆		÷
n NetApp Snap	oCenter®				• =	8-	1 administrator	SnapCenterAdmin	🗊 Sign O	lut
> Managed Hos	sts									×
Search by N	lame	Host Details								
Name	e IL		FSXCLUSTER.Demoaz.co 172.30.5.143	m		Aler	ts			
	Ð	Host Type System	Cluster &) 0 ckage 4.7.0.877 for		No A	lerts			
Total 1		More Options : Po Submit Cancel	ort, gMSA, Install Path, Ad	id Plug-Ins				Activate Windows Ge to Settings to action	s Wordsweg.	

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.

5napCenter	× +			v - 0
← → C ▲ Not se	cure https://snapvm.demoaz	.com:8146/Host#		ie 🖈 🛛 😩
	Configure Plug	r-in for SQL Server	×	
	Configure the log	backup directory for FSXCLUSTER.Demoaz.com		
* 1 C	Configure host lo	g directory		
	Host	FCVM-RDS1.Demoaz.com		
	Host log directory	N:\SNAPFCVM1LOG\	🗄 Apply 🖿 Browse	
	Configure FCI ins	tance log directory		
	FCI Instance	SQLHMCRDS -		
	FCI log directory	G:\SNAPLOGFCI\	Apply Browse	
			Save Close	

The host is now in a running state.

NetApp Sna						•		8-	1 administrate	or SnapGe	nterAdmin	🛿 Sign C
C Deshboard	Managed Hosts Search by Name	Disks Shares	Initiator	Groups E	SCSI Session				Add		4	
Resources	Name		1E	Туре	System	Plug-in				Version	Overal	li Status
Monitor	ESKCLUSIE	B.Demoaz.com		Windows	Cluster	Microsoft We Server	ndows Ser	ver, Micr	osoft SQL	4.7	Ru Ru	inning
											2m	
Reports											9	
Reports											0	
Hosts											0	
Hosts Storage Systems											0	
Hosts											0	
Hosts Storage Systems Settings											0	
Hosts Storage Systems Settings											0	

1. From the **Resources** tab, we have all the servers and databases.

	NetApp Snap(øft SQL Server 💽			• = •-	1 administrator	SnapCenterAdmin Sigr
< Ⅲ	Dashboard	View	Database -	Search by name	7			tertent lissource - New News
	Resources	12	Name	Instance	Host	Last Backup	Overall Status	Туре
	Monitor		INVENTORYDB1	SQLHMCRDS	FSXCLUSTER.Demoaz. com		Not protected	User database
ai	Reports		INVENTORYD82	SQLHMCRDS	FSXCLUSTER.Demoaz.		Not protected	User database
소 24	Hosts Storage Systems		INVENTORYDB3	SQLHMCRDS	PSXCLUSPER.Demoaz.		Not protected	User database
	Settings		INVENTORYD84	SQLHMCRDS	PSXCLUSTER.Demoaz. com		Not protected	User database
▲	Alerts.		INVENTORYD85	SQLHMCRDS	PSXCLUSTER.Demoaz. com		Not protected	User database
			master	SQLHMCRDS	PSXCLUSTER.Demoaz. com		Not protected	System database
			model	SQLHMCRDS	PSXCLUSTER.Demoaz.		Not protected	System database
		Total 9	-					

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.

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**.

NetApp Snap	Center®							0-	1 dusteradmin	SnapCenterAdmin	Sign Out
C Dashboard	Global Settings Policies Utiers and Microsoft SQL Server			Software -				÷			
Monitor	Name		Backup Type		Schedule Type	Replication			Verification		
a Reports	There is no match for your search or data is View or export activity reports	not availabl	le.								
liosta											
Storage Systems											
E Settings											
CONTRACTOR OF A DECK											

2. Enter the policy name and a description. Click **Next**.

New SQL Serve	er Backup Policy			×	
1 Name	Provide a polic	/ name			
2 Backup Type	Policy name	TestDB-fullbackup-policy		0	
3 Retention	Details	TestDB-fullbackup-policy			
4 Replication					
5 Script		⊗ ∘		100	
6 Verification				1.00	
7 Summary					
			Previous	Net	

3. Select **Full backup** as the backup type.

New SQL Serve	er Backup Policy	×
1 Name	Select SQL server backup options	ŕ
2 Backup Type	Choose backup type	
3 Retention	O Full backup and log backup	
4 Replication	Full backup	
5 Script	○ Log backup ⊗ ○	
6 Verification	Maximum databases backed up per Snapshot copy: 100	
7 Summary		
	Availability Group Settings	~

4. Select the schedule frequency (this is based on the company SLA). Click **Next**.

New SQL Serve	er Backup Policy	×		
Name	Copy only backup			
2 Backup Type	Maximum databases backed up per Snapshot copye 100			
3 Retention	Availability Group Settings			
4 Replication				
5 Script	Schedule frequency 🛞 o			
6 Verification	Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling yo stagger your start times. O On demand	ou to		
7 Summary	Hourly Daily			
	Previous	Nhin		

5. Configure the retention settings for the backup.

	New SQL Server	r Backup Policy	×
	1 Name	Retention settings	
	2 Backup Type	Retention settings for up-to-the-minute restore operation 🚯	
	3 Retention	Keep log backups applicable to last 7 full backups	
	4 Replication	O Keep log backups applicable to last 14 days	
	5 Script	Cull backup retention activities	
	6 Verification	Full backup retention settings 0 Weekly	
	7 Summary	Total Snapshot copies to keep 7	
		O Keep Snapshot copies for 14 days	
			Previous
6	Configure the r	replication options	Previous Next
6.	Configure the r	replication options.	Previous Next
6.	Configure the r	replication options.	Previous Next
6.	Configure the r	replication options.	Previous
6.	Configure the r	replication options.	Previous
6.	Configure the r	replication options.	Previous
6.	Configure the r	replication options.	Previous
6.	Configure the r	replication options.	Previous
6.	Configure the r	replication options.	Previous
6.	Configure the r	replication options.	Previous
6.	Configure the r	replication options.	Previous
6.	Configure the r	replication options.	Previous
6.	Configure the r	replication options.	Previous

New SQL Serve	r Backup Policy	
1 Name	Select secondary replication options ()	
2 Backup Type	Update SnapMirror after creating a local Snapshot copy.	
3 Retention	Update SnapVault after creating a local Snapshot copy.	
4 Replication	Secondary policy label Choose ~	0
5 Script	Error retry count 3	
6 Verification		
7 Summary		
		Previous Next

7. Specify a run script to run before and after a backup job is run (if any).

Name	C							
Manne	Specify optional so	cripts to run b	etore perto	rming a back	up job			
Backup Type	Prescript full path							
Retention	Prescript arguments	Choose option	nal argument	s				
Replication	Specify optional so	cripts to run a	fter perforn	ning a backup	o job			
	Postscript full path							
Script	Postscript arguments	Choose option	nal argument	S				
Verification	Script timeout	60	secs					
Summary		-						
Juninary								
							Previous	Next

8. Run verification based on the backup schedule.

) Name	Select the options	to run back	up verificatior				
2 Backup Type	Run verifications f	or the follow	ing backup se	hedules			
3 Retention	Select how often you enabling you to stagg				ecific verification times	are set at backup job creation	
Replication	Weekly						
5 Script	Database consiste	ncy checks o	options				
6 Verification	 Limit the integrity Suppress all information 	l			HYSICAL_ONLY)		
7 Summary	 Display all reporte Do not check non Limit the checks a 	-clustered inde	exes (NOINDEX)		atabase Snapshot copy	(TABLOCK)	
	Verification script	settings					
	Script timeout	60	secs				
	Prescript full path Prescript	Choose optic	onal arguments				
	arguments Postscript full path			1.0			
	Postscript arguments	Choose optic	onal arguments	**		1	

9. The **Summary** page provides details of the backup policy. Any errors can be corrected here.

New SQL Serve	er Backup Policy		×
1 Name	Summary		Í
2 Backup Type	Policy name	TestDB-fullbackup-policy	
(3) Retention	Details	TestDB-fullbackup-policy	
Hetenbon	Backup type	Full backup	
Replication	Availability group settings	Backup only on preferred backup replica	
0	Schedule Type	Hourty 🛞 o	
S Script	UTM retention	Total backup copies to retain : 7	
6 Verification	Hourly Full backup retention	Total backup copies to retain : 7	
	Replication	none	
O Summary	Backup prescript settings	undefined Prescript arguments:	
			Previous

Configure and protect MSSQL Server database

1. Set up the starting date and expiration date of the backup policy.

let/	App SnapCenter®		•	a ()-	1 administrate	or SnapCe	nterAdmin	Sign Out	
1	Microsoft SQL Server 💽	Database - Protect Resource							×
	search by name					Dane Lifecycle	i) Details	F	
	Name Name								
	DEMODB1				5				
	DEMODB2	Resource Policies Verification	Notificatio		Summary				
8	DEMODB3	Televice Product Televice	The second second	340	sammary				
	INVENTORYDB1	Salast and a second set of a final second							1
	INVENTORYDB2	Select one or more policies@nd configure	Const In an						1
5	master	TestDB-fullbackup-policy -	+ 0						
	model								
8	msdb	Configure schedules for selected policies							
	tempdb	Policy IE Applied Schedules			ure Schedules				
	TestD6	TestDB- None fultbackup- policy		+	J				
							1		i
Create	ed backup policy TestD6-fullbackup-pol	Ŷ							×

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.

SnapCenter	× +					8	× -	(3	×
← → C ▲ Not	secure https://snapvm.dem	oaz.com:8146/inventory50	QLProtect/ProtectIndex				A A		4	1
		-								
		Add schedules f	or policy TestDB-fullba	ackup-policy	×					
		Hourly			1 and 1					
		Start date	11/18/2022 03:42 pm	8	- 100					
		Expires on	12/18/2022 03:42 pm	=	- 1000					
		Repeat every	6 C hours 0	mins						
			8 0							
		i The sche time zon	dules are triggered in the Snap e.	oCenter Server	×					
				Cancel	ok and					

3. Configure the verification server. From the drop- down menu, select the server.

4	e c	A Not secure https://snapvr	m.demoaz.com/8146/InventorySQLProte	ect/Protectindex					6 4		
n Ne	etApp	SnapCenter®			٠	⊠ 0 •	administrator	SnapCer	nterAdmin	🗊 Sign Out	
>	Micro	osoft SQL Server 👻	Database - Protect Resource								×
	sea	arch by name					a	ine Lifecycle	Details	Migrate	
U		Name					L. 100				
•		DEMODB1		0 0	4		5				
		DEMODB2	Resource	Policies Verifica	tion Notific	ation	Summary				
iii		DEMODB3	Resource	roncies vernica	ion roomo	DLISH.	Summiny				
4		INVENTORYDB1	6 J								4
h.		INVENTORYDB2		ation serves o							
		master	Verification server	SQLHMCRD5(14.0.1000)	· *					
華		model		SQLHMCRDS(14.0.1	000)						
▲		msdb	Configure verific	ation schedules							1
		tempdb	Policy JE So	chedule Type	Applied Schedu	ules	Configure	Schedules			
		TestOB	There is no match fo	or your search or data is no	t available.						
									Pres	rious Net	í

- 4. Confirm the configured schedule by clicking the plus sign and confirm.
- 5. Provide information for email notification. Click Next.

tAp	p SnapCenter®		0		0-	L clusteradmin	s SnapCe	nterAdmin	Sign Ou
	roxoft SQL Server 💽	Database - Protect Resource							
	earch by name							i	5
	Name	If you want to send notifications for scheduled or on demand jobs, an SMTP server must be configured. Continue to the summary page to save your information,	and then go to S	ettings	Global	1.4			
	master	⁴⁴⁵ server.							
	model								
	msdb	0-0-0-3							
	tempdb	Resource Polices Venfication Notification Summary							
	testdb								
	master	Provide email settings							
	model	Select the service accounts or people to notify regarding protection issues.							
	madb	Email preference Never +							
	ktroa	From From email							
	tempdb	To Email to							
		Subject Nonhiston							
		C Attach job report							
				Δ	ctiv	ate Wind	Inves		
Tota	al 10				or too	Settings to a	CHUNG	VALUE Pro	neus N

The summary of the backup policy for SQL Server database is now configured.

	tApp	o SnapCenter®						•		8-	1 administrator	SnapCen	terAdmin	Sign Out
	Mico	osoft SQL Server	Databas	se - Protect Resou	irce									
	se	arch by name										ne Lifecycle	i) Details	•
	19	Name									t be configured. Co		e summary	page to sav
		DEMODB1	you	ar information, an	id then go to se	congs>	Global Settings (Notification	Server 3	settings	o configure the SM	TP server.		
		DEMODB2									-			
		DEMODB3		0			-0-	-()—	_	- 5			
6		INVENTORYDB1		Resource	Policies		Verification	Notific	cation		summary			
		INVENTORYD82				\otimes	0							
		master		Database nam	e		TestDB							
8		model		Policy			TestDB-fullb	ackup polic	y: Hourl	y				
8		msdb		Verification ser	vers		SQLHMCRD	5						
		tempdb		Verification en	abled for policy		None							
		TestDB		Send email			No							

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**.

Net	Арр	SnapCenter®					2 0	- Ladministra	ator Snapo	CenterAdmin	🖉 Sign	Ou
	Micro	osoft SQL Server 🕒	"TestD8" Topology									
	\$01	srch by name		Magnetic Detailbase	Clone Lifecycle	Bernove Protection	Linck up New		Mertenarce	Details		tree
		Name	Manage Copies					hr				
		DEMODB1							2014-00			
		DEMODB2	0 Backu							iry Card		
		DEMODB3	Local copies	с».					0 Back 0 Clone			
		INVENTORYDB1	cotal copies		-				0 cioni	6		
		INVENTORYDB2		9	8 0							
		master										
		model	Primary Backup(s)									
		msdb	Csearch	7						1 1	- 41	
		tempdb	Conner							inner die		
		TestDB	Backup Name				Cou	int Type	17 E	ind Date	Verified	
			There is no match for your	search.								
	Total	10	Backups not found									

2. In the **Backup** dialog box, click **Backup**.

SnapCenter	× +				× :	0
← → C ▲ Not	t secure https://snapvm.do	Backup	for the selected resource TestDB	×	₽ \$	2
		Policy	TestDB-fullbackup-policy	· ·		
				Cancel Backup		

3. A confirmation screen is displayed. Click Yes to confirm.

Cont	firmation	×
	The policy selected for the on-demand backup is associated with a backup schedule and the on- demand backups will be retained based on the retention settings specified for the schedule type. Do you want to continue ?	
	Yes No	

Monitor backup job

1. From the **Monitor** tab, click the job and select **Details** on the right to view the jobs.

letApp SnapC Dashboard Resources	Jobs	Schedules	Events Logs			
	sear	ch by name				
Resources						websatilitys
	lobs -	Filter			J	
Monitor	ID	Status	Name	Start date	End date	Owner
Reports	1	-	Backup of Resource Group 'SQLHMCRDS_TestDB' with policy 'TestDB-fullbackup-policy'	11/18/2022 3:43:12 PM 🗃	11/18/2022 3:47:05 PM 🛱	Administrator
Hosts	12	4	Create Resource Group 'SQLHMCRDS_TestDB'	11/18/2022 3:42:32 PM 🛱	11/18/2022 3:42:55 PM 🛱	Administrator
Storage Systems	11	×.	Create Policy 'TestDB-fullbackup-policy'	11/18/2022 3:42:06 PM 🛱	11/18/2022 3:42:06 PM	Administrator
	10	*	Discover resources for host 'FCVM-RD52.Demoaz.com'	11/18/2022 3:40:27 PM	11/18/2022 3:40:29 PM	Administrator
Settings	9	~	Discover resources for host 'FSXCLUSTER.Demoaz.com'	11/18/2022 3:39:32 PM	11/18/2022 3:39:55 PM	Administrator
Alerts	8	4	Discover resources for host 'FCVM-RD51.Demoaz.com'	11/18/2022 3:39:13 PM 🗖	11/18/2022 3:39:15 PM 🗖	Administrator
			ackup of Resource Group 'SQLHMCRDS_TestDB' with p	- Andrewski - Andrewski - Andrewski - Andrewski - Andrewski - Andrewski - Andrewski - Andrewski - Andrewski - A	cy.	
			✓ ▼ fcym-rds1.demoaz.com			
			 Preparing for Backup 			
			 Creating SQL Backup 			
			✓ ► Finalizing Backup ⊗ 0			
			 Data Collection 			
			✓ ► Send EMS Messages			
		0	Task Name: Data Collection Start Time: 11/18/2022 3:47:05 PM Er	d Time: 11/18/2022 3:47:05 PM	•	
				View Logs Cancel jo	o ciffin	
					6	

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.

NetApp Snapt	Center®				• = •	Ladmonitoriae SnapCenterAdmon #Sig
	Monut SQ Sever					
Dechinged	ver Detablise • Detablise	(v)				transformation to the
Resources	U Pe Deblere	Instance	Hust	Last Backup	Overall Status	type
Monitor	Availability Group	SQURDSHIMC	PS/CUSTER.DemuA2.com 11	/11/2022 43425 PM 🗃	Backup posseeded	ster database
	0	SQLRDSHMC	PERCUSTER DemoA2.com		Not prunerbed	User database
Reports	nd mbes	SQURDSHIMC	PERCLUSTER DemoA2.com		Not protected	Wowr distabaue
Hats	Indicard para	SQLROSHIMC	PS/CLUSTIR.SemuA2.com		tiot printerned	ther database
Storage Systems	INSIGHTORS	SQLROSHINC	PSICLUSTOR DemokZuom		hiot protected	utar database
	mailler	SQURDSHIMC	PS/CLUSTUR.Demu42.com		Not protected	fysten detabase
	model	SOLROSHING	PSRCUUSTER.DemoAZ.com		Not protected	bystem database
Alerta	muttu	SQLRDSHINC	PSrCU7FTIR.DemuA2.com		Not protected	Bystem statabase
	tempdb.	SQURDSHINC	PERCUPTION Demok2.com		nor available for backup	System Gatabasa
	Olampiorshig081	stanta	STRUMS.Demo4Z.com(Q) O		Not protected	User databate
	Championichg/0610	\$14/M2	s184M3.Demo42.com		two protected	Oser databade
	Olumpionship2/8109	IDGMU	TTAME Demona com		Not protected	Ober itatabase
	(hamponship0810)	STAVM3	STAVAS DemoA2.com		Not protected	Liver statutulue
	ChampionshipOB102	steens	STRAMS.Demo42.com		hist protected	User database
	Of ampionship 09103	\$54/4/3	STRVM3.Dem64Z.com		NOT DISTRIBUTE	Over database
	(Trampionship038104	STALM3	STAVAD DemoA2.com		Not protected	the database
	ChampionshipOle105	STAVAD	STRVM3.Demo42.com		Not protected	User database
	ChampionsNpDB106	304/40	STRVM3.Dems42.com		Not protected	Ster database
	Championahp09187	SDAMD	STRAMS Demonit com		Not protected	utrer database
	Champorichg08108	STAWA3	1TX/M2.DemoA2.com		tiol printited	ster database
	Championt/hg/08109	5260903	STAVM3.DemoA2.com		Not protected	Uter database
	Champonshp0811	10440	stavas period com		Not protected	stow database
	championship08118	554/M3	STRIMS.DemoAZ.com		hot protected	User database
	Chamoion(heO@111	554900	115VM3.Demo42.com		tioc orsteized	
	100(313					

2. Click (+) for a new resource group.

					-
		(Protectindes/ResourceType = Database IIS-DI-Serverinstance = cullId-lo	attenul	8 \$ 0	
NetApp Sna				• E • tadministrator SnapCenterAdmin IS	Sign O
	Microsoft SQL Server				
Deshboard	Mew Resource Group Search by nume			8 ⁻¹¹	J
Resources	17 💌 Name	Resource Count Tags	Policies	Last Backup Overall Status	0
Monitor	There is no match for your search or data is not available.				
Reports					
Hosts					
Storage System					
E Settings					
Alers					
		8	0		
	Resources are not found. Click Refresh Resources to discover dat	abases in the database view or create new resource group on the disco	wered databases from the resource view.	Activate Windows Go to Settings to activate Windows	
tivity The S	most recent jobs are displayed	i Completed 🔒 0 Warnings 🔀 0 Fasled 🔗 0	0 Canceled 👩 0 Running 🌀 0 Queued		

3. Provide a name and tag. Click **Next**.

tApp SnapCenter®								• E O- Ladmanatrator	SnapCenterAdmin	USign Ox
Microsoft SQL Server	New Resource Group									
sausth by name	1 To configure an Da	UTP Server to send ema	il notifications for	scheduled or on-de	mand jobs, go to Settler	po Gladari Semengserini ficar	tion Server Settings			
7 Name				~						
There is no match for your search or data it not available.	1 Hame	Resources	Prices	Verification	Notification	6 Summery				
	Provide a	name and tags for	the resource g	roup						
	Name	RG-CHAMPIC	ONSHIPDB	20.5.10			0			
	Tegs	HG-CHAMPIC	ONSHIPDE				0			
	C) Use custo	om name format for Sila	apshot copy							
					⊗ ∘					
					-77 C - 173					
Resources are not found. Cick Refrects										
Resources are not found. Cick Refrech Resources to discover distabase view or create new resource group										

- 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.

pp SnapCenter®		🌒 🗃 🕢 - 土 administrator SnapCenterAdmin 🗊 Sign Di
Maronoft SQL Server 💽	New Resource Group	
search by name		
Name	0-0-0-0-0	
There is no match for your search or data s not available.	Name Resources Policies Venification Notification Summary	
	Add resources to Resource Group	
	Host Resource Type SQL Server Instance	
	STAVW3.DemoAZ.com • Databases • STAVW3 •	
	Available Resources Selected Resources	
	(search available resources Q	
	Auto select all the resources from the same storage volume ChampionybioDRT (STAVAR)	
	ChargiorthoBill (CAWA) ChargiorthoBill (CAWA)	
Rebources are not found. Click Refresh Resources to discover databases in the Jacabase view of create new resource group in the discovered databases from the resource view.		Activate Windows

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**.

n Net	App SnapCenter®		● ■ O- Lickasteradmin SnapCenterAdmin ISign Out
	Microsoft SQL Server		
> #	search by name	I it is recommended to move all databases on the same blorage volume at the same time.	5
0	T Name		
-	master		
0	model	Name Resources Polices Verification Notification Sommary	
ส่	mado		
Λ.	tempdb	Add resources to Resource Group	
	testób	Host Resource Type SQL Server Instance	
8 J	SalvDB	rdp-vm3.Demp2.com + Detabases + rdp-vm3 +	
ŧ	NorthwindDB	Available Resources Selected Resources	
Δ	DB1	(swarch available resources) Q	
-	DB10	Auto select all the resolution from the same soorage volume	
	082	NorthwindDB (rds-ym3)	3
	083	SaleDB (rdsvm3) DB10 (rdsvm3) restdb.rdsvm3) DB2 (rdsvm3)	
	D84	D83 (rds-vm3)	1
	085	> DB4 (rds-vm3) DB5 (rds-vm3)	
		DB6 (rds-vm3)	
	066	DB7 (rds-vm3) DB8 (rds-vm3)	
	067	089 (rds-vm3) D8100 (rds-vm3)	
	DBS	DB11 (rds.wm3)	
	089	DB12 (rds-vm3) DB13 (rds-vm3)	
	06100	DB14 (rds-wm3)	
	0611	DB35.(chemit)	*
	0812		
	0813		
	0914		
	0815		A
	DB19 Total 114		Activate Windows Go to Settings to activate Windows

5. On the policies, click (+).

NetApp SnapCenter®		🌒 🔳 🔮 🔹 🛓 ckasteradmin SnapCenterAdmin 🛛 Sign Out
Microsoft SQL Server	New Resource Group	×
III Cearch by name		
Name I Dellable Dellable Dellable Matter madbe Matter madbe NorthwindDill KtProble Saleo5 tempdb tempdb	Image: Construction Vertification Summary Select one or more policies and configure schedules Image: Configure schedules Image: Configure schedules Select one or more policies and configure schedules Image: Configure schedules Image: Configure schedules Configure schedules for selected policies Image: Configure schedules Image: Configure schedules Pare: Image: Configure schedules for selected policies Image: Configure schedules Image: Configure schedules There is no match for your search or data is not available Image: Configure schedules Image: Configure schedules	
testab Total 34	R	Activate Windows Go to Settings to activite Windows

6. Enter the resource group policy name.

Internation Internation	New SQL Serve	r Backup Policy	1	×.	
3. Metarino Betain 4. Medication: 5. Sirget: 5. Wethcaron: 7. Summary	1 Name	Provide a polic	y name		
A. Nellicolo S. Strut S. Wethorson J. Summary	2 Backup Type	Eolicy name	FULLBACKUP_CHAMPSONSHIPD8	0	
S strat € WethCation 2 Symmary	3. Autorition	Details	FULLBACKUP_CHAMPICIGH(#08		
(€) Writhouton (2) Symmary (2) Symmary (2) Symmary	4 Replication				
∑ Summary ⊗ ○	3 Script				
	6 Verification				
	7. Summary				
			Ø 0		
			0 0		
				Trema Prot	

7. Select Full backup and the schedule frequency depending on your company's SLA.

New SQL Serve	r Backup Policy ×	
	Select SQL server backup options	
2 Backup Type	Choose backup type	
3 Retention	C Full beckup and log backup	
(E) Replication	Full backup O Log backup	
5 Surget	Capy only backup	
	Maximum databases backed up per Snapshot copy: 100	
, F. Summary	Availability Group Settings 🗸 🗸	
	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 Mounty	
	O Daty O Weekly	
	O Monthly	
	Prevous Mf Im	
-		
	(4, Replication	Information for the backup and the backup in the backup is a strategy of the back

8. Configure the retention settings.

New SQL Serve	er Backup Policy	x
1 Name	Retention settings	
2 Backup Type	Retention settings for up-to-the-minute restore operation 🚯	
3 Retention	Keep log backups applicable to last 7 full backups	
4 Replication	O Keep log backups applicable to last	
5 Script		
6 Verification	Full backup retention settings 1 Weekly	
7 Summary	Total Snapshot copies to keep 7 Keep Snapshot copies for 14 days	
		Previous Next

9. Configure the replication options.

New SQL Serve	r Backup Policy			×
1 Name	Select secondary repl	ication options 🚯		
2 Backup Type	Update SnapMirror af	fter creating a local Snap	shot copy.	
3 Retention	🗍 Update SnapVault afte	er creating a local Snaps	hot copy.	
4 Replication	Secondary policy label	Choose	- 0	
5 Script	Error retry count	3		
6 Verification				
7 Summary				
				Previous Next

10. Configure the scripts to run before performing a backup. Click **Next**.

New SQL Serve	r Backup Policy				3
1 Name	Specify optional s	cripts to ru	before performing a backup job		
2 Backup Type	Prescript full path				
3 Retention	Prescript arguments	Choose op	ional arguments		
4 Replication	Specify optional s	cripts to ru	after performing a backup job		
5 Script	Postscript full path				
	Postscript arguments	Choose op	ional arguments		
6 Verification	Script timeout	60	secs		
7 Summary					
				Previous	Next
				Flevious	Next

11. Confirm the verification for the following backup schedules.

Name	Select the options	to run backup verification				
Backup Type		or the following backup schedules				
Retention	Select how often you		The specific verification times are set at backup job creation			
Replication	Hourly					
Script	Database consiste	ncy checks options				
5 Verification	🗹 Limit the integrity	structure to physical structure of the data	base (PHYSICAL_ONLY)			
- Termedian	Suppress all information message (NO_INFOMSGS)					
7 Summary	Display all report	rted error messages per object (ALL_ERRORMSGS)				
	Do not check non-clustered indexes (NOINDEX)					
	Verification script	⊗ ⊂	ernal database Snapshot copy (TABLOCK))			
	Script timeout	60 Secs				
	Prescript full path	<scripts_path></scripts_path>				
	Prescript arguments	Choose optional arguments				
	Postscript full path	<scripts_path></scripts_path>				
	Postscript arguments	Choose optional arguments				

12. On the **Summary** page, verify the information, and click **Finish**.

Name	Summary		
Backup Type	Policy name	FULLBACKUP_CHAMPIONSHIPDB	
Retention	Details	FULLBACKUP_CHAMPIONSHIPDB	
Recention	Backup type	Full backup	
Replication	Availability group settings	Backup only on preferred backup replica	
-	Schedule Type	Hourly	
Script	UTM retention	Total backup copies to retain : 7	
Verification	Hourly Full backup retention	Total backup copies to retain : 7	
	Replication	none	
Summary	Backup prescript settings	undefined Prescript arguments:	
	Backup postscript settings	undefine O Postscript arguments:	
	Verification for backup schedule type	Hourty	
	Verification prescript settings	undefined Prescript arguments:	
	Verification postscript settings	undefined Postscript arguments:	

Configure and protect multiple SQL Server databases

1. Click the (+) sign to configure the start date and the expire- on date.

()	C A Not secure Hitps://snapym.de	moaz.com 8144/meentorySQLProtect/Protectinder/ResourceType+Datatased/SQLServerInstance+cv288Host+nu8#				년 ☆	0 🔅 i
II Ne	tApp SnapCenter®		•	. 0.	1 administrator	SnapCenterAdmin	Sign Out
5	Microsoft SQL Server	New Resource Group					>
	search by name						
0	Name						
۲	There is no match for your search or data is not available.	Name Rebourses Policies Verification Notification Summary					
-		Select one or more policies and configure schedules					
4		PULLBACKUP, CHAMPONSHIPDB					
30		Configure schedules for selected policies					
		Policy II Applied Schedules Configure Schedules					
		илланоми, снаменованов нане Ди					
92. – E		× •					
		Total I					
		D Use Microsoft SQL Server scheduler 0					

2. Set the time.

Add schedules fo	r policy FULLBACKUP	_CHAMPIONSHIPDB ×
Hourly		
Start date	11/11/2022 05:30 pm	
Expires on	12/11/2022 05:27 pm	
Repeat every	a 🌲 hours o	mins
i The schede zone.	ules are triggered She Snap	Center Server time
		Cancel

I Ne	tApp SnapCenter®		🔹 🗃 🥥 • 🔒 administrator ShapCenterAdmin 🗍 Sign Out
	Microsoft SQL Server	New Ilmource Group	
	Salech by Hime		
,	7 Name	0-0-0-0-0-0	
,	There is no match for your search or data in not available.	Name Resources Policies Ventication Notification Summary	
1	COMPARING THE	- Select one or more policies and configure schedules	
		КULLBACKUP_СНАМРЮНИНГОВ - + 0	
		Configure schedules for selected policies Policy Is Applied Schedules Configure Schedules	
Δ		Plikt BACKUP_CHAMPIONSHPDB Houry: Appear every 8 hours:	
-			
		8 0	
		Total I	
		Use Microsoft SQC Server scheduler	
	Resources are not found. Click Refresh Resources to discover databases in the		5
	database view or create new resource group on the discovered databases from the		3
	Arsource view.		Activate Windows Previous Ne
tivi	ty The 5 most recent jobs are displayed	⊘ 3 Completed 🔒 0 Warnings 🔗 D Failed 🔗 D Canceled 💿 0 Burnsing 🔕 0 Gueund	the second of second determined

3. From the **Verification** tab, select the server, configure the schedule, and click **Next**.

	etApp SnapCenter®					• = 6	• Ladministrator SnapCenterA	dmin 🛛 Sign Ou
\$	Microsoft SQL Server	New Resource Group						
	search by name		5 mg	-				
0	7 Name	0-0-0-	-0-		۲			
0	There is no match for your search or data is not available.	Name Resources Policies	Venfcation	Notification Su	mmary			
ai i	- CALMAN SOL	Select the verification servers						
		Verification server STAVM3(13.0.5026)	•					
1		SQUEDSHMC(13.0.5020)						
*			chedule Type	Applied Schedules	Configure Schedules			
▲		FULLBACKUP_CHAMPIONSHIPOB	oury	hone	*			
		Total I						
	Resources are not found, Click Refresh Resources to discover databases in the database www or create new resource group on the discovered databases from the resource view.							Previous No

4. Configure notifications to send an email.

App SnapCenter®		🔹 🗃 🕘 - 1 administrator: SnapCenterAdmin 🗷 Sign Du
Microsoft SQL Server	New Resource Group	
search by name	4 If you want to send notifications for scheduled or on demand jobs, an SATTP server must be configured. Continue to the Summary page to save your infor server.	rmation, and then go to Settings-Global Settings-Hotthcation Server Settings to configure the f
T Name		
There is no match for your search or data is not available.	Latre Resources Polices Verification Notification Summary	
	Provide email settings o select the service accounts or people to notify regarding protection issues.	
	Email preference Never •	
	From From email	
	fo Insulto	
	Subject toohcarion O	

The policy is now configured for backing up multiple SQL Server databases.

		🗧 🛤 😝 🛨 administrator SnapCenterAdmini 🛙 Sij
Microsoft SQL Server		
search by name	If you want to send nutifications for scheduled or on demend jobs, an SATP server must be configured. Commune to the Summary page to knew your server.	ir information, and then go to Settings-Global Settings-Notification Server Settings to configure
R Name		
There is no match for your search or data is not available.	Nathe Resources Policies Wentcastern Rostification Summary	
	Resource group name RE-CHAMPORTUPOII	
	Tags RE-DIAMPORSHIPOR	
	Policy POLIDACION_COMPORTATION: Neurly	
	Plugen ShapCenter Plug in for Microsoft SQL Senter	
	Verification Server STRUKU Verification enabled for policy RULEBACKUP CHAMPIONSHIPDE Heavily	
	Send entail to & 0	

Trigger on-demand backup for multiple SQL Server databases

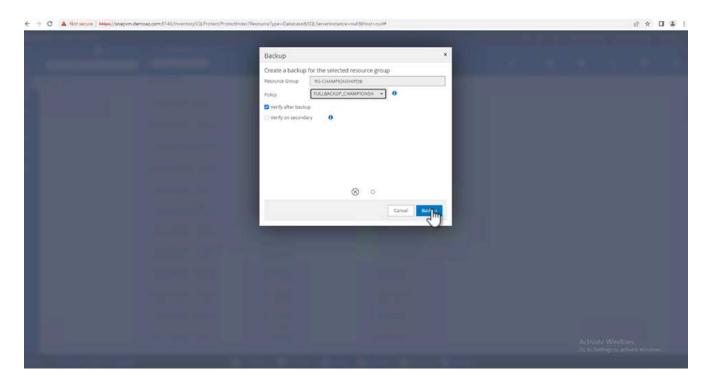
1. From the **Resource** tab, select view. From the drop-down menu, select **Resource Group**.

App SnapCenter®					0- 1=	ministratur SnapCo	enten Admin	SgaDut
Michael SQL Server	NG CHAMPORSHIPDE Details							
search by risting	Smith -				<u> </u>		-	<u> </u>
74 Name	Resource Name	Туре	Host	and a second second second second second second second second second second second second second second second	վեղ			
RG-DHAMPONSHPO8	Championuhg081 (STAVA3)	SQL Detailate	STAND. Demok2.com		C			
	Championalsp0810 (574/40)	SQL Detailure	\$369M3.Demok2.com					
	ChampionshipD(8100((328/463)	SOL Database	\$78yM3.DemoA2.com					
	Ownpionalsp201101 (\$144143)	SQL Dataliana	STKM/J.DemokZ.com					
	ChampionshipOB102 (STWM1)	SQL Datatiana	STAVM3.DampAZ.com					
	Championitsp08103 (554/463)	SQL Database	118xM3.DemoA2.com					
	ChampionshipDB104(STR/M3)	SQL DataBasia	\$3WM3.Demok2.com					
	ChampionshipDB105 (ET#/M3)	SQL Database	\$78VM3.DemoA2.com					
	Champioralsp08106 (578/WI)	SQL Database	\$78xMd1Demok2.com					
	ChampsonshipD3107 (ClavM3)	SQL Database 🛞	33MMd3.Demok2.com					
	ChampionshipD8108 (STM/H/3)	SQL Database	\$249MLDsmo42.com					
	ChampionalhipDB109(ST#44/3)	SQL Database	\$78yM3.Demok2.com					
	Championahip0811(37A/M3)	SQL Database	\$749443.Demo42.com					
	Championahip08110 (CTAVMS)	SQL Database	\$350M3.DamoA2.com					
	ChampionalvgDB111 (CTx4W3)	SQL Database	\$7xxM3.Demol/2.com					
	Championship08112 (STAin43)	SQL Database	\$749343.Damo42.com					
	Champion(h)p08113 (374/H2)	SQL Database	\$10483.Demok2.com					
	Champsonihydd8114 (STAVM3)	SQL Dataliate	17XVM3.Demo42.com					
	ChampionshipDi8115 (\$154Md)	SOL Detailuse	STAIM3 DemokZ zom					
	Championship08116 (27#/943)	SQL Datatiana	\$300M3.Demoit2.com					
	Champion/hip28117 (STAVM3)	SOL Database	\$7AVM3.DemoA2.com					
	ChampionshipDE118 (\$154543)	NOL Calabase	STRIARS DamokZ zom					
	Olumpionalve/DB119 (STAMA9)	SQL Database	stividd Demok2.com					
Total 1	Overnaminalsis/0812137AvM31	SOL Database	\$13/MJ.Damo42.com					

- 2. Select the resource group name.
- 3. Click **Backup now** in the upper right.

	onnell SQL Server 🔒		NG-1 Details						
0	search by name	\supset	(stan)			 	Destinate	<u>×</u> <u>0</u>	
	Name		Resource Name	Туре	Hoat	 Chert and			
	mader		001 (Hds-wei0)	SQL Database	nts-smit DemoZ.com				
	model	- 1	DB10 (rds-emil)	SQL Database	rdu-vm3.Demo2.com				
	maste	- 1	082 (rds-uni3)	5QL Database	rds-um3.Demo2.com				
	teropdo	- 1	003(101-003)	SQL Database	nts-well.DemsZ.com				
-	Metab	- 1	0.64 (rds-ent(3)	SQL Datatiase	rds-ym3.Demo2.com				
	Select	- 6	085(00-400)	SQL Database	rds-vm3.0wmo2.com				
-	NorthwestDB		001041-000	SQL Datatiase	rds-emb Demož.com				
	081		0407 (rds-em/3)	SQL Database	nts-vm3.DemoZ.com				
-	0618		Dist (rds-ored)	5QL Dauttase	hts-west-barried.com				
	082		2093 (Hds-emb)	SQL Datatiase	rds-sm3.Demo2.com				
-	083		0/01/00 (nds-vm3),	SQL Database	rds-vm3.0emoZ.com				
	064		DB11(down3)	SQL Database	nts vin3.DemoZ.com				
	085		0612 (rd) vm3)	SQL Database	nts web DemoZ.com				
	066		D013 (rds-emd)	SQL Database	nts-vm3.Demo2.com				
	047		0614 (rds-emilij	SQL DAVABLEY	rds-em3.Demo2.com				
	DBB		0615 (rd)-em3)	SQL Database	nts-sm3.0emoZ.com				
-	001		0014(10)((0.0))	5QL Database	rds-ym3.0emo2.com				
	08100		DB17 mbi vmity	SQL Database	wb.vm3.DemoZ.com				
	0811		0818 (rds-end)	SQL Database	rds-und-Demo2.com				
	0813		0019 (rds-end)	SQL Database	nty-vm3.2emo2.com				
	0813		0020 (rds-emil)	SQL Database	rds-im3.2emo2.com				
	0814		0821 ads-writ3)	SQL Database	nts-emil.Demo2.com				
	0015		0022(Hts-em3)	SQL Database	nts-en3.0emoZ.com				
	DB16		D623 ods-em/li	SOL Database	nts-im-3.5emoZ.com	Activ	rate Wi	ndows	

4. A new window opens. Click the **Verify after backup** checkbox and then click backup.



5. A confirmation message is dsiplayed. Click Yes.

Conf	irmation	×
	The policy selected for the on-demand backup is associated with a backup schedule and the on- demand backups will be retained based on the retention settings specified for the schedule type. Do you want to continue ?	
	Yes No	

Monitor multiple-database backup jobs

From the left navigation bar, click **Monitor**, select the backup job, and click **Details** to view job progress.

C A Not secure Helper/Anaport demoac.com/#140/http:		교 호
	and the second second second second second second second second second second second second second second second	
	job Details *	
	Backup of Resource Group: RG CHAMPIONSHIPDB with policy FULLBACKUP_CHAMPIONSHIPDIF	
	0 • Takisp of Nessana Group 16-CHAMPORGHIDTE web party TULLIACE P. CHAMPONGHIDTE	
	✓ + SIMMEDemAZeam	
	 (3)0.361/Werty Sacoup PG-DWW/PGH3HPDBLSDWMQ_11-H-2022_1128.51-4332 of 10 ddts TstWM2/Dumpon/WpDB185 	
	(pp. 31) (verty labor) = 0.0 CAMPOREDED (1996), (1-11-2)22, (7.2).81 A322 of 100 Ox80(3) 376403/Champoning00132, (376403/Championing00177), (376403/Championing0017) (376403/Championing0012)	
	(1)(5) 25 yourly backed with Counterposed and the counterposed of	
	⊗ ∘	
	O An de verflasion server is surrently universitable, the job is governil.	
	Weer Logs Thread () Cone	
	and the second se	

Click the **Resource** tab to see the time it takes for the backup to be completed.

tApp Snap	Center®							🕴 🗃 😔 🛨 admin	strator SnapCenterAdmin	U Sign
	Microsoft S	QLServer •								
Dashboard	View Re	source Group • search by name	V							
Resources	17.14	Name	Resource	Count	Tags	Po	licles	Last Backup	Overall Status	
Monitor		RS-CHAMPIONSHIPDII		200	RS-OHAMPONSHFDB	H.	LLEACKUP, CHAMPIOLESHIPDIE	11/11/2022 5:01:05 PM 🛱	Queued	
Reports								C		
Hosts										
itorage Systems										
emes										
lerts.										
					× •					
	Total 1								vate Windows Settings to activate Winds	

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**.

tApp SnapCer	ter®					Lasimentitudur ScapContecAdmin Sign
	Noronalt SQL Server					
-	Decadada - Decadada	na 🗸				
aurus' 1	Detatore Instance	Instance	Hunt	Last Backup	Overall Status	Туре
water	Austability Group	SQUADSHINC	PSRCLOSTER Demo42.com	11/11/2022 43#25 PM 🖨	Beckup succeeded	Citier database
(112)	0	3QUADSHIMC	PSXCU/STER Demo42.com		Not protected	User database
10155	man and the	SQURDSHMR:	PSRUISTIK2enoA2.com		Not protected	Uner datalitane
-	NSCHOOL	SOLKOSHIAC	FEXCLOSTER.Demo42.com		Not protected	User database
rage Systems	INSIGHTORS	SQLADSHUC	PSRCLUSTER.Demok2.com		Not protected	the database
	matter	SQLROSHIAC	PERCEIPTER.Demo42.com		Not protected	System database
inte a	hoom	SQUADSHIMC	PSRCLUSTER.DemoA2.com		teat protected	System database
90)	muth	SQUADSHIAC	PINCLUSTER Demokit.com		Not protected	System database
	tempdb	SQURDSHIMC	PIRQUISTER Demok2.com		Not evalable for backup	System database
	Otempionship081	\$74/943	Stanta DeniA2.com	15/11/2022 5:30:25 PM 🛱	Backup succeeded	User database
	Championihqi0010	(0.640)	USAVM3.DemoA2.com	11/11/2022 5:20:26 PM	Bachup succeeded	Use/ database
	Championship/DB100	TOOM3	STAVM3.DemoA2.com	11/11/2022 5:30:25 PM 🛱	Bachup succeeded	Stor database
	Champomhp00131	Staunds	STAVM3.DemoA2.com	11/11/2022 5:30:26 PM 🖨	Balkup succeeded	thei dataltate
	Changiomhp36102	STAVAD	STAVM3.DemoA2.com	11/11/2022 5:30:28 PM 🛱	Backup succeeded	Uter database
	Champiorship08103	354403	STAVM3.DemoA2.com	11/11/2022 5:30:26 PM 🛱	Backup successful	User database
	Champoninhp08154	134/43	514MLDevoAZ.com	11/11/2022 5:30:26 PM 🗖	Backup turneeded	Uper database
	champ-onthipDB105	STAV43	STAVALDEMOAZ.com	11/11/2022 5:00:26 PM 🗇	Backup succeeded	titler datafiane
	Championship/DB104	156563	STAVM3.DemoA2.com	11/11/2022 5:30:25 PM 🛱	Beitup succeeded	User database
	Championship(08107	ST#W0	STAVMLDemoA2.com	11/11/2022 5:30:26 PM 🗖	Baciup succeeded	Uter database
	Championship08108	514/40	STAVA3.DemdA2.com	10102622 530.26 PM 🛱	Backup succeeded	Vier database
	Championship20109	124440	(DAVM3.DwmuA2.com	11/11/2022 5:30:26 PM	Being succeeded	User decebase
	Championship:0011	106/40	[TAVM].DemoAZ.com	11/11/2022 5:30:26 PM	Backup successfed	Oser database
	Champomhp0011E	STAUMS	STAVM3.DemoA2.com	11/11/2022 500/28 PM 🛱	Ballup scounded	Uner database
	Outcomto00111	104/40	STAVM3.DemoA2.com	13/11/2022 5:30:26 PM 🛱	Beckus succeeded	Access Mer detabase
1	wai 213					

- 2. Select the resource group backup policy created.
- 3. Select Modify Resource Group in the upper right.

etApp SnapCenter®					0. Tep	matrian Snap	CenterAdmin	150
Microsoft SQL Server	RG-CIMMPONERFORDER/IN							
inerth by form	Seet.			1		. 🔺	0	
F Name	Resource Name	Туре	Hust.	and the second second second second second second second second second second second second second second second	A DUM CO POTATO			
NG CHANTCHENTON	ChampionshipDB1 d3N4M3s	SQL Database	TAVM3.0emoAZ.com					
	Champonship0810 (STXVM2)	SQL Catabase	STAVAS SemiAZ com					
	Championich(pDIE1023STAVM3)	SQL Database	11AvM3.DemoA2.com					
	Championuhig0(8101 (STAVM3)	SQL Database	1TAVM3.DemoAZ.com					
	ChamponshipOlitit2 (STAyM3)	SQL Deceptere	TAVMJ.DemoAZ.com					
	Championship0(8103/STAvM3)	SQL Database	\$74/M3.Demo42.com					
	Championship08104 (STAVK3)	SQL Database	1fAVM3.DemuAZ.com					
	Champonship0@101(S7Xv90)	SQL Datablete	TRAVMS DemoAZ.com					
	Championuhip0@105(STAVM3)	SQL DaGidase	\$74xM3.DemoA2.com					
	ChampionshipDi8107 (STAVM3)	SQL Database	diwwi.0enok2.com					
	Championshipole (18 (572/383)	SQL Detabase	TTAVMD.DemoAZ.com					
	Championuhip0@103.6TAVM3)	SQL Database	114/43.Demo42.com					
	ChampionshipDIE11 (57K/M3)	SQL Database	11AVM3.DemoAZ.com					
	Champonuhg@8110(\$1xv93)	SQL Ownativate	STAVM3-DemoAZ.com					
	Championuhip0@111 (STAVM3)	SQL Database	11A/M3.DemiA2.com					
	ChampionshipDill112 (STAVM3)	SQL Database	11AVM3.DemoA2.com					
	Overgionshipbili112 (STWW2)	SQS OWGROUP	\$fAVM3.0emoA2.com					
	Championship0(8114(574vM3)	SQL Database	STAVAS DemoA2.com					
	Championship2(8155451AVM3)	SQL Detablate	11AVM3.0emoA2.com					
	Championship0@116.csTxv3r3)	SQL Database	STAVES Common Star Star Star Star Star Star Star Star					
	Champeniship0(8117.7 (STAVM3)	SQL Database	11AvM3.DemoAJ.com					
	Championul/sp0iii113 d514VM31	SQL Database	114VM3.0emok2.cz#					
	Oumponship@lt111GTAVM31	SQL Detetate	STAVM2.0+maA2.com					
704/1	Championshie@iii12751AvM3i	SOL Database	17AVM3.DemoA2.com					

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.

Verach by name Name </th <th>NetApp SnapCenter®</th> <th></th> <th>🔹 🗃 🥹 🛨 administratur StapCenterAdmin 📲 Sign Out</th>	NetApp SnapCenter®		🔹 🗃 🥹 🛨 administratur StapCenterAdmin 📲 Sign Out
verch by same verch by s	Microsoft SQL Server	Wodfy lineaute Group	\$
It Name Resources Fulces Winflaction Summary Add resources to Resource Group Name Resources Selected Resources Available Resources Selected Resources Selected Resources Sourpoint/DBID GTXMAB Ourmoont/DBID GTXMAB Ourmoont/DBID GTXMAB Ourmoont/DBID GTXMAB Ourmoont/DBID GTXMAB <	search by name		
Nume Resources Pulces Winflaction Summary Add resources to Resource Group Notification Selected Resources Selected Resources Available Resources Selected Resources Selected Resources Selected Resources Status select all the resources for the same strage volume Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Ourmoon/bpDBI GTAVAB) Our		- 0 - 0 - 0 - 0 - 0	
Add resources to Resource Group Hot Resource Type All Instances Selected Resources Selected Resources Selected Instances Selected Resources Sourpoint/politi (StrvMa) Oumpoint/politi (StrvMa) Oumpoint/politi (StrvMa) Oumpoint/politi (StrvMa)	RECHARGED BARD	Name Resources Policies Verification Notification Summary	
Hots: Reductrik Type All Instances Available Resources Selected Resources Secrit available resources Selected Resources Secrit available resources ChamponityDBI (StAVAB) OumponityDBI (StAVAB) OumponityDBI (StAVAB) OumponityDBI (StAVA		Add resources to Resource Group	
Al Instances Available Resources Selected Resources Selected Resources Selected Resources ComponshipEdBell (Stawall) ChemponshipEdBell (Stawall) ChemponshipEdBell (Stawall) Selected Resources </td <td></td> <td>Host Resource Type</td> <td></td>		Host Resource Type	
Available Resources Selected Resources Selected Resources Selected Resources Selected Resources Componsigned resources Componsi		A8 • instances •	
Auto selicit all the resources from the same storage volume Sources from the same storage volume from the same storage volume Sources from the same storage volume Sources from the same storage volume Sources from the same storage volume Sources from the same storage volume Sources from the same storage volume Sources from the same storage volume Sources from the same storage volume Sources from the same storage volume Sources from the same storage volume Sources from the same storage volume Sources from the same		Available Resources Selected Resources	
SQLBDJANKC ChamponshuD01 (STAVA3) - STANAS ChamponshuD01 (STAVA3) - ChamponshuD01 (STAVA3) - - Champons		(search available resources)	
:51V/N3 Oumpoint/bg0810 (51V/N3) Oumpoint/bg0810 (21V/N3) Outpoint/bg0810 (21V/N3)			
		Champonship0610 (STAMA3) Champonship0610 (STAMA3) Champonship0101 (STAMA3) Champonshi	

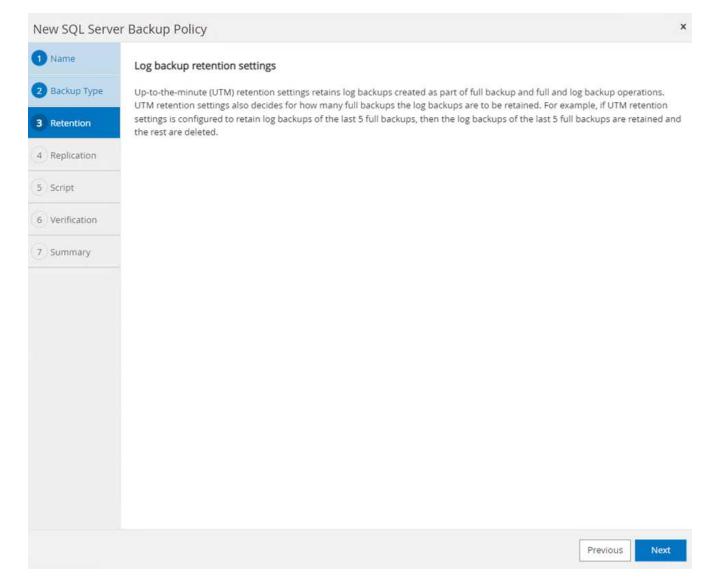
5. Enter the policy name.

And the second se	eaz.com/3146/InventorySQUProtect/Protect				18 ¥
	New SQL Serve	r Backup Polic	6	*	
	(1) Name	Provide a polic	y name		
	2 Backup type	Policy name	LOGBACHUP_CHAMPIONSHIPDB	0	
	3 Retention	Ortails	LOSBACKUP, CHAMPORSHIPDB		
	4 Replication				
	(s.) script				
	(6) Verification				
	2 Summary				
			⊗ ∘		
			0.0		
				President Next	
		_		President Next	
				13	

- 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 Serve	er Backup Policy
1 Name	Select SQL server backup options
2 Backup Type	Choose backup type
3 Retention	O Full backup and log backup
4 Replication	 Full backup Log backup
5 Script	Copy only backup 0
6 Verification	Maximum databases backed up per Snapshot copy: 100
7 Summary	Availability Group Settings
	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
	O Daily O Weekly
	○ Monthly
	Previous Next

9. Configure the log backup retention settings. Click Next.



10. (Optional) Configure the replication options.

534

A Not secure Mitais://snapvm.	amoaz.com/1146/InventorySQLPictect/ProtectIndex/ResourceType=nwIthSQLServerInstance=nw8bHost=ex/8*	id \$
	New SQL Server Backup Policy	*
	Select secondary replication options 0	
	Indixed Type Drugdate Snaphtvror after creating a local Snapshot copy	
	Retainetion Update SnapVault after creating a local Snapuhot copy.	
	A Replication Secondary policy label	
	Error retry count. 3 0	
	(6) WestMation	
	7 Summary	
	⊗ ∘	
		Previous.
		Activate Windows Grito Settispile activity Week

11. (Optional) Configure any scripts to run before performing a backup job.

constrained and a second second second second second second second second second second second second second s	condex resource (yee *)	suit&SQLServerInstance+mult&Host+mult#		18 A
The subscription of the local division of the local division of the local division of the local division of the				
New SQL Serve	er Backup Policy		×	
1 tame	Specify optional s	cripts to run before performing a backup job		
O factor Type	Prescript full path	<scripts_fath></scripts_fath>		
C Ratincon	Prescript arguments	Choose optional arguments		
O Replication	Specify optional s	cripts to run after performing a backup job 0		
5 Script	Postscript full path	<scrp15_path></scrp15_path>		
a vertication	Postscript arguments	Choose opponal arguments		
	Script briegout	éo zecs		
7 Summary				
		× •		
			Previous Notice	
	_			

12. (Optional) Configure backup verfication.

🗧 🕘 C 🔺 Not secure | Hillipe//tanapom.demosz.com/0146/teventorySQLProtect/ProtectIndee/Resource?ype=nu885QLServerinetance=nu883foct=nu8#

the second second second second second second second second second second second second second second second se		
New SQL Se	rver Backup Policy *	
O have	Select the options to run backup verification	
🕤 Barchap Typ	Log backup	
Retartion	Wirlying backup.	
O Parliador		
0		
6 Verification		
7. Summary		
	⊗ ∘	
	Pressue	

13. On the **Summary** page, click **Finish**.

New SQL Serve	er Backup Policy		*	
1 tame	Summary			
O Backup Type	Policy name	LOGEACIOUP, CHAMPIONSHPOIL		
	Oetaits	LOGBACKOP_CHAMPIONSHIPDB		
O Retenciols	Backup type	Log transaction backup		
O Replaced	Availability group settings	Backup only on preferred backup replica		
	Schedule type	Howry.		
Script.	Replication	(NUTH)		
Verification	Backup prescript settings	undefined Procept arguments		
7 Summary	Backup postscript settings	undefined Postscript arguments:		
	Verification for backup schedule type	····* ⊗ o		
	Venfication prescript settings	Undefined Prescript arguments		
	Venfication postscript settings	underforeid. Postscript enguments:		
			Previous Frees	
			Pressous and In	

Configure and protect multiple MSSQL Server databases

1. Click the newly created transaction log backup policy.

tApp SnapCenter®		●
Mcrosoft SQL Server	Modify Resource Group	
aearch by name		
Marne Name		
AC-CHAMPIONSHIPOR	Name Resources Policies Verification Notification Summary	
	Select one or more policies and configure schedules	
	тицьяскил_снамловылов, Logваскі •	
	Configure schedules for selected policies	
	Policy IL Applied Schedules Configure Schedules	
	RULLIACKUP_CHAMINONSHIPDE Hourly: Reprat every 8 hours	
	Tent J	
	Use Microsoft SQL Server scheduler 0	
		& theate Woodpart

- 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.

SQLProtect/Protectindes/PersourceType=multitisQLServerInstance=multitificat=mult#	e \$ 0
Add schedules for policy LOGBACKUP_CHAMPIONSHIPDB *	
Hourty	
Start date 11/11/2022 07:30 gm	
Explores on 12/11/2022 06:30 pm	
Repeat every 2 \$ bours 0 mms	
1 The schedules are triggered What SngoCarter Server time	
2016	
2006	
2016	
2016	
2016	
2016	
2016	
2016	

4. You can see both policies. Click **Next**.

Net/	App SnapCenter®		🌒 🗃 🥹 🛨 administrator SnapCenterAdmin 🛛 Sign Out
	Microsoft SQL Server	Modily Insource (insup	
	search by name		
12	T Name		
0	RE-CHAMPIONSHIPDO	tiame Resources Policies Verification Summary	
a		Select one or more policies and configure schedules	
		PULLBACKUP_CHAMPIONSHIPDB_LOGBACK+	
E		Configure schedules for selected policies Policy IE Applied Schedules Configure Schedules	
Δ.		FULLBACKUP_CHAMPONSHIPDE Hourty: Repeat every 8 hours / x	
111		LOGBACKUP_CHAMPIONSHIPOB Hourly: Repliat every 2 hours 🖌 🗴	
		⊗ ∘	
		Total J	
		Use Microsoft SQL Server scheduler	
			Activate Windows

5. Configure the verification server.

	pern.demoaz.com/ll146/thventory/SQLProtect/ProtectIndex/Resource/spe=nu/RSQLServerbratance=nu/R8/Heat+nu/#	18 🖈 🗌 🛓 1
NetApp SnapCenter® Microsoft SQL Server	Mindally Messaurize Circlup	I I vidministrator SnapCenterAdmin P sign Out
a second and a second s		
Search by name		
Name	Name Resources Policies Wellfication Notification Summary	
IIG CHAMPICII GHIPOB	THE TRANSPORT TO A TRANSPORT CONTRACT (ACCOUNT)	
4	Select the verification servers	
	Verification server TEN/MU(13.0.5026) +	
k • ≡		
	Configure verification schedules	
	Policy II Schedule Type Applied Schedules Configure Schedules	
<u> </u>	PULLBACKUP, CHAMPIONSHIPD8 Hourly Venity after backup	
	⊗ °	
	Total 1	
Total 1		Activate Windows
rthebu	erland	٢.

6. Configure email notification.

II N	etApp SnapCenter®		🌲 🗃 😜 🛓 clusteradmin SnapCenterAdmin 🖉 Sign Out
>	Microsoft SQL Server	Modify Resource Group	×
	search by name	If you want to send notifications for scheduled or on demand jobs, an SMTP server must be configured. Continue to the Summary page to save your i server.	nformation, and then go to Settings-Global Settings-Notification Server Settings to configure the SATP
0	Mame		
ତ ଲା	1 1894	Name Resources Polices Verification Notification Summary	
A 34		Provide email settings 0 Select the server account or people to notify regarding protection mours.	
=		Email preference toever •	
A		From International	
		Ta Email to	
		Subject Notification	
		C Attach job report	
	Total 1		Activate Windows

7. On the **Summary** page, click **Finish**.

Macrosoft SQL Server Modely Heacured Group Learch by name • Single	R152: Server Modely Hencarce Group Projerowska Fyree work to seed noofhulations for scheduled or on demand jobs, with SHTP server must be configured. Control-to the Skermary page to save your information, and then go to Skettings-Solidad Stettings-Hotefulation Server Settings to configure the Skermary page to save your information, and then go to Skettings-Solidad Stettings-Hotefulation Server Settings to configure the Skermary page to save your information, and then go to Skettings-Hotefulation Server Settings to configure the Skermary page to save your information, and then go to Skettings-Hotefulation Server Settings to configure the Skermary page to save your information, and then go to Skettings-Hotefulation Server Settings to configure the Skermary page to save your information. Name	tApp SnapCenter®			🏮 🖬 🥶 💄 administrator SnapCenterAdr	min ØSgn
Name Image: CARAMINODSSWORD I	Inter Image: Severe: Inter Image: Severe: Image: Severe: Seve		Modify Resource Group			
Neil CHAMINGUESHIPOR Image: Polices Yellices Yellices Yellices Summary Resource group name RE-CHAMINGUESHIPOS Summary Resource group name RE-CHAMINGUESHIPOS Tags NG-CHAMINGUESHIPOS NG-CHAMINGUESHIPOS Resource group name RE-CHAMINGUESHIPOS Poly RES-CHAMINGUESHIPOS RES-CHAMINGUESHIPOS RES-CHAMINGUESHIPOS Resource group name Poly RES-CHAMINGUESHIPOS Resource group name RES-CHAMINGUESHIPOS Resource group name Poly RES-CHAMINGUESHIPOS Resource group name RES-CHAMINGUESHIPOS Resource group name Poly RES-CHAMINGUESHIPOS Resource group name RES-CHAMINGUESHIPOS Resource group name Poly RES-CHAMINGUESHIPOS Resource group name RES-CHAMINGUESHIPOS Resource group name Poly Resource group name RES-CHAMINGUESHIPOS Resource group name Resource group name Poly Resource group name Resource group name Resource group name Resource group name Poly Resource group name Resource group name Resource group name Resource group name Resource group name Resource gro	CAMPOOLSHIPDE	search by reame		d or on demand jobs, an SMTP server must be configured. Continue to the Summary page to save your m	formation, and then go to Settings-Global Settings-Notification Server Settings	to configure th
Name Resources Pelification Notification Summary Resource group name RIS-OHAMPIOHSHIPD8 International Control Contr	Name Resource group name RS-CHAMPIOLSHIPDE Resource group name RS-CHAMPIOLSHIPDE Tags NG CHAMPIOLSHIPDE Policy RULBACCUE_OWNPOINSHIPDE Policy RULBACCUE_OWNPOINSHIPDE Regin Smallcatter Verification Server STNVA3 Verification Server STNVA3	7 Name				
Resource group name RE-CHAMPIONSHIPSE Tags NO-CHAMPIONSHIPSE Policy POLISACCEP_COMMPOSHIPSE Policy POLISACCEP_COMMPOSHIPSE Policy ShapCenter Plage in for Marqueth SQL Senser	Resource group name RS-CHAMPIORSHIPDS Tags ND-CHAMPIORSHIPDS Pelicy RLLBACKIP_OHAMPIORSHIPDE Page N ShapCenter Plagin for Matrixon SQL Server Verification Server STAVA3 Verification Server STAVA3	RECHAMPIONSHIPDE	0-0	0-0-0-0		
Tags NG-CHAMPIONSHIPDI Policy FULLENCRUP, COMMPIONSHIPDI: Hourly LODBACKUP, COMMPIONSHIPDI: Hourly Majsin SinapCenter Flag in for Microsoft SQL Senser	Tags ND-CHAMPODIGHPDD Policy RULLBACKDP CHAMPODIGHPDD Policy COGRACIUP, CHAMPODIGHPDD Plag-IN Sub-Center Plag-IN for Mutmith SQL Gener Verification Server STAVA3. Verification Server STAVA3.		Name Resources	Polices ventication Notification Summary		
Policy PULIANO2P, SOMMPONISHIPSE Hourly COGRACIUP, CHAMPONISHIPSE Hourly Plagetin ShapCenter Plag in for Microsoft SQL Sever	Pekcy RULLBUCKIP_CHAMPORCHPOELHoung LOOBACOUL_OWAPPORCHPOELHoung Plag-in Subjection Pragin for Mitroph 10(L Server Verification Server STAVA3 Verification Server STAVA3		Resource group name	NS-CHAMPIONSHIPDS		
LOGBACKUP, DAAMPONSHPDB: Inverty Plag-th ShapCenter Plag in for Microsoft SQL Sever	ICOGRACIUP, DUALINOVICIP/DBL Hourly Plag-in SnapCenter Plag in for Microsoft SQL Server Verification Server STX/M3 Verification relatived for policy PULLIMCKIP, CHAMPONERIPDBL Hourly		tags	NG-CHEMPIONSHIPOR		
	Wentflation Server STAVM3 Wentflation enabled for policy FULLSLOKUP_CHAMPONEFIETDIB. Hourling		Policy			
	Verification enabled for policy FULLBACKUP, CHAMPIONSHIPDB, Hours		Plag-in	SnapCenter Plug in for Microsoft SQL Server		
	Verification enabled for policy PULLIACKUP_CHAMPSONEHEPDB Hourig O		Verification Server			
Verification enabled for policy FULLIACKUP_CHAMPOLICHIPDB Hourig	Send email iso		Ventication enabled for policy	FULLIACKUP_CHAMPONSHEPDE: Hours		
Send email to			Send email	tio .		
	Activate Windows Press					

Triggering an on-demand transaction log backup for mutiple SQL Server databases

To trigger an on- demand backup of the transactional log for multiple SQL server databases, complete the following steps:

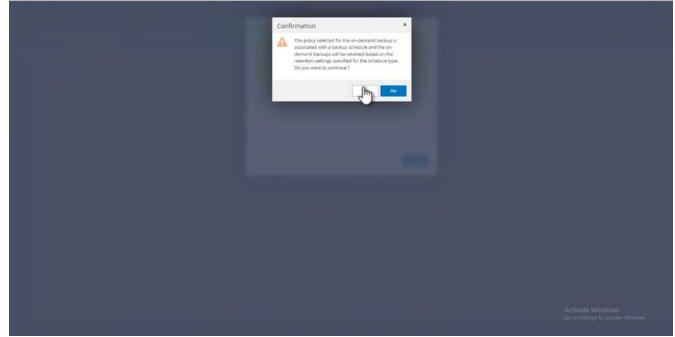
1. On the newly created policy page, select **Backup now** at the upper right of the page.

App SnapCenter®				• # O- Latimistrator SnapCenterAdmin #59
Monsell SQL Sever	RG CHAMPORTHYOD Database			
pearth by riama	Caut			
P Name	Resource Name	Type	Host	0
RG-CMMINONSHIPOB	Championiship561 (STAUM3)	SQL Database	STAVALS DemoA2.com	
	Championship5010 (STAVA0)	SQL Detailuise	STANSE3.DemoA2.com	
	ChampionshipDRTIC (STAVAG)	SQL Database	STRUMS.DemokZ.com	
	Ownproship58101 (STAVM3)	SQL Database	STAVAS Demok2.com	
	ChampionshipD0102(STAVM3)	SQL Detabate	\$1AVALI DemoAlizam	
Secury your pullcies, global securitys, createritue	Championship09103 (578vM3)	SQL Database	STRAMD.Diemok2.com	
and surfigure roles, permissions, and toall tailoncing	ChampionshipD8104 (\$1A\/M3)	SQL Databale	STRIM3.DemoA2.com	
	Championshipb@105(5769M3)	SQL Database	STAMALDOWICK2 com	
	Champion(hipD9108(STAVM3)	10% Database	STAVAL DenicA2.com	
	ChampionishipDib107 (STAUMO)	SQL Datathake	JTAMA DemoA2 com	
	ChampionshipD81081555v9d8	SQL Database	STAIM0.DemoA2.aum	
	ChampionshipD8109 (STAVM3)	SQL Detativitie	STAVM3 DemoA2 com	
	Champiotehip2(811)(574/M3)	SQL thatabase	STAIM3 DemoA2.com	
	ChampionshipOII110(STAVM3)	SQL Database	STAIM3.0emoAZ.com	
	Championehip2(8111(S58VM3)	SQL Detabase	STAGM3.DemoAZ.com	
	Championship08112(578VM2)	SQL takabase	STABAGI DemoAd zom	
	ChampionshipD8113(STAVM3)	SQL Database	STAVM3.DemoA2.com	
	Championship20114(STAVM3)	SQL Database	STAGAD DemoAZ com	
	Championthip08013 (05kVM/3)	SQL Database	17Mht3.DemoAZ.com	
	Champions/sp08114 (STAVM3)	SQL Detabase	STAVM3 DemoA2 com	
	Championship2(9117 (\$14VM3)	SQL Detailarse	STAVA2.Demo/2 rpm	
	Championship08718/074/M/b	3QL Database	STRAFE DemoAz com	
	ChampionshipD8119 (STAVM3)	SQL Detailable	STAVAG SweepAZ zone	
Total 1	ChamaionshipDB12(57AVM3)	TOIL DWINDARM	STAVAT DemoA2 spin	Activate Weddows An to Second to Second Streets

2. From the pop-up on the **Policy** tab, select the drop-down menu, select the backup policy, and configure the transaction log backup.

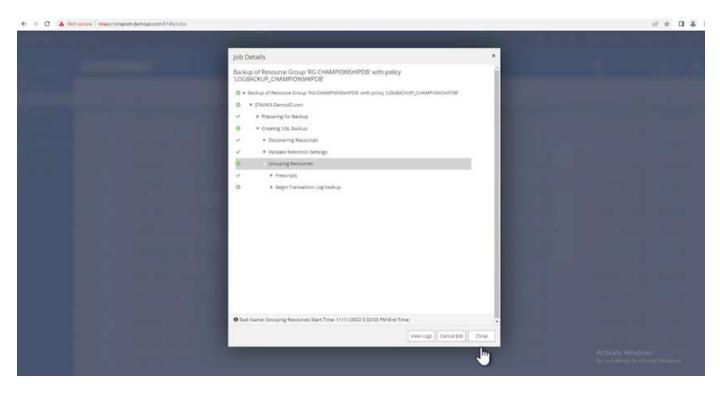
C A Not secure Negec/Unapvm.demoaz.com/l146/InventorySQEProtect/Protect	tindex?ResourceType+nullSSQLServerInstance+null8HSost+null#		2 × 0
	Backup		
	Create a backup for the selected resource group		
	Resource Group RG-CHAMPONSHIPDE		
	Policy PULLBACKUP_CHAMPIONSH - 0		
		nol Mackap	

- 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.

di Yam Query Project	Tools Window Page	
0 2 . 0 . 0	Dem Carry あかかかか (美の古(フ・マ・田)・戸 ・ □ チェロ・。	
W9X00081 -	b look = V 日前日 丁田県 当前の 3 3 4 5 5 3 4 5 4 5	
 The Part of the P	(Sector Social Statistications connect from SUBS - *****/ (Sector (Sector) (Author, and) (Author, and) (Author, and) (Author, and) (Rec (Information).(Aut).(%)Authors)	
ia di Yabut ia di Estenut Ressurces ia di Systempris ia di Programmabity	⊗ °	
I Mill Service Broker	#5 *1	
in 10 Security	III fresh 18 Manaper	
AnonProperty Server Organis Server Organis Deployment Profiles Profiles Profiles Profiles Profiles Profiles Profiles Management December Server Enterprofiles SG, Server Agent Stocker Profiles	B Addregation Interf 1 Marker 1 Compty - Some 2 2 Adirs - 2 Compty - Some 3 3 Adirs - 3 Compty - Some 4 4 Adirs - 3 Compty - Some 5 5 Adirs - 4 Compty - Some 6 6 Adirs - 5 Compty - Some 7 7 Adirs - 7 Compty - Towne 8 6 Adirs - 7 Compty - Towne 9 6 Adirs - 16 Compty - Towne 10 14 Adirs - 16 Compty - Towne 11 14 Adirs - 16 Compty - Towne 12 2 Adirs - 16 Compty - Towne 13 14 Adirs - 11 Compty - Towne 13 14 Adirs - 11 Compty - Towne 13 14 Adirs - 11 Compty - Towne 14 14 Adirs - 11 Compty - Towne 15 14 Adirs - 16 Compty - Town	

The screenshot shows that 20 rows were deleted from the table.

SCORENCEDERMENTORIDET - downsilw			Radies (Administration					Gask Lauren (DA-Q	P - 1
· 0 0 · 0 · 0 0 0 30	- Carry 20 21 3	02010			· 2741	3+, (2 = ≈ 11 Our	piteri te da Grille te a		
a statement of the second second second second second second second second second second second second second s				SIGDIS.		110000			_
d lighter + 1 X	SCEPHACKES MAKE	Author name	Analy and a second seco	SQ. Administrative (TV)	93.0 myfael: 90, Adminis	tation (\$20			
mate # 19 11 1 6 4		Author - 21	Country - 21 Ma.						
SOLHMORDS (SQL Server 14.0.1000.188) III IIII Detelanesi	and the second second	Author - 22	Country - 22 Adv.	N					
an mill Sectem Catabases		Battor - 23	Country - 21 ma	42					
a III Dutation Srephots			Level of the second second	10					
II G DENACORT	25	Author - 25	Country - 25 mil-						
# B DEMODE2	28	Author - 25	Churthy - 26 Ab.						
a B DEMODES	127	diathere = 27	Country - 27 ma.						
E B NGMTON/001	28	Autor - 28	Country - 28 mil-						
III III Databasa Diegrame	29	Author (28	Country - 25 mil-						
10 MI Tables 10 MI Tables	10	Author-30	Country - 30 rok						
an and FileTables	10	Author - 11	Country - 31 ers						
in this External Tables	12	Author - 32	Country - 32 ma						
I III Graph Tables	100	Author - 33	Centerly - 32 etc.						
is 📅 dba.Impleyee	14	distant - M	Country - 34 mar.						
i = ₩ 350.m/Tati+	100	Author-35	Country - 35 ma.						
a m dha thiladhari	10	Author - 10	Country - 30 ins						
a 🖬 die miliooke	12	Adder - 17					 		
a di Litrai Previn			Courtery - 37 mb				4		
at the Sumaryona	10	Autor-18	Country - 30 est-			8 0			
a si Programmatility	. 20	Author-38	Country - 32 ne.						
at 10 Service Broker	-	Autor - 40	Cautany dina.						
an sit Storage	41	Author-47	Country - 41 ma						
22 HE Security	40	Autor 42	Courtey + 42 etc.						
B INVENTORIORI	40	Ruther + 43	Country - KI ma						
III III Security III III Server Otjects	44	Author-Mi	Coursey-Mine.						
i il Aplicates	48	Author - 45	Country - 45 mm.						
III III PulyEste	145	Author-46	Courtry - 80 ma						
III III Always On High Availability	47	Audion: 47	Country - 47 mau						
🔬 📾 Management	100	Autor - 42	Country 41 etc.						
👷 📶 Integration Sensors Catalogs	-	Author - 45	Country - Alima						
■ 昌 SQL Server Agent									
I T XIvent Profiler	*	Jutter - 58	Country - Stime.						
	M	Author - 51	Ceantry - 51 etc						
	9	Author - 52	Casarity - S2 est.						
	50	Author-55	Country - 52 ma						
	54	Author - 54	Country - 54 Hal-						
	46	Authors 55	Country- 35 ma.						
	- 56	Authors 55	Courtey - 35 na.						
	12	Juttor-57	Country - 57 ma					Letisate Windows	
	30	Author: 58	Caurery - 50 Ha-						
	30	Author: 58	Country-West					Gan for Sarbingst his airbriade We	
	14 4.11	uf 180 & bi b							

2. Log into SnapCenter Server. From the **Resources** tab, select the database.

	Simpe	enter®					• = •-	Ladministrator SnapCenterAdmin VSign
		Microsoft	SQL Server					
Dashbisa	rd	View 0	search by name					Administration instant
Resource	e6 :	17 1	Name	Instance	Host	Last Backup	Overall Status	Туре
Monitor			INVENTORYDE1		FSICLISTER.Domoac.com	11/16/2022 10:32:44 PM	Backup succeeded	Over detailuise
			INVENTORYDB2	SQLHMCRDS	PSXELUSTER.Demoac.com	11/16/2022 10:50:49 PM	Backup succeeded	User database
Reports	1		master	SQUHMCROS	PSXCLUSTER.Demokz.com		Not protected	System database
Hosts			model	SQUHMCADS	PSXCLUSTER.Demoar.com		Not protected	System database
Storage	Systema		midti	SQUHMCROS	FSRCLUSTER.Demoar.com		Not protected	System database
Settings			tempdb	SQUHMCRDS	PSXCLUSTER.Demoar.com		Not available for backup	System database
	2		DEMODB1	SQLHMCRDS	PSXCLUSTER.Demoat.com	11/16/2022 11:03:57 PM	Backup succeeded	User database
Alerts			DEMODB2	SQUHMCHDS	PSACLUSTER Demoac.com	11/16/2022 11:08:57 РМ 렀	Backup succeeded	User database
			DEMOD83	SQLHMCROS	PSXCLUSTER.Demoat.com	11/16/2022 11:03:57 PM	Backup succeeded	User database

- 3. Select the most recent backup.
- 4. On the right, select **Restore**.

÷ • /				SnapCenter®	Арр		
			RVXINTORYDR1 (SQLIMCROS) Topplagy	soft SQL Server	More		
Researchmenter Rading-Stee Marky 1				ron by name	-		
A second s			Manage Copies	Name	17		
				WHENTORYDE!			
Summary Card			Cones	INVENTOR/DB2			
a Backups				master			
© Clones			Local copies	model			
				madb			
				tempdb			
			DEMODB1 Primary Backup(s)				
			(sketh	DEMODES			
				DEMOD83			
4F End	Count Type	⊗ ° ∰	Backup Name				
(15/05/2022-452)267	1 Pull backup	Υ.	SQLHM0005_04V0V(001001_femines1_11-78-2022_16.36-87.9645				
11/15/2022 6:48:02 /	1 Full backup		SQLHMCRDS_UVENTORYDB1_ftvm-rin1_11+15-2022_18.44.24.9788				
11/15/2022 4:41:10 #	1 Full backup		SQLHMCROS_INVENTORYDB1_form-rds1_11-15-2022_16.31.08.0038				
11/15/2022 4:31:167	1 Full backup		\$QUHMCRD5_INVENTORYDB1_fevm-rds1_11-15-2022_16.27.32.3928				
11/15/2022 648:02 / 11/15/2022 441:10 /	1 Fullbackup 1 Fullbackup		SQLHMCRD5_0WENTORYDB1_fivm=rb1_11+15-2022_18.44.24.9788 SQLHMCRD5_1WVENTORYDB1_fivm=rb1_11+15-2022_16.31.08.0038				

- 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.

T SnapCenter X +			~ - 0
← → C A Not secure here//snapsmidemcaz.com/8140/linen	torySQUProtect/Protectl	ndee/RenouiceType i multBSQLServerInstance i multB.Host i mult	e a 🗆 🛎
	in a state of		
and the second se	Restore	×	
	1 Restore scope	Restore database files using	
	2 Recovery Type	Restore the database to the same host where the backup was created.	
	3 Pre Ops	O Restore the database to an alternate host 0	
	4 Post Ops	C Restore the database using existing database files: 0	
	5 Notification		
1000	6 Summary	A.	
		⊗ ∘	•
	A To configure a	in SWTP Server to send email notifications for Restore jobs, go to Settings-Skotal Settings-Natification Server Settings.	
		Pressent Next	
		Activate	

7. For the Recovery type, select All log backups. Click Next.

Previous Next Invertion/SCLProtect/Index/MesourceType ===ulfdS/QLenvertextance===ulfdAtair.mull Restore x Invertex source Choose logs Invertex source Choose logs Invertex source Choose logs Invertex source Choose logs Invertex Image: Decision of the source o
Previous Next Invertion/SCLProtect/Index/MesourceType ===ulfdS/QLenvertextance===ulfdAtair.mull Restore x Invertex source Choose logs Invertex source Choose logs Invertex source Choose logs Invertex source Choose logs Invertex Image: Decision of the source o
Previous Next Invertion/SCLProtect/Index/MesourceType ===ulfdS/QLenvertextance===ulfdAtair.mull Restore x Invertex source Choose logs Invertex source Choose logs Invertex source Choose logs Invertex source Choose logs Invertex Image: Decision of the source o
Previous Next Invertion/SCLProtect/Index/MesourceType ===ulfdS/QLenvertextance===ulfdAtair.mull Restore x Invertex source Choose logs Invertex source Choose logs Invertex source Choose logs Invertex source Choose logs Invertex Image: Decision of the source o
Previous Next Invertion/SCLProtect/Index/MesourceType ===ulfdS/QLenvertextance===ulfdAtair.mull Restore x Invertex source Choose logs Invertex source Choose logs Invertex source Choose logs Invertex source Choose logs Invertex Image: Decision of the source o
Previous Next Invertion/SCLProtect/Index/MesourceType ===ulfdS/QLenvertextance===ulfdAtair.mull Restore x Invertex source Choose logs Invertex source Choose logs Invertex source Choose logs Invertex source Choose logs Invertex Image: Decision of the source o
Previous Next Invertion/SCLProtect/Index/MesourceType ===ulfdS/QLenvertextance===ulfdAtair.mull Restore x Invertex source Choose logs Invertex source Choose logs Invertex source Choose logs Invertex source Choose logs Invertex Image: Decision of the source o
Previous Next Invertion/SCLProtect/Instances/Types inulfib2(2).ServerInstances inulfib.Host inulfi Restore * * * * * * * *
Previous Next Invertion/SCLProtect/Instances/Types inulfib2(2).ServerInstances inulfib.Host inulfi Restore * * * * * * * *
InvertorySQUProtect/IndextResourceType insultaSQUSenvertestance insultaNost insult Restore × Choose logs Choose logs Choose logs backups duration Choose log backups duration Choose log backups duration Choose log backups duration
Restore × Restore kode Choose logs Restore kode Choose log backups duration High bedues 0
Restore × Restore kode Choose logs Restore kode Choose log backups duration High bedues 0
Austance kooper Choose logs Choose log backups duration All tig backups O
Restrict Kope Chaose logs Choose log backups duration Alting backups O
Receivery Type Choose log backups duration Alling backups O
Allogbaciups 0
3 Provides
 By lag backups emil 11:116/2022 5:01:25 PM O
a Prat Ope O By specific date until 11/16/2022 (x07/25 PM
3 Nutflatton O None 11/7/0/2022 5/37/21 PM
11.116/2022 5-52-26 PM
 Summary 11/16/2022 #07/25 PM 11/16/2022 02:25 PM
11/16/2022 4:37:25 PM
11/1/6/2022 6 2/2 PM (8) · · · · · · · · · · · · · · · · · · ·
To configure an SMTP Server to send email notifications for Restore Jobs, po to : Settimes-Silotal Settimes/Hotelsation/Settimes/Hotelsation/Settimes.

Pre- restore options:

1. Select the option **Overwrite the database with same name during restore**. Click **Next**.

nter x +					
C A Not secure Heren/Innepum.demoaz.com	r8148/inventorySQUProtect/Protect	Index/HesourceType : r	null65QESenerinstances null6thtats null		ピ 幸
				*	
	Restore			-	
		Pre restore optic			
	Recovery Type	C Retain SQL data	database with same name during restore abase replication settings		
	3 Pre Ops	Quit restore if t	tion log backup before restore transaction log backup before restore fails		
	4 Post Opt	Specify optional	scripts to run before performing a restore job 0		
	5 Notification	Script full path	<scripts_path+< td=""><td></td><td></td></scripts_path+<>		
	6 Summary	Script arguments	Choose optional arguments	1	
		Script timeout	60 9853		
			⊗ ∘	~	
	To configure	an SMTP Server to serv	id email notifications for Restore jobs, go to <u>Settimes-Groba</u>	Contract Protification, Second Section,	
				Previous Next	
	18				

Post- restore options:

1. Select the option Operational, but unavailable for restoring additional transaction logs. Click Next.

StapCenter K +		v - a
← → C ▲ Not secure Helper//snapvm.demcaz	ccom 8144/eventurySQLProtect/ProtectIndes/ResourceType=multiSSQLServerinstance=multiHoct=mult	@ * O #
	Restore	
	Restore score Post restore options	
	Recovery Type Select the database state after restore completes	
	Operational, but unavailable for restoring additional transaction logs Operational control operational control operation operat	
	Non-operational, but available for restoring additional transaction logs	
	Pool Opy Tead-only mode, and available for restoring additional translation logs	
	3 Noofication Unde file directary. Intrechilder	
	6 turning Specify optional scripts to run after performing a restore job 0	
	Script full path escellating anti-	
	Script arguments Choose optional arguments.	
	🛕 To configure an SMTP Server to send email notifications for Restore jobs, go to: Second Second Sections restrictions Second Second Sections	
	Previous	
	0	
anna llanannan daaraan oon 1920 (naanno 1971 footar) Sontarioatar	Parama Tanana (A. C. Kananana ang Kanana ang Kanana)	

2. Provide the email settings. Click Next.

Restore	×	
Restore stope Provide email	settings 0	
The second state of the se	a have -	
O Pre Cas	from enail	
To Page Day	Unal to	
Subject	Nonfication	
6 Summary		
T Semanal		
A If you want to send notifications	for Rectore jobs, an SMTP server must be configured. Continue to the Summary page to save your	
A If you want to send notifications information, and then go to Sett	for Restore jobs, an SMTP server must be configured. Continue to the Summary page to Save your ings-Global Settings-HottiCation Server Settings to configure the SMTP server.	

3. On the **Summary** page, click **Finish**.

Disalieter X 4			5
€ → C A hist secure seems//unap-mide	enced constitution and static framework (CDP inters), Promechades Phase and	pper en AREQ.Serverbetance en alliteten en al	8 A O & 1
	Restore	*	
	terministicopie Summary		
	Discourse Type Beckup that	 SQLAWCROS, WWW/D00001, Assessing 1, 11-10-2022, 10-20-07 20-05 	
	Eachup typ	Reltake	
	O recogning data	1016203743231294	
	O Past Dev Restore typ	117509	
	Austoire Ing	Unit 10160022 EB226 PM	
	Send email	ber .	
	(6 summary		
		⊗ o	
	A Council is suit weter	end for Restore poly, an SMTP carver must be configured. Continue to the Summary page to save your	
	information, and then go to	entrops - Cicital Sectorys - Hamiltanian Server Settings to configure the SMTP Server.	
		Previous 2000	
	in the second second second second second second second second second second second second second second second		
Miss. Totagoon demonstrational Md. Inventor/SCI Present Present	Trebs Resource Experimental Social Services and Brings and Ministry		

Monitoring the restore progress

1. From the **Monitoring** tab, click the restore job details to view the progress of the restore job.

■ StapComer		* +	m.demoas.com/1140/lobs			e o 🗆 4
NetApp SnapC	-				0. Lainminer Success	1000 C 4400
reetapp snape			Territy Line		A Constanting The Arrive	Arrent Bulace
	100					7.44
Detthord	(1445N	29.1016	T ²		9 . .	<u> </u>
Answeren .	Jobs - F	iter			0	
Monitor	10	Status	Same	Start date	End date	Owner
Reports	124		Rennes 50, HeCRONWINTORIORY	11/16/2020 11:11:12:04 (0)		Administration
Concession of the local division of the loca	355		Budiup of Resource Group SQUHMCRDS, INVENTOR/OR1 with policy InventoryOR_JogSediup_policy'	11/16/2022 11:00:01 PM		Adminuterto
1400	114	-	Bachup of Reloute Group 361-DIMODE with policy BernoDB logBackup policy	11/16/2023 10/36/02 PM 🖨	11/16/2022 11/10/54 PM 🛤	Administrator
Sounge Systems	127	1.	Backup of Resource Group, SQLHMCRDS_WVII/109/0082 with parky Sweetary002_N/Backup	11/16/2022 10:55:01 PM	11/16/2022 10:55:50 PM	Almostuto
Settings	152	4	Backup of Resource Group SQUEMCRDS, HV18110910ET with policy TrivenonyOR, high-adual, pullicy	11/16/2022 10:45/01 PM E	11/16/2022 11:10:54 PM	Aprovation
	131	100	Backup of Resource Group WGI-DEMODEL with policy demoDE-lightackup policy.	11/18/2022 10:44:02 PM	11/16/2022 10:55:50 PM	Administrato
Alinta	150	14	Backup of Resource Group SQLHMCRDS, HV5N704YDB1' with policy Inversity/DB, ingliaickup, policy'	11/16/2022 10:30:01 PM	11/16/2022 10:55:54 PM 🖨	Abrolite
	148	~	Backup of Resource Group HSI 0633008 with policy terms08-highactup policy	11/16/2022 10:29:02 PM	11/16/2022 10:40:53 PM 🗖	Administration
	140		Backup of Resource Group SQUHACROS_AVUM/DITIOR1 with policy InventoryOIIL/legtackup_policy	turia/2022 teriseir PM 🛱	11/16/2022 Tel:40:53 PM 🖨	ASPAROUN
	147.	191	Rectup of Resource Group WOL-OSMODIT with policy demoDB-tugbeckup policy	11/16/2022 10:14:02 PM	11/16/2022 10/25/53 PM 🖽	Administratio
	140		Backup of Resource Group SQLHMORDS_moth/tom/DBT with policy trivenosryDB_logbackup_policy'	11/16/2022 10/00/01 PM 🛱	11/16/2022 10/25:53 PM 🖨	Adminuteget
	145	1940	Backup of Resource Group WSI-DEMODIF with policy demoDifilingBackup,policy	11/16/2022 3:58:02 PM	71/16/2022 10:10:53 PM	Administratio
	144	190	Because of Resource Group SQUHMCRDS_NVINTORYD91 with policy InventoryD8_logDeckup_policy	T1/16/2022 1945/31 PM	TM/16/2222 10:10/31 PM	Administration
	143	1	Bactup of Resource Group 761-01M008 with policy temu08-logteckup policy	11/16/2022 9:44:02 PM	11/10/2022 9:35:54 PM	Administratio
	142	1	Backup of Resturce Group SQUARCROS, WVINITORY081" with policy Truentity/OB_logBackup_policy	11/16/2022 #38001 PM 🛱	11/16/2022 9/35/34 PM 🗖	Administratio
	548	191	Bactup of Kessuros Group 301-DEMODIF with policy demoDillinghackup policy	11/16/2022 R/28/52 PM	11/16/2022 3 40/58 PM	Administration
	140	×	Bachup of Resource Group SOLHMCRDS, WVIII/108/DB1 with policy Invence/OB, logbachup, policy	11/15/2022 915/01 PM	11/16/2022 9:40:53 PM 🗖	Administratio
	139	1	Backup of Resource Group 'RGF-DEMODB' with policy demodB-logbackup, policy	11/16/0022 3:14/02 PM 🗖	11/16/0022 9/25:54 PM 🗖	Administratio
	138	9	Rectup of Resource Group SQLHMCRDS_WV0x100Y087 with policy Trivenosy08_logBeckup_policy/	11/16/2022 1000/21 PM	11/16/2022 9/25/54 PM	Administrati
	137	30	Rectup of Resource Group WDL-DRMODR with policy themoDB-higherkiep,policy	11/15/2022 8/59/02 PM	11/16/2022 9(10:53 PM 🛱	Administration
	138	1	Backup of Resource Group SQUARCEDS, WHINTOFFIDET with parky InventoryOB, Inglandus, pailing	11/16/2022 3145/01 PM 🗖	11/18/2722 10:003 PM 🖨	Atministratio
	135	1.00	Backup of Resource Group WG1-DEMOOBT with policy demoDB-loghackup_policy	11/16/2022 0-44/02 PM	11/14/2022 6:56:54 PM	Administrate
	134	1.0	Backup of Resource Group SQLHMCRDD, INVENTORYOBY with policy InvenoesyDB, inglackup, policy	11/15/2022 8:30/01 PM	11/56/2022:635:54 PM 🛱 🖉	Attention
	135	1	Backup of Resource Group WSI-DEMODD' with policy demoDB-legbackup, policy'	11/16/2022 8/29/52 PM	11/18/2922 8-4053 PM 🗖	Almostyle

2. Restore the job details.

I Server A 14		
F I C A Not secure 1 Here (Interpret demost condition)		19 A D &
	job Details *	
	Restore SQLHMCRDSWIVENTORYDB1"	
	O * Ressue SQUEMCODSUMENTORIDET	
	 Use 157 LTM legbackup of SQUMCR09/INVENTOR/DB11 	
	w. * KVM-RDS1.Demost.com	
A second s	✓ ➤ Presaring for Backup	
	vr + Creating SQL Backsip	
	✓ + Initing takip	
	e Sect IVI Metaget	
	FCVM4DS: Demost.com	
	⊗ °	
	O faits name: Send BMS Messages Start Time: 11/16/2022 11:18:54 PM Evo Time: 11/16/2022 11:18:54 PM	
	versign cruze	

3. Return to SQL Server host > database > table are present.

C 172.30.4.152 - Remote Desktop Connection	R/DB1 (FCVM-RDS1/Administrator (IRI) - Microsoft SQL Server Management Studio (Administrator)		Cuick Emptoh (Chil-Q)	-	× ø
	terology あいないない X 0 の フ・マ・四 - ● - beode - ✔ 27 週日 27 25 単 週間の 3ち 45 ちょ	· 印 / 主 曰 · .			
Constant	SGCAury Rad, SGL Adenciature (BD) = >K /******* [SELECT IOF (1996) [14] .[Author_mem] .[country] .Rade [InvEnDence]].[dbu].[tblAuthors]	⊗ 0			
■ Soviet Broker ■ Soviet Broker ■ Soviet Broker ■ Soviet Broker ■ Soviet Broker ■ Soviet Broker ■ Soviet Broker ■ Soviet Broker ■ Soviet Broker ■ Soviet Broker ■ Soviet Broker ■ Soviet Broker ■ Soviet Broker ■ Soviet Broker ■ Soviet Broker ■ Soviet Broker	LO %				

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/protectscsql/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 datamanagement 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 todays 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.

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

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 missioncritical 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.

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.

Quick create	Standard create
Use recommended best-practice configurations.	You set all of the configuration options, including
Most configuration options can be changed after	specifying performance, networking, security,
the file system is created.	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 AP	4.
O 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.	
O 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.

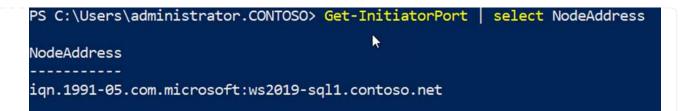
```
Start-service -Name msiscsi
Set-Service -Name msiscsi -StartupType Automatic
```

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



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)
www.sql::> 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
/server Path
                                         State Mapped
                                                          Туре
         /vol/vol_awssqlprod01_data/lun_awssqlprod01_data
svmsql
                                          online unmapped windows 2008
                                                                       700GB
svmsql
         /vol/vol awssqlprod01 log/lun awssqlprod01 log
                                         online unmapped windows_2008
                                                                       100GB
 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_awssqlprod01lun map -path
/vol/vol_awssqlprod01_log/lun_awssqlprod01_log -igroup
igrp_awssqlprod01
```

svmsq1::> Vserver	lun show Path	State	Mapped	Туре	Size
svmsql	I /vol/vol_awssqlprod01_data/lun	_awssqlpr	od01_data	 1	
		online	unmapped	d windows_20	08 700gB
svmsql	/vol/vol awssqlprod01 log/lun a	awssqlprc	d01 log		700GB
				d windows_20	08
	and the second second				100GB
entries	were displayed.				
vmsql::>	lun map -path /vol/vol_awssqlp	rod01_dat	a/lun_aws	sqlprod01_d	ata -igroup igrp_awssqlprod01
umcal	lun map -path /vol/vol awssqlp:	nod01 100	(1)	alprod01 lo	a ignoup igno succedenedol
ovinoq1>	iun map -pach /voi/voi_awssqip.	10001_100	/IUI_dwba	dibrogor_10	g -igioup igip_awssqipiodoi
svmsql::>					
svmsql::>	lun show Path	Ctate	Mannod	WITTE O	Cinc
/server	Pauli	State 	Mapped	Туре	Size
svmsql	/vol/vol_awssqlprod01_data/lun	_awssqlpr	od01_data	1	
vmsql	/vol/vol_awssqlprod01_data/lun		od01_data mapped	windows_20	
		online	mapped		 08 700GB
	/vol/vol_awssqlprod01_data/lun_	online awssqlpro	mapped		700GB
svmsql		online awssqlpro	mapped d01_log	windows_20	700GB

- 14. 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.
- 15. Connect Windows Server to an SVM with an iSCSI target. Find the target IP address from AWS Portal.

Summary		
SVM ID svm-09e98ab33a31b724a SVM name svmsql UUID ea00ea2d-1b1d-11ec-9de1-6f9cef731025 File system ID fs-0ab4b447ebd6082aa Resource ARN arr:aws:fsx:us-west-2:139763910815:storage-virtual-machine/fs-0ab4b447ebd6082aa/svm-	Creation time 2021-09-21T13:19:34-07:00 Lifecycle state O Created Subtype DEFAULT	
09e98ab33a31b724a D Endpoints Wanagement DNS name svm-09e98ab33a31b724a.fs-0ab4b447ebd6082aa.fsx.us-west-2.amazonaws.com D NFS DNS name svm-09e98ab33a31b724a.fs-0ab4b447ebd6082aa.fsx.us-west-2.amazonaws.com D		Management IP address 198.19.255.153 ① NFS IP address 198.19.255.153 ①
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

16. 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.

	perties			Discover Target Portal	
gets Discover	Y Favorite Targe	ts Volumes and Devices	RADIUS Configuration	Enter the IP address or DNS name a want to add.	and port number of the portal you
Farget portals				To change the default settings of th	e discovery of the target portal, di
The system will	look for <u>T</u> argets o	n following portals:	Refresh	the Advanced button.	
Address	Port	Adapter	IP address	IP address or DNS name:	Port: (Default is 3260.)
				10.2.1.167	3260
				Advanced	OK Cancel
To add a target	t portal, click Disco	ver Portal.	Discover Portal		
i e ann a miger					
To remove a ta	rget portal, select	the address above and	Advanced Settings		2
To remove a ta	rget portal, select	the address above and	Advanced Settings General IPsec		?
To remove a ta then click Remo	rget portal, select	the address above and			2
To remove a tai then dick Remo SNS servers	rget portal, select	the address above and	General IPsec	Microsoft iSCSI Initiator	?
To remove a tai then dick Remo SNS servers	rget portal, select		General IPsec	Microsoft iSCSI Initiator 10.2.2.232	
To remove a tai then click Remo SNS servers The system is re	rget portal, select		General IPsec Connect using Local adapter:		
To remove a tai then dick Remo SNS servers The system is re Name	rget portal, select	ollowing įSNS servers:	General IPsec Connect using Local adapter: Initiator IP:		

- 17. Repeat step 12 for the second iSCSI IP from the SVM.
- 18. Select the Targets tab, select Connect, and select Enable muti-path.

argets	Discovery	Favorite Targets	Volumes and Devices	RADIUS	Configuration
Quick C	Connect			olumes and Devices RADIUS Configuration a basic connection, type the IP address or ick Connect. Quick Connect Quick Connect Quick Connect Connect To Target Target name: 32-08.com.netapp:sn.ea00ea2d1b1d11ec9de16f9cef731025:vs.3 Add this connection to the list of Favorite Targets. This will make the system automatically attempt to restore the connection every time this computer restarts. Image: Connect Image: Connection every time this computer restarts. Image: Connect OK Cannetter Connect	
) on to a target usir arget and then click		ype the IP	address or
<u>T</u> arget	:			Qu	ick Connect
Discove	ered tar <u>q</u> ets				Connect To Target
Name	:			Status	Target name:
ign. 19	992-08.com.i	netapp:sn.ea00ea2	2d 1b 1d 1 1ec9de 16f9	Inactive	92-08.com.netapp:sn.ea00ea2d1b1d11ec9de16f9cef731025:vs.3
					This will make the system automatically attempt to restore the
					∑ <u>E</u> nable multi-path
				Devices RADIUS Configuration nection, type the IP address or t. Quick Connect Connect To Target Target name: 92-08.com.netapp:sn.ea00ea2d1b1d11ec9de16f9cef731025:vs. Add this connection to the list of Favorite Targets. This will make the system automatically attempt to restore the connection every time this computer restarts. Enable multi-path Advanced OK	

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}
```

<pre>\$TargetPortals = ('10.2 foreach (\$TargetPortal</pre>	2.1.167', '10.2.2.12') in \$TargetPortals) {New-IscsiTargetPortal -TargetPortalAddress \$TargetPortal}
Initiator InstanceName Initiator Portal Address IsDataDigest IsHeaderDigest TargetPortal Address TargetPortal PortNumber PSComputerName	: : False : False : 10.2.1.167
InitiatorInstanceName InitiatorPortalAddress IsDataDigest IsHeaderDigest TargetPortalAddress TargetPortalPortNumber PSComputerName	: : False : False : 10.2.2.12

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)
InitiatorNodeAddress InitiatorPortalAddress InitiatorSideIdentifier IsConnected IsDataDigest IsDiscovered	<pre>NONE ROOT\ISCSIPRT\0000_0 iqn.1991-05.com.microsoft:awssqlprod01.cloudheroes.dom 0.0.0.0 400001370000 True False True False True i false i True i false i True i false 2 1 ffff9988350ff010-4000013700000012 i qn.1992-08.com.netapp:sn.ea00ea2d1b1d11ec9de16f9cef731025:vs.3 0200</pre>

2. Initialize disks with the following PowerShell command.



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.

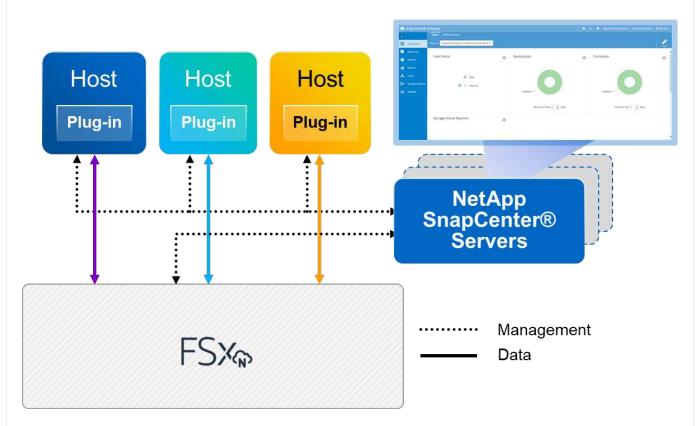
SnapCenter overview

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. 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	Windows Server 2012
	Windows Server 2012 R2
	Windows Server 2016
	Windows Server 2019
Software packages	.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 re	equirements.

For version compatibility, see the NetApp Interoperability Matrix Tool.

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.

i

(i)

 (\mathbf{i})

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.

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.

<	Global Settings Policies Users and Access	Roles Credential Software			
Dashboard	Search by Credential Name				New
Resources	Credential Name	Authentication Mode		Details	
C Monitor	There is no match for your search or data is not availab	sle.			
ार्गे Reports		Credential	×		
📥 Hosts		I			
Storage Systems	A.	Credential Name SCAdmin			
Settings		Authentication Mode Windows Username rdccustomval/Administrator Password]		
		Cancel	ок		

2. Add the storage system by selecting Storage Systems > New and the provide the appropriate FSx for ONTAP storage information.

<		ONTAP Storage						
	Dashboard	Type ONTAP SVMs Search by N	ame					<u>t</u>
	Resources	ONTAP Storage Connections						
۲	Monitor	Name		IL IP	Cluster Name	User Name	Platform	Controller License
ai	Reports	There is no match for your search or data is not available.						
۸	Hosts							
24	Storage Systems		Add Storage System					
			Add Storage System					
			Storage System	198.19.255.71				
			Username	fskadmin				
			Password					
			Event Management S	stem (EMS) & Aut	oSupport Settings			
			Send AutoSuppo	t notification to st	orage system			
			Log SnapCenter					
			More Options : Pla	tform, Protocol, Pr	referred IP etc			
			Submit Cancel	Reset				

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.

<	Managed Hosts Disks		Initiator Groups								
Dashboard	Search & Name	Y									<u></u>
Resources	Name Name		Ц Туре		Syster	n	Plug-in			Version	Overall Statu
Monitor							There is no mate	h for your search or data	is not a	available.	
Reports			Managed Hosts								
Hosts			Search by Name	e:		Add Host					
				lame	1 in	Host Type	Windows	•			
		📀 T		for your search or d available.		Host Name	10.0.1.85				
		11	10	overbore.		Credentials	SCAdmin	.*	+		
		Δ.				Select Plug-ins to Inst	all SnapCenter Plug-ins Paci	age 4.6 for Windows			
		20					Microsoft Windows Microsoft SQL Server				
		=					Microsoft Exchange Ser	rer			
		A					SAP HANA	n Plug-Ins			
						· mara spinning / 0					
						Submit Cancel					

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 $\% \times$ (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.

Manag	ed Hosts			
Searc	ch by Name		Host Details	
	Name	IL.	Host Name	RDSAMAZ-FFIDFMR.rdscustomval.com
	RDSAMAZ-		Host IP	10.0.1.56
	FFIDFMR.rdscustomval.com		Overall Status	Configure log directory
			Host Type	Windows
			System	Stand-alone
			Credentials	SCAdmin 🖋
			Plug-ins	SnapCenter Plug-ins package 4.6.0.6965 for Windows
				✓ Microsoft Windows
				Microsoft SQL Server <u>Remove</u> <u>Configure log directory</u>
			More Options : Po	ort, gMSA, Install Path, Add Plug-Ins
			R	
			Submit Cancel	Reset

If the LUNs have already been provisioned, you can select the mount point to represent the host log directory.

Configure host log	directory	
Host log directory	dedicated disk directory path	Browse
Choose directory of	on NetApp Storage	
RDSAMAZ-FFIDF	MR.rdscustomval.com	
D:\FSxN\	Data\	
D:\FSxN\	HLD	
D:\FSxN\	Log\	

Now you are ready to perform backup, restore and clone operations for SQL Server.

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**.

	Microsoft	KQL Server 🕒					
Detboard	Wey D	table . Heriti fa tarre	7				
Records	17.14	Name	instance	Most.	Last Galogo	Overall Status	Type
Monther		DWCooffguration	RDSAMAZ-HODEMR	RDSAMR2 INDEMR.ebuschemial.com		twist availables for thickup	1/ker database
	-	Diffiguring -	And And And And And And And And And And	(and the second s		this automa his home	Liter database
al Augurta		DWQueue	REFERENCE FREETENRY	RDSAMA2 FTIDPMR:rdurumo-sal.com		him available for backup	Dorr database
S Heats		History	RESEARAZ FEIDEMIR	BDSAMA2 IT OF MILLION INVALUES		feat available for backup	Typers database
I Storage Systems		model	RETEARAD FILEFARR	Richald Antibelik Antibelik Antibelik		NUT evaluative for backup	Sydem-database
		reads	RDLAMAZ 4TIDEMR	RDSAMA2 #TID/ MR.rduumbewat.com		tust available for backup	System detailatere
E Settings		Segnatural	RDSAWAG FEDERAR	RDSWARD PROPORT putations all com-		Yest primation?	they database
A ANITE		tempdb	RDIGAMAZ PTIDEMR	RDSAMA2 PEDFARL/bitumomail.com		faint available for tackup	Typer 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.

U.	Name								
	DWConfiguration								
	DWDiagnostics	Resource Policies Verification Notification Summary							
	DWQueue	Resource Folkles Verification Folklation Summary							
	master	Colort one or more policies and configure schedules							
	model	Select one or more policies and configure schedules							
	msdb	Full Backup -							
	SeattleRetail								
	tempdb	Configure schedules for selected policies							
		Policy Lie Applied Schedules Configure Schedules							
		Full None To schedule operations select a policy that has the appropriate schedule associated, or modify Backup the selected policy to allow schedules.							
		13							

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**.

E N	ame	40							
D	WConfiguration		0			4	5		
D	WDiagnostics		Resource	Policies	Verification	Notification	Summary		
D	WQueue		Resource	1 Officies	venneadon	Nouncation	Summary		
m	naster		Coloretable	10					
m	nodel			rification serve					
	nsdb		Verification serv	er Select one o	Select one or	more servers			
Se	eattleRetail								
te	empdb		Configure ve	rification sched	ules				
			Policy 11	Schedule Type		Applied Schedules		Configure Schedules	
			There is no mat	ch for your search o	or data is not availat	le.			

4. Click **Back up Now** to test the backup. In the pop- up windows, select **Backup**.

Backup			
Create a backup	for the selected reso	urce	
Resource Name	SeattleRetail		
Policy	Full Backup	- 0	
🗌 Verify after back	up		

5. Select **Monitor** to verify that the backup has been completed.

	jobs	Schedules	Dwnts Logs			
Determi	-	this same				
Manager 1	Jobs	Filter.				
. Martine		Status	Name	Start date	End date	Dener
a Regards	54	10	Backup of Nanicova Group NDGANOA FEDENIK SeattleNets? with policy YUE Backup	0.0225/2022 1547/30 AM	@25/2022 (1474) AM	RESCOTORNAL/administration
Stateset.	88.1	10	Criste Resource Group ROSAMA2 PROFAM, SeattleHetal?	03/25/2022 11/2024 AM	10091/2022 (145/25 AM 🛱	RESOUTERWACKERSTEINER
None -	32	× .	EmakePolicy Yull Backage	BARANDER STEREFERE	UNON DOLL THE AD AM	RDSEUSTOMVALUE: UNIT
Storage Systems	31	10	Deameer resources for all heads	03/25/2022 1100:12 AM: 🛱	101010-0022 1(10117 AM	RDSCUSTORVACUMHISISTIATIE

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.

>		osoft SQL Server 🕞	SeattleRetail (RDSAMAZ-FFIDFMR) Topology								
	Sei	arch by name			F		Ť	<u>_</u>	1		i
					Migrate Database	Clone Lifecycle	Remove Protection	Back up Now	Modify M	intenance D	Details
	1	Name	Manage Copies								
-		DWConfiguration	1 Backup					Su	mmary Card		
		DWDiagnostics	0 Clones								
â		DWQueue	Local copies						:kup		
Δ.		master	cocar copies					0 Clo	nes		
34		model									
		msdb									
		SeattleRetail	Primary Backup(s)								
A		tempdb	(search)							Ì. Perurre	- Cerr
			Backup Name	Count	Туре	15			End Date	Verified	
			RDSAMAZ-FFIDFMR_SeattleRetail_RDSAMAZ-FFIDFMR_03-29-2022_01.47.31.3117	1	Full backup			03/29/20	022 1:47:37 AM 🗮	Unve	erified

- 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.

2	jobs	Stredules	Ivents Logi					
Contributed	(in its name						
Resources	Jobs	Fitter						
Mantan .	iD.	Status	Name	Start date	first date	Owner		
d Reports	76	×.	Petiter RDAMA24EQ1M/Geatteleta7	10/29/2022 1 SH131 AAT	83/29/2022 1:58:26 AM	RESCUR	rowni, tedin	
	54	1	Backup of Resource Group NEEAMA2 Proceeds, Seattly/Repair with putry Trull Backup	DEGRAPHICS LATED AM 5	ERVERT AND A REPORT OF A REPOR	ADSCUST	CAPPAL LAD	autoritie.
🗛 Hota	.00	. 4	Onlin Residuce Group 105AWA2 (FIGENR) Seattleheta?	00/28/2022 1-45/24 AM	00000000014504 AM	RESCUS	OANNOLUMEN	witrator.
a Darage Dystams		10	Create Pulsy Pul Backage	03/29/2022 1-01:127 AM	-03/29/2022 1.4140.4M	RESCUST	CANNEL WORK	with the
E Sellings	ų.e	4	Discover resources for all hosts	05/25/0022 1:38:12 AM 5	05/29/2022 1.38-17 AM 🛱	RESCUE!	CREWKL WAR	entitle
	88	1	Discoverentation for hour HDDAW2 Proprietationsmall.com	33/29/2022 10:55 (7 HM 🗖	01/01/2022 10:55 18 PM	HISCUST	Chirol, Section	nitiatur
A Nata	87		Drussee resources for hold: HDGAMAD PFOHMErdousticesed aren	000992022104111894	00/28/2022 10:41 19 PM		Children and the	

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.

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.

The following step 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 follow 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.

	Microsoft SQL Server	SeattleRetail (RDSAMAZ-FFIDFMR) Topology								-		
	search by name			Migrate Database	Cione Lifecycle	Remove Protection		Modify Maire		1 Details		
2	🐺 Name	Manage Copies										
	DWConfiguration	1 Backup					Sum	mary Card				
1	DWDiagnostics	0 Clones					1 Backu					
4	DWQueue	Local copies					0 Clones					
۱	master						U CIONES					
	model											
-	msdb											
=	SeattleRetail	Primary Backup(s)										
2	tempdb	(search)) Ž Recame	Cone	H Restore	ti Del
		Backup Name	Count	Туре	17			End Date	Verified			
		RDSAMAZ-FFIDFMR_SeattleRetail_RDSAMAZ-FFIDFMR_03-29-2022_01.47.31.3117	1	Full backup			03/29/2022	1:47:37 AM	tim	rerified		

- 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.

Dashboard	searc	h by name	<mark> </mark>			
Resources	Jobs -	Filter				
S Monitor	ID	Status	Name	Start date	End date	Owner
Reports	108	~	Clone from backup 'RDSAMAZ-FFIDFMR_SeattleRetail_RDSAMAZ-FFIDFMR_03-29-2022_01.47.31.3117'	3/30/2022 6:09:10 PM 🛱	3/30/2022 6:09:55 PM 🛅	rdscustomval\administrato
and reports	107	~	Discover resources for all hosts	03/30/2022 6:06:40 PM 🛱	03/30/2022 6:06:54 PM 🛅	RDSCUSTOMVAL\administra

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. 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:

```
• _EnableMPI0.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
SschTrigger -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 = $hostname.Replace('-',' ')
    volsluns = Q()
    for ($i = 1;$i -lt 10;$i++) {
        if ($i -eq 9){
            $volsluns
+=(@{volname=('v '+$hostname+' log');volsize=$fsxN.loqvolsize;lunname=(
'l '+$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 aggr1
-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) {
        $sshcmd = "lun create -path /vol/" + $vollun.volname + "/" +
$vollun.lunname + " -size " + $vollun.lunsize + " -ostype Windows 2008
" #-vserver " +$vserver
        $return = Invoke-SSHCommand -Command $sshcmd -SSHSession
$session
        #map all luns to igroup
        $sshcmd = "lun map -path /vol/" + $vollun.volname + "/" +
$vollun.lunname + " -igroup " + $igroup
        $return = Invoke-SSHCommand -Command $sshcmd -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 EnableMPIO.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-asservice (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 <u>memory-optimized</u> 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.

Х

Create a volume

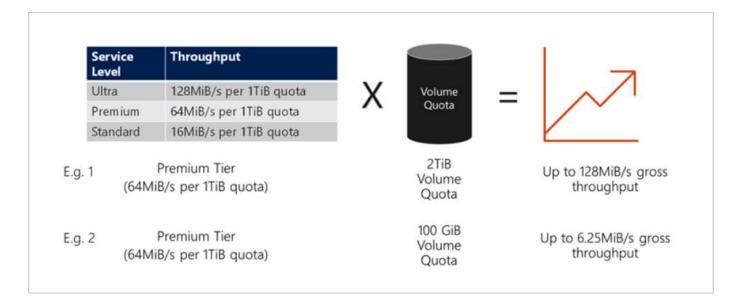
Basics Protocol Tags Review	+ create	
Configure access to your volume.		
Access	○ NFS ● SMB ○ Dual-protocol (NFSv3 and SMB)	
Protocol type		
Configuration		
Active Directory * ①	10.0.0.100 - anf.test/join	
Share name * 🕕	SQLDB	
Enable Continuous Availability ①		
Review + create < Preview	vious Next : 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

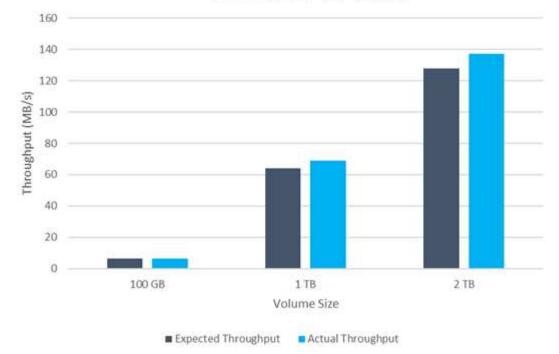
For more information, see Service levels for Azure NetApp Files.



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.

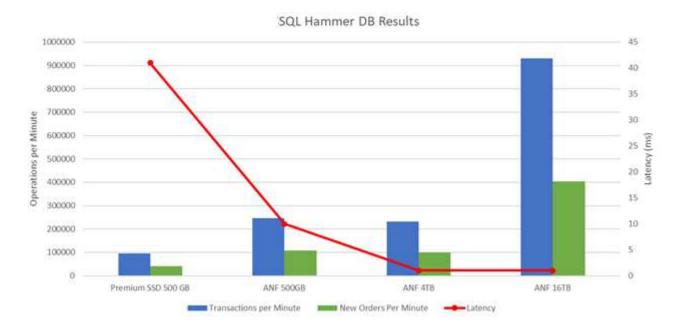


ANF Premium Tier Quotas

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.



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.

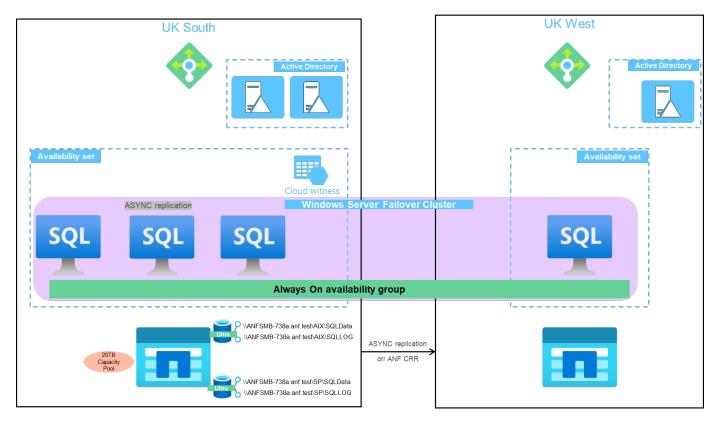
Real-time, high-level reference design

This section covers a real-time deployment of a SQL database estate in an AOAG configuration using an Azure NetApp Files SMB volume.

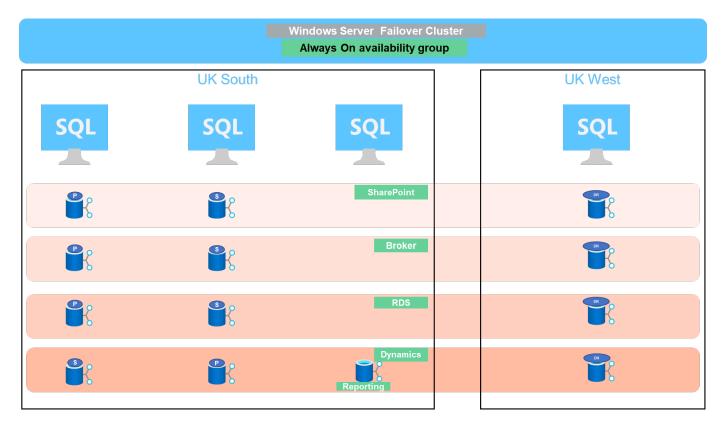
- Number of nodes: 4
- Number of databases: 21
- Number of availability groups: 4
- Backup retention: 7 days
- Backup archive: 365 days



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.

Add metric 🍾 Add	filter 🛛 👷 Apply splitting	🖄 Line chart 🧹	🕞 🕞 Drill into Logs 🧹 😲 New alert rule 🔗 Pin to dashboard
Scope	Metric Namespace	Metric	Aggregation
volume1	NetApp Volumes stand 🗸	Total throughput 🗸 🗸	🖌 Avg 🗸 🏈
100B/s		Percentage Volume Consume	ed Size
		Read iops	
80B/s		Read throughput	
60B/s		Total throughput	
40B/s		Volume allocated size	
20B/s		Volume Backup Bytes	

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 copied to Blob container of 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

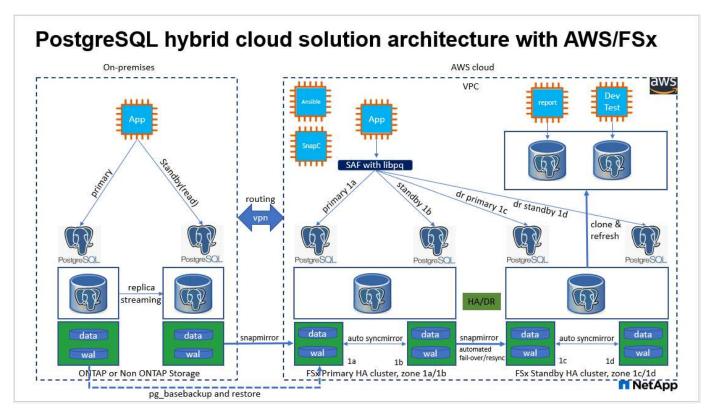
This solution is intended for the following people:

- 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	
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 alway 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-premisses 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 onpremises 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_01ps.pem
```

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.posgresql[1]}}" represents DB server 2, "{{groups.posgresql[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_archivelog_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

#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.postgresg1[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_kwcid=AL!4422!3!467723097970!e!!g!!aws%20ec2&ef_id=Cj0KCQiA54KfB hCKARIsAJzSrdqwQrghn6I71jiWzSeaT9Uh1-vY-VfhJixFxnv5rWwn2S7RqZOTQ0aAh7eEALw_wcB:G:s&s_kwcid=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-premisses 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/nsbb/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 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
```

```
# 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

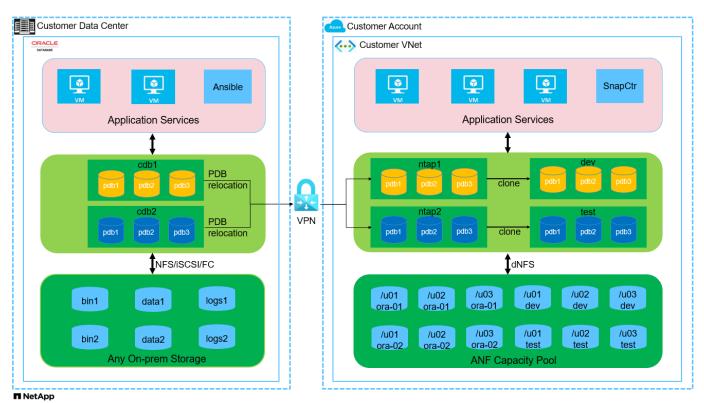
This solution is intended for the following people:

- 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	A capacity pool with Premium
	Microsoft	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_PD B2,NTAP1_PDB3)	/u01, /u02, /u03 NFS mounts on ANF capacity pool
ora-02	NTAP2(NTAP2_PDB1,NTAP2_PD B2,NTAP2_PDB3)	/u01, /u02, /u03 NFS mounts on ANF capacity pool

Key factors for deployment consideration

- **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 /tmplv and /rootlv OS partition accordingly. Ensure the database volume naming follows the VMname-u01, VMname-u02, and VMname-u03 convention.

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

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

- 3. From the Azure cloud portal, provision a Windows server to run the NetApp SnapCenter UI tool with the latest version. Refer to the following link for details: Install the SnapCenter Server.
- 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-premisses 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:

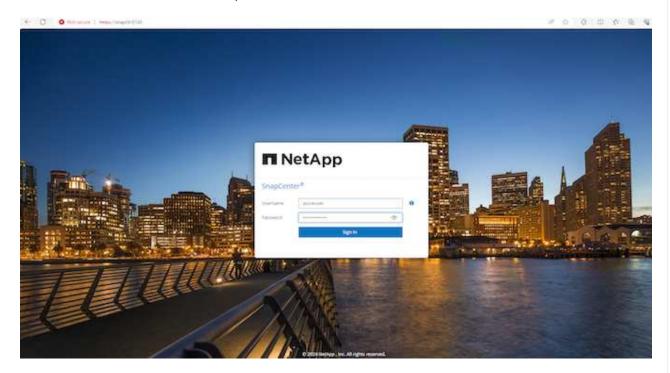
Oracle Database Backup, Recovery, and Clone on ANF with SnapCenter

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.

	Status Get Started		> ×
		Get Started	~ ^ ^
Dashboard		Add storage connections and licensing	~
Resources		L Configure user credentials	~
Monitor		Add a host & install plug-ins	
Reports		Reate policies	
Hosts	· · ·	Protect resources	
Storage Systems		Back up now	
Settings	Unable to connect to YouTube. You can use the playlist	Restore a backup	
Alerts	(https://www.youtube.com/playlist? list=PLdXI3bZJEw7nofM6lN44eOe4aOSoryckg) to view the videos.	Clone a backup	
	· · · · · ·	CA Certificate Settings	
		Backup to Object Store	~
		Learn more	

5. In Settings-Global Settings, check Hypervisor Settings and click on Update.

NetApp Snap	Center@	۰	? -	1 azureuser	SnapCenterAdmin	🖡 Sign Ou
	Clobal Settings Policies Users and Access Roles Credential Software					
Dashboard						
Resources	Global Settings					
S Monitor						
Reports	Hypervisor Settings					
📥 Hosts	VMs have ISCSI direct attached disks or NFS for all the hosts Update					
Storage Systems	Notification Server Settings ()					
Settings	Configuration Settings 0					
Alerts	Purge Jobs Settings 0					
	Domain Settings 0					
	CA Certificate Settings 0					
	Disaster Recovery 0					
	Audit log Settings 0					
	Multi Factor Authentication (MFA) Settings					

6. If needed, adjust Session Timeout for SnapCenter UI to the desired interval.

NetApp SnapC	Eenter® 🔮 🚱 🕹 aureuser SnapCenterAc	lmin 🖡 Sign
<	Calobal Settings Policies Ubers and Access Roles Credential Software	
Dashboard		
Resources	Global Settings	
Monitor		
Reports	Hypervisor Settings 0	
Hosts	Notification Server Settings 0	
- Storage Systems	Configuration Setting: 0	
E Settings	Session Timeout (in minutes) 20 Save	
Alerts		
	Purge Jobs Settings 0	
	Domain Settings	
	CA Certificate Settings 🚯	
	Disaster Recovery 0	
	Audi log settings 0	
	Multi Factor Authentication (MFA) Settings	

7. Add additional users to SnapCenter if needed.

П	NetApp Snap	Center®						•	2	9 · Lazureuser	SnapCenterAdmin	🖡 Sign Out
<		Global Settings		Users and Access		Software						
	Dashboard	Search by Nam	e				Type All •				+	
0	Resources		Name		11	Туре	Roles			Domain	Add	Ranow
✨	Monitor		azureuser			User	SnapCenterAdmin			localhost		
ñí	Reports											
Å	Hosts											
ł	Storage Systems											
	Settings											
▲	Alerts											

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.

<		Giobai	Settings Policies Users and Access Roles Credential Software		
	Dashboard	Searc	ch by Name		1 🕂 🦷 1 🦷 1
0	Resources		Name 1	Details	Add Copy R Members
۲	Monitor		SnapCenterAdmin	Overall administrator of SnapCenter system	1 User, No Groups
.	Reports		App Backup and Clone Admin	App Backup and Clone Admin	No Members
*	Hosts		Backup and Clone Viewer	Backup and Clone Viewer	No Members
h	Storage Systems		Infrastructure Admin	Infrastructure Admin	No Members
-	Settings				
▲	Alerts				

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.

NetApp SnapCe					•	🔤 😲 🗜 azur	euser SnapCe	nterAdmin	🗊 Sign (
	Global Settings Policies Users and Acce	ess Roles Credential	Software						
Dashboard	Search by Credential Name						+		
	Credential Name		Authentication Mode		Details		New	Mostly	
	azure_anf		AzureCredential						
	azureuser		Linux		Userid:azureuser				
Hosts									
Storage Systems									
Settings Alerts									
Crede	ential					×			
C	redential Name	azureuse	er						
			16						
Authe	ntication Mode	Linux			•				
Auth	entication Type	O Passwor	rd <mark>Ba</mark> sed 🔘 SSH	Key Based 🚺					
	Username	azureuse	er			0			
		XRIrK1QCa	EOHg==						
	SSH Private Key	<mark>END R</mark> S	5A PRIVATE KEY	ą.	*	0			
Use	e sudo privileges	6			4 (2009)				
USE USE	e sudo privileges	U							
				Cancel		ок			

Credential Name azure_anf Authentication Mode Azure Credential Azure Details Tenant ID Enter Tenant Id	
Azure Details ()	
Tenant ID Enter Tenant Id	
Client ID Enter Client Id	- C
Client Secret Key Enter client secret key	

10. From Storage Systems tab, add Azure NetApp Files with credential created above.

							•	S 0-	l azureuser	SnapCenterAdmin	🚺 Sign (
	ONTAP Sto	rage Azure NetApp Files									
Dashboard	Search b	y NetApp Account									
Resources		NetApp Account			11	Resource Group	Credential				
Monitor		ANFAVSAcct				ANFAVSRG	azure_anf				
Reports											
Hosts											
Storage Systems											
E Settings											
Alerts											
											_
Add Azure NetApp Ad	count										
			•	F 0							
Cradanti		re anf									
Credentia	azu	re_anf									
Credentia Subscriptio		re_anf prid Cloud TME Onprem		0							
Subscriptio	n Hyb	orid Cloud TME Onprem	-	0							
	n Hyb		-								
Subscriptio	n Hyb	orid Cloud TME Onprem	-	0							
Subscriptio	n Hyb	orid Cloud TME Onprem	-	0							
Subscriptio NetApp Accour	n Hyb	orid Cloud TME Onprem	-	0							
Subscriptio NetApp Accour	n Hyb	orid Cloud TME Onprem	-	0							
Subscriptio NetApp Accour	n Hyb	orid Cloud TME Onprem	-	0							
Subscriptio NetApp Accour	n Hyb	orid Cloud TME Onprem	-	0							
Subscriptio NetApp Accour	n Hyb	orid Cloud TME Onprem	-	0							

11. From Hosts tab, add Azure DB VMs, which installs SnapCenter plugin for Oracle on Linux.

Monitor Ora-01.hr2z2nbmhnoutdsxgscjt	uxizd.jx.internal.cloudapp.net	Type Linux	System Stand-alone	Plug-in UNIX, Oracle Database			Version 5.0	Overall St
Reports ora-02.hr2z2nbmhnoutdsxgsgit	uxrzd jy internal kloudago net	Linux	Stand-alone	UNIX, Oracle Database			5.0	Runn
Settings Alerts								
Add Host								
Host Type	Linux				•			
Host Name	ora-01							
Credentials	azureuser				+	0		
	L							
Select Plug-ins to Inst	all SnapCenter Plug-ins Pa	ckag	e 5.0 for	Linux				
Select Plug-ins to Inst	all SnapCenter Plug-ins Pa	ckag	e 5.0 for	Linux				
Select Plug-ins to Inst	Oracle Database	ckag	e 5.0 for	Linux				
Select Plug-ins to Inst	 Oracle Database SAP HANA 	ckag	e 5.0 for	Linux				
	 Oracle Database SAP HANA Unix File Systems 			Linux				
	 Oracle Database SAP HANA 			Linux				
More Options : Po	 Oracle Database SAP HANA Unix File Systems 			Linux				
	 Oracle Database SAP HANA Unix File Systems 			Linux				
More Options : Po	 Oracle Database SAP HANA Unix File Systems 			Linux				
More Options : Po	 Oracle Database SAP HANA Unix File Systems 			Linux				
More Options : Po	 Oracle Database SAP HANA Unix File Systems 			Linux				
More Options : Po	 Oracle Database SAP HANA Unix File Systems 			Linux				
More Options : Po	 Oracle Database SAP HANA Unix File Systems 			Linux				
More Options : Po	 Oracle Database SAP HANA Unix File Systems 			Linux				
More Options : Po	 Oracle Database SAP HANA Unix File Systems 			Linux				
More Options : Po	 Oracle Database SAP HANA Unix File Systems 			Linux				
More Options : Po	 Oracle Database SAP HANA Unix File Systems 			Linux				

Port	8145	(
Installation Path	/opt/NetApp/snapcenter	
	Skip optional preinstall checks	
	Add all hosts in the oracle RAC	
Custom Plug-ins	Choose a File	
	Browse	
	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-Polices, 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.

<		Global Settings Policies Users and Access Oracle Database	Roles Credential Software			
III D	Dashboard				-	/ 6
😨 R	Resources	Search by Name			+ New	
😍 N	Aonitor	Name IL	Backup Type	Schedule Type	Replication	Verification
ណ៍ R	Reports	Oracle archivelogs backup	LOG, ONLINE	Hourly		
		Oracle full online backup	FULL, ONLINE	Hourly		
A 1	losts					
l-1 s	torage Systems					
	iettings					
S S						
	Verts					
	Verts					
	Verts					

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.

<		Oracle	Database -									
	Dashboard		Database	Search database							\$	+
	Resources	lin.	Name	Oracle Database Type	Host/Cluster	Resource Group	Policies		Last Bac	1.0	Refresh Resources Overall Status	Add S
8 N	Monitor		NTAP1	Single Instance (Multitenant)	ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.interna						Not protected	d
M P	Reports		NTAP2	Single Instance (Multitenant)	l.cloudapp.net ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.interna						Not protected	
A (losts		NIAP2	single instance (Multitenant)	I.cloudapp.net						Not protected	,
- I s	itorage Systems											
🗉 s	iettings											
	Verts											

2. Click on View drop-down to change to Resource Group. Click on Add sign on the right to add a Resource Group.

NetApp Snap	Center®				•	•	- Lazu	euser SnapCen	terAdmin 🛛 🗍 Sign Ou
<	Oracle Database 👻								
Dashboard	View Resource Group	Search reso	urce group						t.
Resources	Name	Resources	Tags	Policies		L	ast Backup	Overall Status	
Monitor	There is no match for your sea	arch or data is not available.							Resource Group
Reports									
Hosts									
Storage Systems									
E Settings									
Alerts									

3. Name your resource group, tags, and any custom naming.

	esource Group	
	Name Resources Policies Verification Notification Summary	
	Provide a name and tags for the resource group	
	Name full_online_bkup f	
	Tags oradata 0	
	Use custom name format for Snapshot copy	
	\$HostName x	
	Backup settings	
	Exclude archive log	
	destinations from the second sec	
	Pr	evious
٨dd	l resources to your Resource Group. Grouping of similar resources can simplify database	
200	nagement in a large environment.	
nar		
nar		

	1 3 4 5 6 Name Resources Policies Verification Notification Summary
	Add resources to Resource Group
	Host
	All
	Available Resources Selected Resources
	Image: Addition of the control of t
	Previous Next
	T EVIDUS TEXT
5. Sele	ct the backup policy and set a schedule by click on '+' sign under Configure Schedules.

	3 4 5	6	
Name Resources P	olicies Verification Notificat	tion Summary	
Select one or more policies and co			
Oracle full online backup	• • 0		
Configure schedules for selected			
Policy Oracle full online backup	Li Applied Schedules	Configure Schedules	
Total 1			
			Previo
Add achodulas	for policy Ore	clo full opling hadres	×
Add schedules	for policy Ora	icle full online backup	~
Jourty			
Hourly			
Start date	02/06/2024	05:55 pm 🛛 🛱	
Start date	02/06/2024	05:55 pm	
Start date	02/06/2024		
<mark>Expires</mark> on	03/06/2024	05:51 pm	
	03/06/2024		
<mark>Expires</mark> on	03/06/2024	05:51 pm	
<mark>Expires</mark> on	03/06/2024	05:51 pm	
<mark>Expires</mark> on	03/06/2024	05:51 pm	
<mark>Expires</mark> on	03/06/2024	05:51 pm	
<mark>Expires</mark> on	03/06/2024	05:51 pm	
<mark>Expires</mark> on	03/06/2024	05:51 pm	
<mark>Expires</mark> on	03/06/2024	05:51 pm	
<mark>Expires</mark> on	03/06/2024	05:51 pm	
Expires on Repeat every	03/06/2024 2	05:51 pm	
Expires on Repeat every	03/06/2024 2	05:51 pm	ne
Expires on Repeat every	03/06/2024 2	05:51 pm	ne
Expires on Repeat every	03/06/2024 2	05:51 pm	ne
Expires on Repeat every	03/06/2024 2	05:51 pm nours 0 mins	
Expires on Repeat every	03/06/2024 2	05:51 pm	ne

6. If backup verification is not configured in policy, leave verification page as is.

0	2		<u></u>	•	5	6			
Name	Resour	es	Policies	Verification	Notification	Summary			
	e verification								
Policy	LE Sche		ta ia mataviaila	Applied Schedule	s	Configure Schedules			
i nere is no	match for you	search or dat	ta is not availa	ole.					
Total 0									
Total 0									
Total 0									
Total 0									
Total 0									
Total 0									
Total 0									
Total 0									
Total 0									
Total 0									
Total 0									
Total 0									

7. In order to email a backup report and notification, a SMTP mail server is needed in the environment. Or leave it black if a mail server is not setup.

	New Resource Group								×
	0	-			-	6			
	Name	Resources	Policies	Verification	Notification	Summary			
	Provide email : Select the service		le to notify regardir	ng protection issues.					
	Email preference	Never		•					
	From	From email							
	То	Email to							
	Subject	Notification							
	🗍 Attach Job repo	ort							
									Previous Next
~	•	~							
8.	Summary of	r new re	esource	group.					

				-			
Name	Resources	Policies	Verification	Notification	Summary		
Resource grou	ip name	full_onli	ne_bkup				
Tags		oradata					
Policy		Oracle fr	ull online backup: Hou	irly			
Plug-in		SnapCer	nter Plug-in for Oracle	Database			
Verification er	abled for policy	None					
Send email		No					

9. Repeat the above procedures to create a database archive log only backup with corresponding backup policy.

	Oracle Database					
Dashboard	View Resource Group	Search resource g	tront A			
Resources	Name	Resources	Tags	Policies	Last Backup	Ac Overall Status
Monitor	full_online_bkup	2	oradata	Oracle full online backup	02/06/2024 6:00:44 PM 📋	Completed
Reports	archivelog_bkup	2	oralog	Oracle archivelogs backup	02/06/2024 5:59:25 PM	Completed
I webore						
Hosts						
Hosts	5					
Hosts	5					

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.

Oracle Database 🕞	full_online_bkup Details						
Search resource groups	search			Modily Resource Group	L Back up Now	Maintenance	Defet
	Resource Name	Туре	Host				
Name	NTAP1	Oracle Database	ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net				
full_online_bkup	NTAP2	Oracle Database	ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net				
archivelog_bkup							

Resource Group	full_online_bkup		
Policy	Oracle full online backup	• 0	
🗌 Verify after back	up		

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'

- Internal.cloudapp.net
- ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net

 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

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.

Close

>	Oracle Database 🔹	full_online_bkup Details	NTAP1 Topology								
	Search resource groups	search								·	
0		Resource Name							Backup te	o Object Store Protect	
•	Name full_online_bkup	NTAP1	Manage Copies								
a il	archivelog_bkup	NTAP2	3 Backups						Summa	ary Card	
			0 Clones						3 Backu		
*			Local copies							ta Backup	
֥									2 Loj 0 Clone:	g Backups	
華										hots Locked	
A											
			Primary Backup(s)								
			search	7					0 Ž	Tin 1 A	
			Backup Name	Snapshot Lock Expiration	Count	Typel;	End Date	Verified	Mounted	RMAN Cataloged	
			ora-01_02-06-2024_18_00_ 06_0582_1		1	Log	02/06/2024 6:00:41 PM 🛱	Not Applicable	False	Not Cataloged	
			ora-01_02-06-2024_18_00_ 06_0582_0		1	Data	02/06/2024 6:00:26 PM 🛱	Unverified	False	Not Cataloged	3
			ora-01_02-06-2024_17_59_ 01_1158_1		1	Log	02/06/2024 5:59:18 PM 🛱	Not Applicable	False	Not Cataloged	
	Total 2		Total 3								

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	Cibiject Store Protect	Refresh
Manage Copies								
3 Backups						Summa	ary Card	
0 Clones						3 Backu	os	
Local copies						1 Dat	a Backup	
							g Backups	
						0 Clones		
						0 Snapsł	nots Locked	
Primary Backup(s)	Y				G	• t Lalog Rename	Cone Restore Mount	A To Unmount Delets
Backup Name	Snapshot Lock Expiration	Count	Typel;	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	Restore Scope 🚯	
2 Recovery Scope	All Datafiles	
3 PreOps	O Pluggable databases (PDBs)	
4 PostOps	O Pluggable database (PDB) tablespaces	
5 Notification	Control files	
6 Summary	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.	
		ext

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.

Restore NTAP1		×
1 Restore Scope	Choose Recovery Scope	
2 Recovery Scope	All Logs	
3 PreOps	 Until SCN (System Change Number) Date and Time No recovery 	
4 PostOps		
5 Notification	Specify external archive log files locations 🛛 🕒 🗐	1
6 Summary		0
	Previous Next	

 $\mbox{4. The $\tt PreOps$ allows execution of scripts against database before restore/recovery operation. } \label{eq:preOps}$

1 Restore Scope	Specify optional s	cripts to run l	before performing	a restore job 🚯		
2 Recovery Scope	Prescript full path	/var/opt/sn	apcenter/spl/scripts/	Enter Prescript path		
3 PreOps	Arguments					
4 PostOps	Script timeout	60	secs			
5 Notification						
6 Summary						
					Previous	Nex

	Restore NTAP1					×
	1 Restore Scope	Specify optional sc	ripts to run after performing a	restore job 🚯		
	2 Recovery Scope	Postscript full path	/var/opt/snapcenter/spl/scripts/	Enter Postscript path		
	3 PreOps	Arguments		a.		
	4 PostOps	Open the databas	e or container database in READ-W	RITE mode after recovery		
	5 Notification					
	6 Summary					
					Previous	xt
6.	Notification via e	mail if desired.			Previous	xt
6.	Notification via e	mail if desired.			Previous	xt
6.	Notification via e	mail if desired.			Previous	xt
6.	Notification via e	mail if desired.			Previous	xt
6.	Notification via e	mail if desired.			Previous	xt
6.	Notification via e	mail if desired.			Previous	xt
6.	Notification via e	mail if desired.			Previous	xt
6.	Notification via e	mail if desired.			Previous	xt
6.	Notification via e	mail if desired.			Previous	xt
6.	Notification via e	mail if desired.			Previous	xt
6.	Notification via e	mail if desired.			Previous	xt
6.	Notification via e	mail if desired.			Previous	xt
6.	Notification via e	mail if desired.			Previous	xt

1 Restore Scope	Provide email set	tings 🚯	
2 Recovery Scope	Email preference	Never -	
3 PreOps	From	From email	
4 PostOps	То	Email to	
5 Notification	Subject	Notification	
6 Summary	🗌 Attach job report		
a sumary			
		estore jobs, an SMTP server must be configured. Continue to the Summary page to save your Slobal Settings>Notification Server Settings to configure the SMTP server.	

7. Restore job summary

Restore NTAP1			>
1 Restore Scope	Summary		
2 Recovery Scope	Backup name	ora-01_02-06-2024_18_00_06_0582_0	
3 PreOps	Backup date	02/06/2024 6:00:26 PM	
o neops	Restore scope	All DataFiles	
4 PostOps	Recovery scope	All Logs	
5 Notification	Options	Change database state if necessary , Open the database or container database in READ-WRITE mode after recovery	
6 Summary	Prescript full path	None	
o cumury	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.

Resto	re 'ora-01.hr2z2nbmhnqutdsxgscjtuxizd	.jx.internal.cloudapp.net\NTAP1'
v v	Restore 'ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.	internal.cloudapp.net\NTAP1'
4	v ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.intern	ial.cloudapp.net
~	Prescripts	
~	Mount log backups	
~	Pre Restore	
~	▶ Restore	
~	▶ Post Restore	
~	Unmount log backups	
~	Postscripts	
~	Post Restore Cleanup	
~	Data Collection	
	Name: ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.ir /06/2024 4:08:42 PM	nternal.cloudapp.net Start Time: 02/06/2024 4:04:55 PM End
		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.

NTAP1 Topology								×
						Backup to	Cibject Store Protect	Refresh
Manage Copies								
3 Backups						Summa	ry Card	
0 Clones						3 Backup	05	
Local copies							a Backup	
							Backups	
						0 Clones		
						Ushapsi	iots Locked	
Primary Backup(s)								
search	v				Car	alog Rename	Cone Restore Mount	≜ ∰ Delete
Backup Name	Snapshot Lock Expiration	Count	Typel₹	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 NT/			
1 Name	Capacity Pool Max. Throughput (MiB/s)		0
2 Locations	O Complete Databa	se Clone	
3 Credentials	Clone SID	ntap1dev	
4 PreOps	Exclude PDBs	Type to find PDBs	
5 PostOps	O PDB Clone		
6 Notification			
7 Summary			
			Previous
			· · · · · · · · · · · · · · · · · · ·

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.

	Select the host to create	e a clone						
Locations	Clone host ora	-02,hr2z2nbmh	nqutdsxgsc	tuxizd.jx.inter •				
Credentials	⊙ Datafile locations 1							
PreOps	/u02_ntap1dev					÷		Reset
PostOps								
Notification	⊙ Control files ()							
Summary	/u02_ntap1dev/ntap1dev	/control/control	101.ctl			\mathbf{x}	*	+
-	/u02_ntap1dev/ntap1dev	/control/control	102.ctl			×	~	Reset
	💮 Redo logs 🚯							
	Group		Size	Unit	Number of files			
		×	Size 200	Unit	Number of files		+	*
	Group	×	-				+	+ Reset
	Group • RedoGroup 1		200	МВ	1			
	Group RedoGroup 1 RedoGroup 2	×	200 200	MB	1		+	
	Group RedoGroup 1 RedoGroup 2	×	200 200	MB	1		+	
	Group RedoGroup 1 RedoGroup 2	×	200 200	MB	1		+	

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 N	TAP1		×
1 Name	Database Credentials fo	or the clone	
2 Locations	Credential name for sys user	None - + 3	
3 Credentials	Database port	1521	
4 PreOps			
5 PostOps	Oracle Home Settings	0	
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. 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.

	Specify scripts to r	run before clone op	eration 🕕			
Locations	Prescript full path	/var/opt/snapcente	r/spl/scripts/ Enter Presc	ript path		
Credentials	Arguments					
PreOps	Script timeout	60 secs				
PostOps	🕑 Database Parame	ter settings				
Notification	processes		320		×	*
Houncadon	remote_login_pass	wordfile	EXCLUSIVE		×	+
Summary	sga_target		3G		×	Reset
	undo_tablespace		UNDOTBS1		\times	v

6. 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).

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

7. Notification via email if desired.

1 Name	Provide email set	tings 🕕	
2 Locations	Email preference	Never •	
3 Credentials	From	From email	
PreOps	То	Email to	
5 PostOps	Subject	Notification	I
6 Notification	🗌 Attach job report		
7 Summary	4		
		or Clone Jobs, an SMTP server must be configured. Continue to the Summary page to save your gs>Global Settings>Notification Server Settings to configure the SMTP server.	

8. Clone job summary.

1 Name	Summary					
2 Locations	Clone from backup	ora-01_02-06-2024_18_00_06_0582_0				
	Clone SID	ntap1dev				
3 Credentials	Capacity Pool Max. Throughput (MiB/s)	none				
4 PreOps	Clone server	ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net				
	Exclude PDBs	none				
5 PostOps	Oracle home	/u01/app/oracle/product/19.0.0/NTAP2				
6 Notification	Oracle OS user	oracle				
	Oracle OS group	oinstall				
7 Summary	Datafile mountpaths	/u02_ntap1dev				
	Control files /u02_ntap1dev/ntap1dev/control/control01.ctl /u02_ntap1dev/ntap1dev/control02.ctl					
	Redo groups	RedoGroup =1 TotalSize =200 Path =/u02_ntap1dev/ntap1dev/redolog/redo01_01.log RedoGroup =2 TotalSize =200 Path =/u02_ntap1dev/ntap1dev/redolog/redo02_01.log RedoGroup =3 TotalSize =200 Path =/u02_ntap1dev/ntap1dev/redolog/redo03_01.log				
	Recovery scope	Until Cancel				
	Prescript full path	none				
	Prescript arguments					
	Postscript full path	none				
	Postscript arguments					
	Send email	No				
		Previous Fir	nis			

9. 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.clouda	app.net Start Time: 02/06/2024 6:21:5	9 PM End
Time: 02/06/2024 6:28:10 PM		
	View Logs Cancel Job	Close

10. Cloned database registers with SnapCenter immediately.

<	Oracle Dat	abase 🝷						
Dashboard	View Da	itabase	Search databases	V				* +
Resources	17 14	Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Refresh Resources Add Overall Status
Monitor		NTAP1	Single Instance (Multitenant)	ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.i	archivelog_bkup	Oracle archivelogs backup	02/06/2024 7:29:18 PM	Backup succeede
Reports	-	ntap1dev	Single Instance (Multitenant)	nternal.cloudapp.net ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.i	full_online_bkup	Oracle full online backup		Not protected
Hosts				nternal.cloudapp.net				
Storage System	5	NTAP2	Single Instance (Multitenant)	ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.l nternal.cloudapp.net	archivelog_bkup full_online_bkup	Oracle archivelogs backup Oracle full online backup	02/06/2024 7:29:19 PM 🛱	Backup succeede
Settings								
Alerts								

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 1.6G 0 1.6G tmpfs 0% /run/user/0 172.30.136.68:/ora-01-u02-Clone-020624161543077 250G 8.2G 242G

```
4% /u02 ntap1dev
[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 $ORACLE SID are not allowed.
#
#
NTAP2:/u01/app/oracle/product/19.0.0/NTAP2:Y
# SnapCenter Plug-in for Oracle Database generated entry (DO NOT
REMOVE THIS LINE)
ntap1dev:/u01/app/oracle/product/19.0.0/NTAP2:N
[oracle@ora-02 ~]$ export ORACLE SID=ntap1dev
[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 -
Production
Version 19.18.0.0.0
```

SQL> select name, open_mode, log_mode from v\$database; NAME OPEN_MODE LOG_MODE _____ ARCHIVELOG NTAP1DEV READ WRITE 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 Fixed Size9100040 bytesVariable Size654311424 bytesDatabase Buffers2550136832 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 pdbs OPEN MODE RESTRICTED CON ID CON NAME _____ ____ READ ONLY NO 2 PDB\$SEED 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 cloudbackup, 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:

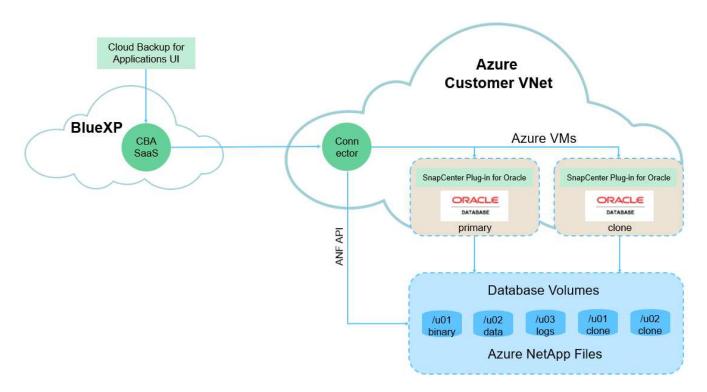
- The DBA who manages Oracle databases running on Azure NetApp Files storage
- The solution architect who is interested in testing Oracle database backup, restore, and clone in Azure
- The storage administrator who supports and manages the Azure NetApp Files storage
- The application owner who owns applications that are deployed to Azure NetApp Files storage and Azure

VMs

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
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.
- 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.

letApp BlueXP	Q BlueXP Search	ch Account V Workspace Automation-te Azure-DB	Connector V 🏚 🌣 🤉 😫
Canvas My working environments My estate			Go to Tabular View
+ Add Working Environment		C Enable Services (1)	Working Environments
			Amazon S3
			Amazon S3 O Buckets
			Azure Blob Storage 20 Storage Accounts
			20 Storage Accounts
Azure Blob Storage	Amazon S3		
Storage Accounts	Buckets		
		-+	

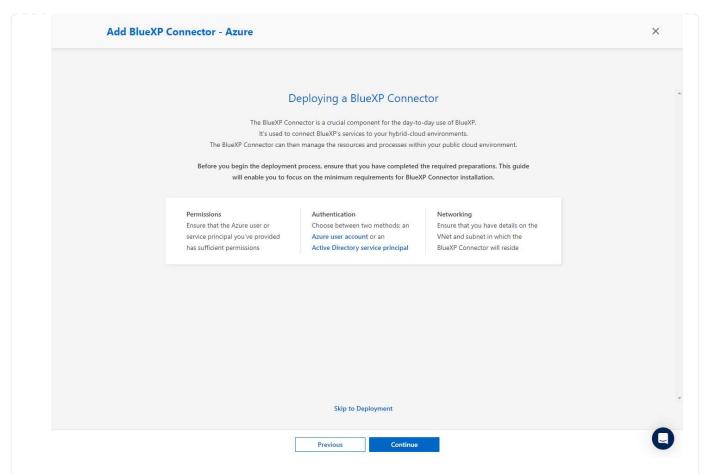
2. Click on **Connector** drop down arrow and **Add Connector** to launch the connector provisioning workflow.

NetA	pp BlueXP			Q BlueXP	Search	Account ~ Automation-te	Workspace Azure-DB	~ 0	Connector 🖌	` 4 0	¢ (9 6
•	Canvas	My working environments	My estate					Connectors	A	dd Connector	Manage Co	onnector
2	+ Add Working Envir	onment				C Enable Service	es O	Q Search F	BlueXP Connect	ors		
2								acao-av	vs-connector		Go to Local U	7
									us-east-1 🔳 Ir	nactive		
9								AzureCo	onnector		Go to Local U	1.2
0									southcentralus		do to cocar o	
		Azure Blob Storage 20		Amazon S3								
		Storage Accounts		Buckets	aws							
						-	+	S	Switch		Cancel	-(

3. Choose your cloud provider (in this case, **Microsoft Azure**).

Add BlueXP Conn	ector			
		Provider		
		oud provider where you want to run the Bl	JeXP Connector:	
		aws	٥	
	Microsoft Azure	Amazon Web Services	Google Cloud Platform	
		Deploy the Connector on your premises	2	
		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."



 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.

VM Authentication (2) Details	3 Network (4) Security Group (5) Review
Virtual Mac	nine Authentication
You are logged in with Azure user: acao@netapp.c	com 🖉 Tenant: Hybrid Cloud TME 🗸
Subscription	
Hybrid Cloud TME Onprem	Authentication Method
Location South Central US	
Resource Group	User Name azureuser
Create New 💿 Use Existing	Enter SSH Public Key
Resource Group	BEGIN RSA PRIVATE KEY MIIGSAIBAAKCA
ANFAVSRG	

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.

VM Authentication 2 Details	(3) Network (4) Security Group (5) Review	
	Details	
Connector Instance Name	0	
AzureConnector		
	Create	
Add Tags to Connector Instance	Role Name	
	BlueXP Operator-5519248	
	Subscriptions to apply with the role	
	Hybrid Cloud TME Onprem \sim	
		C
Frevious	Next	
onnector - Azure		
		More Information X
VM Authentication	3 Network (4) Security Group (5) Review	More Information X
	Network A Security Group S Review	More Information X
		More Information X
VM Authentication 🔗 Details	Network	More Information X
⊘ VM Authentication	Network Proxy Configuration (Optional)	More Information X
© VM Authentication () Details Connectivity VNet	Network Proxy Configuration (Optional) HTTP Proxy	More Information X
♥ VM Authentication ♥ Details Connectivity VNet ANFAVSVal	Network Proxy Configuration (Optional) HTTP Proxy	More Information X
VM Authentication O Details Connectivity VNet ANFAVSVal Subnet VM_Sub VM_Sub	Network Proxy Configuration (Optional) HTTP Proxy Example: http://172.16.254.1:8080 Define Credentials for this Proxy ~	More Information X
VM Authentication Image: Details Connectivity VNet ANFAVSVal Image: Details Subnet VM_Sub Public IP Image: Details	Network Proxy Configuration (Optional) HTTP Proxy Example: http://172.16.254.1:8080 Define Credentials for this Proxy ~	More Information X
VM Authentication O Details Connectivity VNet ANFAVSVal Subnet VM_Sub VM_Sub	Network Proxy Configuration (Optional) HTTP Proxy Example: http://172.16.254.1:8080 Define Credentials for this Proxy ~ Upload a root certificate ~	More Information X
VM Authentication Image: Details Connectivity VNet NFAVSVal Image: Details Subnet Image: Details VM_Sub Image: Details Public IP Disable	Network Proxy Configuration (Optional) HTTP Proxy Example: http://172.16.254.1:8080 Define Credentials for this Proxy ~ Upload a root certificate ~ ty	More Information
✓ VM Authentication ✓ Details Connectivity VNet ANFAVSVal Subnet VM_Sub VM_Sub Public IP Disable Notice: Ensure that the subnet has internet connectivit through a NAT device or proxy server so that the Context or the Context or the Context	Network Proxy Configuration (Optional) HTTP Proxy Example: http://172.16.254.1:8080 Define Credentials for this Proxy ~ Upload a root certificate ~ ty	More Information
✓ VM Authentication ✓ Details Connectivity VNet ANFAVSVal Subnet VM_Sub VM_Sub Public IP Disable Notice: Ensure that the subnet has internet connectivit through a NAT device or proxy server so that the Context or the Context or the Context	Network Proxy Configuration (Optional) HTTP Proxy Example: http://172.16.254.1:8080 Define Credentials for this Proxy ~ Upload a root certificate ~ ty	More Information
✓ VM Authentication ✓ Details Connectivity VNet ANFAVSVal Subnet VM_Sub VM_Sub Public IP Disable Notice: Ensure that the subnet has internet connectivit through a NAT device or proxy server so that the Context or the Context or the Context	Network Proxy Configuration (Optional) HTTP Proxy Example: http://172.16.254.1:8080 Define Credentials for this Proxy ~ Upload a root certificate ~ ty	More Information
✓ VM Authentication ✓ Details Connectivity VNet ANFAVSVal Subnet VM_Sub VM_Sub Public IP Disable Notice: Ensure that the subnet has internet connectivit through a NAT device or proxy server so that the Context or the Context or the Context	Network Proxy Configuration (Optional) HTTP Proxy Example: http://172.16.254.1:8080 Define Credentials for this Proxy ~ Upload a root certificate ~ ty	More Information
✓ VM Authentication ✓ Details Connectivity VNet ANFAVSVal Subnet VM_Sub VM_Sub Public IP Disable Notice: Ensure that the subnet has internet connectivit through a NAT device or proxy server so that the Context or the Context or the Context	Network Proxy Configuration (Optional) HTTP Proxy Example: http://172.16.254.1:8080 Define Credentials for this Proxy ~ Upload a root certificate ~ ty	More Information
✓ VM Authentication ✓ Details Connectivity VNet ANFAVSVal Subnet VM_Sub VM_Sub Public IP Disable Notice: Ensure that the subnet has internet connectivit through a NAT device or proxy server so that the Context or the Context or the Context	Network Proxy Configuration (Optional) HTTP Proxy Example: http://172.16.254.1:8080 Define Credentials for this Proxy ~ Upload a root certificate ~ ty	More Information
ſ	AzureConnector Add Tags to Connector Instance	AzureConnector Cancetor Role Atach existing Manual BlueXP Operator-5519248 Obscriptions to apply with the role Uybrid Cloud TME Onprem

8. Configure the **Security Group** for the connector that allows HTTP, HTTPS, and SSH access.

VM Authentication	Details O Network Sec	curity Group 5 Review	
	Security Group		
The securit	y group must allow inbound HTTP, HTTPS and	SSH access.	
Assign a security group: 💿 Creat	e a new security group O Select an exist	ting security group	
HTTP (Port 80)	HTTPS (Port 443)	SSH (Port 22)	
Source Type	Source Type	Source Type	
Anywhere	Anywhere 🗸	Anywhere 🗸	
Source (CIDR)	Source (CIDR)	Source (CIDR)	
0.0.0.0/0	0.0.0/0	0.0.0/0	

9. 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.

Add BlueXP Connecto	or - Azure		More Information	>
	VM Authentication	🛇 Details 🕢 Network 🔗 Security Group 🚺 Review		
		Review		
	BlueXP Connector Name			
	Subscription	Hybrid Cloud TME Onprem		
	Location	South Central US		
	Resource Group	Existing - ANFAVSRG		
	Role	New - BlueXP Operator-5519248		
	Authentication Method	Password (user: azureuser)		
	VNet	ANFAVSVal		
	Subnet	VM_Sub		
	Public IP	Enable		
	Proxy	None		
	Security Group	HTTP: 0.0.0/0, HTTPS: 0.0.0.0/0, SSH: 0.0.0.0/0		
	Г	Previous Add		

10. After the connector is deployed, the newly created connector appears under **Connector** drop-down.

•	O Canvas	My working environments My estate			Go to Tabular View
2	+ Add Working Environ	iment		C Enable Services 🚯	Working Environments
0				Second of Contraction	
					Amazon S3 O Buckets
					Azure Blob Storage 20 Storage Accounts
:					
		Azure Blob Storage	Amazon 53		
		20 Storage Accounts	0 Buckets		
				(- +)	

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.

NetApp	BlueXP			Q. BlueXP Search Account ✓ Automation-team	Workspace Azure-D8	e Connector AzureConnector	* 0 0
<u>ې</u>	Credentials	Account credentials	User credentials			🚯 Settings	
				BlueXP and the Connector use account-level credentials to deploy and manage resources in your cloud environment.		Connector Settings	
۰ ۵		3 с	redentials	and manage resources in your cloud environment.	Ac	Timeline	
•			aws DemoFSxNCMCredentials Type: Assume Role BlueXP			Credentials	
			982589175402	DhruvCloudManagerRole		Software Update	
			AWS Account ID	Assume Role		HTTPS Setup	
			aws shantanucreds Type: Assume Role BlueXP			Alerts and Notifications Settings	
			210811600188 AWS Account ID	nkarthik, kafka_nfs_role_FSxN Assume Role			
			Managed Service Identity Type: Managed Service Ident				
			1 View Subscriptions	0 Working Environments			
							C
							E

2. Choose credential location as - Microsoft Azure - BlueXP.

n NetApp BlueXP		Q BlueXP Search	Account ~ Automation-team	Workspace 💙 Azure-D8	Connector Y AzureConnector	4 0 🌣	0	8
Add Credentials							×	
9	Choose Creder	ntials Location						
•		aws						
°	Microsoft Azure	Amazon Web Services						
•	Choose how to associ	ate the credentials 🕚						
		(•					
	Connector	BlueXP						
		~						
	Ne	art .					(
	IVE	A1						

3. Define Azure credentials with proper **Client Secret**, **Client ID**, and **Tenant ID**, which should have been gathered during previous BlueXP onboarding process.

III N	etApp	BlueXP		Q BlueXP Search	Account ~ Automation-team	Workspace V Azure-DB	Connector ~ AzureConnector	1 0 🌣	8
	0	Add Credentials	Credentials Type 2 Define Crede	entials (3) Marketpla	ce Subscription (4) Revie	ew			×
9			Define Microsoft	t Azure Credentials					
				re application credentials					
6			Credentials Name	Client Secret					
۲			Azure_Hybrid_TME						
**			Application (client) ID 2fbc9be5-a259-4539-bb57-036b176f5c	Directory (tenant) ID 9bb0aab6-5c98-419b	o-9cfd-7a38bd496				
			 I have verified that the Azure role assigned BlueXP policy requirements. 	to the Active Directory service of the service of t	ice principal matches				
									0
			Previous	Next					U

4. Review and **Add**.

III Ne	etApp BlueXP		Q BlueXP Search Account Automat	✓ Connector ✓ AzureConnector	4 0 🌣	9
	Add Credentials	Credentials	Type ⊘ Define Credentials 3 Review			×
0			Review			
۰						
Ġ		Credentials Type	Azure			
۲		Credentials Name	Azure_Hybrid_TME			
•		Credential Storage	Cloud Manager			
		Application (client) ID	2fbc9be5-a259-4539-bb57-036b176f5cc7			
		Directory (tenant) ID	9bb0aab6-5c98-419b-9cfd-7a38bd496e1f			
			in the second seco			C
		Previous	Add			

5. You may also need to associate a **Marketplace Subscription** with the credential.

n Ne	tApp BlueXP	Q BlueXP Search Account V Workspace Connector Acure-D8 Acure-D8 Acure-D8	¢ ()	6
	Associate Subscription		2	×
9	Credentials > Associate Subscription			
•		Associate a Marketplace Subscription		
Ģ		Choose an Azure subscription to associate with the Azure Marketplace subscription.		
۲		Azure Subscription		
•		Hybrid Cloud TME Onprem (default)		
		Marketplace Subscription		
		dhruv-cvo-subscription		
		Add Subscription		
		Associate Cancel		

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.

INetApp BlueXP		Account Y Workspace Automation-te Azure-DB	Connector Connector
Canvas My working environments My estate			🖽 Go to Tabular View
+ Add Working Environment		C Enable Services 🚯	Working Environments
			Amazon S3 0 Buckets
			Azure Blob Storage 20 Storage Accounts
			20 Storage Accounts
Azure Blob Storage	Amazon S3		
Storage Accounts	Buckets		
		- +	(

2. Choose **Microsoft Azure** as the location and click on **Discover**.

Add Working Environment	Choose a Location	999¢⊗ ×
	Microsoft Azure Amazon Web Services Google Cloud Platform On-Premises	
9 8	Cloud Volumes ONTAP Discover Existing Add new	
	Cloud Volumes ONTAP HA Discover Existing Add new	
	Kubernetes Cluster Discover	

Continue.

NetA	op BlueXP		Connector 🖌 🍋 🔅 ?	8
e 9	Add Azure NetApp Files Wizard The sector of the sector of	Quarter Net App Files Credentials More Net App Files Credentials More Net App Files Credentials More Net App Files Credentials More Net App Files Credentials More Net App Files Credentials More Net App Files Credentials More Net App Files Credentials More Net App Files Credentials More Net App Files Credentials More Net App Files Credentials More Net App Files Credentials More Net Net App Files Credentials More Net Net App Files Credentials More Net Net Net Net Net Net Net Net Net Ne	aveConvector	×
		Continue		C

4. BlueXP console returns to **My working environments** and discovered Azure NetApp Files from Azure now appears on **Canvas**.

NetApp BlueXP	Q Blue20° Search Automation te Workspace Connector AzureConnector Q Blue20° Search Automation te
Canvas My working environments My estate	Go to Tabular View
+ Add Working Environment	Enable Services Working Environments Image: Account of the service of the se
Azure Blob Storage 20 Storage Accounts	0 Buckets

5. Click on **Azure NetApp Files** icon, then **Enter Working Environment** to view Oracle database volumes deployed in Azure NetApp Files storage.

NetApp BlueXP		Q BlueXP Search Automation-team Azure-D8	Connector AzureConnector
9	Azure NetApp Files AzureNfile		(i)
,			
,	ora01-u01 AVAIL	ABLE ···· Ora01-u02	AVAILABLE
•	INFO CAPACITY Service Level Premium Location South Central US Protocols INFSv2	GIB INFO CAPACITY GIB Service Level Premium Location South Central US Provisioned Provisioned	0 GB Used Capacity
	ora01-u03	AILABLE Est	AVAILABLE
	INFO CAPACITY Service Level Premium Location South Central US Protociti NF5/3	GIB GIB d Capacity Capa	0.08 Used Capacity
	Service Level Premium Location South Central US Source Control US	I GIB d Capacity Service Level Standard Location Germany West Central Dono GiB	0 Citi Uned Capacity

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**.

<section-header><section-header><section-header><section-header><section-header><section-header><text><text><text><text><text></text></text></text></text></text></section-header></section-header></section-header></section-header></section-header></section-header>	rts	berneties Job Monituring Ties	mes Restore AppReations Virtual Machi	Backup and recovery	•
Enterprise Applications Integrated Data protection & Copy management service for on-premises and cloud workloads BuektP backup and recovery delivers quick, seamless, and cost effective backup, restore, and copy management capabilities for enterprise databases boated on both on-premises and in the cloud (CMPP, Acure NetApp ONTAP),					
Enterprise Applications Integrated Data protection & Copy management service for on-premises and cloud workloads IhuestP backup and recovery delivers, quick, seamless, and cost- effective tackup, restore, and copy management capabilities for endudt (Cloud volumes, OMTAP), store NetApp files and Amazon FSx for NetApp ONTAP) (storage. Get started with Cloud Backup for Applications by discovering applications. Discover Applications	C ###	for	BlueXP backup and reco		
Integrated Data protection & Copy management service for on-premises and cloud workloads BluetP backup and recovery delivers quick, seamless, and cost- effective backup, restore, and copy management sapablities for enterprise databases hosted on both ni-premises and in the doud (Cloud volumes ONTAP, Azure NetApp files and Amazon FSx for Wedpp CMTAP) storage. Gert started with Cloud Backup for Applications by discovering applications.	1 3 1 2 1 3 1 3 1 3 1 3 1 3 1 3 1 3 1 3	■ ٤.	Enterprise Applications		
service for on-premises and cloud workloads IllustRP backup and recovery delivers quick, seamles, and cost- effective taskup, restore, and copy management capabilities for enterptise databases boated on both on greenises and in the doud (Cloud volumes ONTAP, Azore NetApp files and Amazon F5x for NetApp ONTAP) storage. Get started with Cloud Backup for Applications by discovering applications. Docour Applications		nent	Integrated Data protection & Copy n		
IllueXP backup and recovery defivers quick, seamless, and cost- effective tackup, restore, and copy management capabilities for enterprise databases boated on both on premises and in the doud (Cloud volumes ONTAP, Azure NetApp files and Amazon FSx for NetApp ONTAP) starage. Get started with Cloud Backup for Applications by discovering applications.	And And And And And And And And And And	i interest	and the second second second second second second second second second second second second second second second		
effective backup, restore, and copy management sapabilities for enterprise databases booted on both on premises and in the cloud (Cloud volumes CNTRP, Azure NetKpp files and Amazon FSx for NetApp ONTAP) storage. Get started with Cloud Backup for Applications by discovering applications.		Sector Sector	The set has been and as a set of the set of the		
cloud (Cloud volumes ONTAP, Azure NetApp files and Amazon FSx for NetApp ONTAP) storage. Get started with Cloud Backup for Applications by discovering applications.					
for NetApp ONTAP) starage. Get started with Cloud Backup for Applications by discovering applications.					
Get started with Cloud Backup for Applications by discovering applications. Discover Applications		azon FSx			
applications. Discover Applications			in workpo na na transfer		
Otycover Applications		ring	Get started with Cloud Backup for Applications		
			applications.		
			Discours Accelerations		
S (5)			C Indexedentities Anno construction		
		Ø	0		
	0	(S.)	\odot		
Streamlined data management Save time & resources Protect data in minutes	Protect data in minutes	ave time & resources	Streamlined data management		
Manage your cloud native Automated workflows without Paster backup and restore applicators with one console downtime save organizational operations help you to meet					

7. Select **Cloud Native** as the application source type.

		Q BlueXP Search Account ~ Automation-team	Workspace Connector Azure-D8 AzureConnector	l 🔹 🌣 🛛 🖯
8	Select Appli	ication Source Type		
9	Select the application source to			
•				
ক				
۲	Hybrid	Cloud Native		
*8	Applications hosted within your organization's infrastructure.	Applications that are hosted and run in the cloud using AWS, Azure, GCP, etc.,		
	Cancel	Next		0

8. Choose **Oracle** for the application type, click on **Next** to open host details page.

n N	etApp BlueXP		Q BlueXP Search Account Automation-team	Workspace Azure-D8	Connector Y AzureConnector	4 0 🌣	9
2	Discover Applications						×
6							
٠		Select Ap	plication Type				
\$							
@ •*		ORACLE	SAP HANA				
**		Uracle	SAP HANA				
							0
			Next				Q

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.

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 O Manual 🕕 🐵 Using SSH 🕥	
	Host FQDN or IP Connector	
	172.30.137.142 AzureConnector 💌	
	Username	
	azureuser 🕒 Add SSH Private Key Optional 🕥	
	SSH Port Plug-In Port	
	22 8145	
iscover Applications	1 Host Details (2) Configuration (3) Review	
iscover Applications	Host Details Ornfiguration 3 Review	
iscover Applications	Host Details 2 Configuration 3 Review Select host type	
iscover Applications		
iscover Applications	Select host type	
scover Applications	Select host type Provide the following details to add host and discover applications	
scover Applications	Select host type Provide the following details to add host and discover applications Host Installation Type O Manual Using SSH Using SSH	
iscover Applications	Select host type Provide the following details to add host and discover applications Host installation Type O Manual O O Using SSH O Validate fingerprint	
iscover Applications	Select host type Provide the following details to add host and discover applications Host installation Type Manual Image: Colspan="2">Image: Colspan="2">Colspan="2" Colspan="2">Colspan="2">Colspan="2" Colspan="2">Colspan="2" Algorithm ssh-rsa	
iscover Applications	Select host type Provide the following details to add host and discover applications Host installation Type Manual Image: signarran Algorithm signarran Pingerprint: AAAAE2VjZHNhLXNoYTItbmlzdHAyNTYAAAAIbmlzdHAyNTYAAAAIbmlzdHAyNTYAAAAIbmlzdHAyNTYAAAAIbmlzdHayNTyAAAIbmlzdHayNTYAAAAIbmlzdHayNTyAAAIbmlzdHayNTYAAAAIbmlzdHayNTyAAAIbmlzdHayNTyAAAAIbmlzdHayNTyAAIbmlzdHayNTyAAIbmlzdHayNTyAAIbmlzdHayNTyAAIbmlzdHayNTyAAIbmlzdHayNTyAAIbmlzdHayNTyAAIb	
iscover Applications	Select host type Provide the following details to add host and discover applications Host installation Type Manual Image: Colspan="2">Using SSH Validate fingerprint Algorithm ssh-rsa Fingerprint: AAAAE2VJ2HNhLXNoYTitbmizdHAyNTYAAAAIbmizdHAyNTYAAAB Image: Colspan="2">Image: Colspan="2" Image: Colspan	
iscover Applications	Select host type Provide the following details to add host and discover applications Host installation Type Manual Image: Colspan="2">Using SSH Validate fingerprint Algorithm ssh-rsa Fingerprint: AAAAE2VJ2HNhLXNoYTitbmizdHAyNTYAAAAIbmizdHAyNTYAAAB Image: Colspan="2">Image: Colspan="2" Image: Colspan	
iscover Applications	Select host type Provide the following details to add host and discover applications Host installation Type Manual Image: Colspan="2">Using SSH Validate fingerprint Algorithm ssh-rsa Fingerprint: AAAAE2VJ2HNhLXNoYTitbmizdHAyNTYAAAAIbmizdHAyNTYAAAB Image: Colspan="2">Image: Colspan="2" Image: Colspan	
iscover Applications	Select host type Provide the following details to add host and discover applications Host installation Type Manual Image: Colspan="2">Using SSH Validate fingerprint Algorithm ssh-rsa Fingerprint: AAAAE2VJ2HNhLXNoYTitbmizdHAyNTYAAAAIbmizdHAyNTYAAAB Image: Colspan="2">Image: Colspan="2" Image: Colspan	
Discover Applications	Select host type Provide the following details to add host and discover applications Host installation Type Manual Image: Colspan="2">Using SSH Validate fingerprint Algorithm ssh-rsa Fingerprint: AAAAE2VJ2HNhLXNoYTitbmizdHAyNTYAAAAIbmizdHAyNTYAAAB Image: Colspan="2">Image: Colspan="2" Image: Colspan	

10. Move on to next **Configuration** page to setup sudoer access on Oracle Azure VM.

n Ne	etApp BlueXP	Q BlueXP Search Account * Workspace * Connector * Account Activitianition team Account team
2	Discover Applications	Host Details (2) Configuration (3) Review
9		
•		Configuration
6		Follow the steps to make sure all the configuration expectations are met
0		1. Configure sudoer access for "azureuser".
••		1. Log into the application host. 2. Create following file /etc/sudoers.d/snapcenter with the following content.
		# #
		I have configured sudo access for "azureuser" as per the above steps.
		Previous
		Previous Next

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.

letApp BlueXP			vrkspace V Connector V ire-D8 AzureConnector	40 \$ 6 6
Discover Applications	Host Deta	ils 🕜 Configuration 3 Review		
		Review		
	Follow the steps to ma	ake sure all the configuration expectations are met.		
	Host Details	Configurations		
	Host Installation Type	SSH		
	Host FQDN or IP	172.30.137.142		
	Connector	AzureConnector		
	User name (Sudo)	azureuser		
	Plug-In Port	8145		
	SSH Port	22		
	Fingerprint	AAAAE2VJZHNhLXNoYTItbmlzdHAyNTYAAAAIbmlzdH		
	Кеу Туре	ecdsa-sha2-nistp256		
	Previo	Us Discover Applications		C

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**.

Cloud Native						
	▼ Oracle	*				
3 Hosts	BD 3 ORACLE	Clone	Application Protection Image: Operating the second secon	9.3 Unprotected		
	 Host Name 	Q Policy Name	Manage Databases	Settings 🔻		
NTAP	172.30.137.142		0 Unprotected			
db1	172.30.15.99		0 Unprotected			
db1tst	172.30.15.124		9 Unprotected			
			1 - 3 of :	< < 1 > >>		
	Hosts Hosts Filter By + Name NTAP db1	Hosts ORACLE 3 Databases Fitter By + Name A Host Name NTAP 172:30.137.142 db1 172:30.15.99	Hosts ORACLE 3 Databases Filter By Name NtAP 172.30.137.142 db1 172.30.15.99	Image: Detabases Image: Detabases Filter By ▲ Max Mame Policy Name Protected NTAP 172:30:137.142 Image: Durprotected db1 172:30:15.599 Image: Durprotected db1tst 172:30:15.124 Image: Durprotected	3 BB 3 BB 3 Protected Unprotected 3 Detabases Cone 0	3 BB 3 BB 0 0 3 Detabases BITER FBY ● 0 <td< td=""></td<>

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	08	/dev
tmpfs	7.9G	0	7.9G	08	/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	28	/tmp
/dev/sda1	500M	6.8M	493M	28	/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 ~]\$					

 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.

NetApp	BlueXP		Q BlueXP Search	Account Y Automation-te	Workspace 💙 Azure-DB	Connector ✓ AzureConnector	lo 🌣	?
	Backup and recovery	Volumes Restore Applications	Virtual Machines Kubernetes	Job Monitoring	Reports			
,								
	Cloud Native	Oracle	*					
	-				Application Pro	otection		
	4 Hosts	BB 3 ORACL	E	O	O Protected	9 3 Unprotected		
	3 Databases							
	Filter By			Q	Manage Databases	▼ Settings	•	
	Name	∧ │ Host Name	Policy Name		Protection Status	¢ 1		
	NTAP	172.30.137.142			Unprotected	•••		
	db1	172.30.15.99			() Unprotected	View Details Assign Policy		
	db1tst	172.30.15.124			Unprotected			
						1-3 of 3 << < 1	> >>	

2. When you are happy with the policy configuration, you can then **Assign** your policy of choice to protect the database.

🗖 NetAp	P BlueXP			Q BlueXP Search Account Workspace Connector Automation-te Azure-DB AzureConnector
<i>a</i> (Backup and recovery	y Volumes Restore Application	s Virtual Machines Kubernetes	Job Monitoring Reports
9	Applications > Assign Po	Policy		
•			Assig	n Policy
ବ			Assign a policy to start taking	g backups of the database "NTAP"
۲		4 Policies		
•:		Policy Name	🔨 📔 Backup Type	0 Schedules
		Oracle Full Backup for Bronze	FullBackup	Daily: Repeats Every 1 Day, Keeps 14 copies Weekly: Repeats Every Fri, Keeps 4 copies Monthly: Repeats Every 1st Day of Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, C
		Oracle Full Backup for Gold	FullBackup	Hourly: Repeats Every 6 Hrs, Keeps 16 copies Daily: Repeats Every 1 Day, Keeps 30 copies Weekly: Repeats Every Fri, Keeps 4 copies Monthly: Repeats Every 1st Day of Jan, Feb, Mar, Apr, May, Jun, Jul, Aug. Sep. 0
		Oracle Full Backup for Silver	FullBackup	Hourly: Repeats Every 12 Hrs, Keeps 6 copies Daily: Repeats Every 1 Day, Keeps 14 copies Weekly: Repeats Every Fri, Keeps 4 copies Monthly: Repeats Every 1st Day of Jan, Feb, Mar, Apr, May, Jun, Jui, Aug, Sep, C
		my_full_bkup	FullBackup	Hourly: Repeats Every 6 Hrs. Keeps 3 Days
				1-4of4 << <mark>1</mark> > >
			Cancel	Assign

 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.

NetApp	BlueXP						Q BlueXP Sea		Automation-te	Workspace Azure-DB		AzureConnecto
0	Backup and recovery	Volumes	Restore	Applications	Virtual Machines	Kubernetes	Job Monitoring	Reports				
		Cloud Native		*	Oracle		•					
					_				Application	n Protectio	n	
		3 Hosts			BB 3 ORACLE			one	1 Protected		0 2 Unprote	ected
		2 Databases										
		3 Databases Filter By +		Host Nag	70		Policy Name	Q				ettings T
			^	Host Nan 172.30.13			Policy Name	٩	Manage Databases		¢	ettings v
		Filter By +	^		17.142			Q	Protection Statu	IS View Deta	o I. Is	•••
		Filter By + Name NTAP	~	172.30.13	.99			Q	Protection Statu	IS View Deta On-Demai	ils	•••

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.

n Ne	tApp BlueXP						Q BlueXP Se	sarch	Account ~ Automation-te	Workspace Azure-DB	~	Connector Y AzureConnector
	Backup and recovery	Volumes	Restore Ap	plications	Virtual Machines	Kubernetes	Job Monitoring	Reports				
9	Job Monitoring > Job Name: Backup	of NTAP oracle databa	se on host 172.30.137	7.142 with polic	y my_full_bkup and sc	hedule Hourly						
•		1.	Job Name: Bac	kup of NT			172.30.137.142 wi		my_full_bkup	and schedul	e H	
¢								0050				
۲			0		0				6			
-:			Other Job Type		Jul 11 2023, 2:1 Start Time	7:53 pm	Jul 11 2023, 2 End Time	:21:38 pm	-	Success Status		
		Sub-Jobs(17)										Collapse All
		Job Name		\$	i Job ID	‡	Start Time	‡∣ End	Time	Duration		≎∣ 🖨
		Backup of NT.	AP oracle database or	n host 172.30	🗇 61a12139-33	0e-4390-bc	Jul 11 2023, 2:17:53 pm	Jul	11 2023, 2:21:38 pm	4 Minute	5	
		Apply	ing Retention		27ff9d5f-68f0)-4880-a48	Jul 11 2023, 2:21:38 pm	Jul	11 2023, 2:21:38 pm	0 Second	1	
		Perfo	rming cleanup after b	backup	074c0689-09	7e-41aa-ac	Jul 11 2023, 2:21:36 pm	Jul	11 2023, 2:21:38 pm	2 Second	is	
		Finali	zing Oracle database	log backup	348189d3-90	b5-4cce-97	Jul 11 2023, 2:21:36 pm	Jul	11 2023, 2:21:36 pm	0 Second	1	
		1										

5. From three-dot drop down menu **View Details**, you can view the backup sets created from snapshot backup.

	App BlueXP					Q BlueXP Search	A	utomation-te	Azure-DB		AzureConnector
	Backup and recovery	v Volumes	Restore Applications	Virtual Machines	Kubernetes .	Job Monitoring	Reports				
,		Cloud Native		Oracle		-					
s -		Cloud Native									
)								Application	n Protection	n	
:		4 Hosts		3 ORACLE		Clone		2 Protected		0 1 Unprote	ected
		3 Databases									
		3 Databases Filter By +					٩	Manage Database	s ∣ v	Se	ettings 🔻
			∧ │ Host Name		1	Policy Name	Q	Manage Database		Se 0	ettings 🔻
		Filter By 🕂	 Host Name 172:30.137.1 		1. A.	Policy Name	Q			• 1	ettings ▼
		Filter By +		142			Q	Protection Statu	IS View Deta	• •	••
		Filter By + Name NTAP	172.30.137.1	9		Smy_full_bkup	Q	Protection Statu	5	I	••
		Filter By + Name NTAP db1	172.30.137.1	9		Smy_full_bkup	٩	Protection Statu Protected Protected	View Deta On-Demai	ils	••

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.

	G Backu	p and recovery Volumes	Restore Applications Vir	rtual Machines Kubernet	es Job Monitoring	Reports	
	Appli	cations > Database Details					
•				Database Detai	ls		
<u>ہ</u>		NTAP Database Name	Protected Protection	my_full_bk Policy Nam		Database Type	
•		172.30.137.142 Host Name	ANF Host Storage	Unreachat Database V		zEHlu7vkdyaBnujcxlibkk Connector Id	KELkVXToyNIclients
		- Clones	- Parent Database	<mark>Disabled</mark> RMAN Cata	log	- RMAN catalog repository	0
		14 Backups					
		Filter By +				۹ (Select Timeframe 🔻
		Backup Name	🗧 📔 Backup Type	0 SCN 0	RMAN Catalog	Backup Time 🗸 🗸	1
		my_full_bkup_Hourly_NTAP_2023_07_13_1	2_04_28_8376 Log	29192187	Not Cataloged	Jul 13, 2023, 8:06:22 am	Delete
		my_full_bkup_Hourly_NTAP_2023_07_13_1	2_03_07_4363 Data	29192136	Not Cataloged	Jul 13, 2023, 8:03:40 am	Delete
		my_full_bkup_Hourly_NTAP_2023_07_13_0	6_04_28_5618 Log	29178022	Not Cataloged	Jul 13, 2023, 2:05:50 am	Delete
				29177972			

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.

NetApp	BlueXP			Q BlueXP Search	Account ~ Automation-te	Workspace Connector Azure-DB AzureConnector
Storage	Backup and recovery	Volumes Re	store Applications Virtual Mag	chines Kubernetes Job Monitoring F	Reports	
9 Health						
Protection						
Backup and rec		Cloud Native	• Oracle	*		
					Applicatio	n Protection
Governance		4 Hosts		RACLE Clone	≥ 2 Protected	1 Unprotected
Mobility						
 Mobility Extensions 		3 Databases				
		3 Databases Filter By +			Q. Manage Database	
			Host Name	Policy Name		ss ∣ ¥ Settings ¥
		Filter By 🕂	 Host Name 172,30,137,142 		Q. Manage Database	ss ∣ ¥ Settings ¥
		Filter By +		Policy Name	Q Manage Database	ss ▼ Settings ▼ us
		Filter By + Name NTAP	172.30.137.142	Policy Name ≅ my_full_bkup	Q Manage Database Protection State Protected	ss ▼ Settings ▼ us 0 View Details On-Demand Backup
		Filter By + Name NTAP db1	172.30.137.142	Policy Name ≅ my_full_bkup	Q Manage Database Protection State Protected Protected	ss ▼ Settings ▼ us 0 View Details On-Demand Backup

2. Choose your **Restore Point** by time stamp. Each time stamp in the list represents an available database backup set.

etApp BlueXP		connector 🖌 🎝 🎝 🤹 🤉 😫
Restore "NTAP"	Restore Point and Location (2) Configuration (3) Review	
	Restore Point and Location	
	Specify the restore point to which the database should to be restored.	
	Restore Point 0	
	Jul 13, 2023, 8:03:40 am	
	Jul 13, 2023, 8:03:40 am	
	Jul 13, 2023, 2:03:43 am	
	Jul 12, 2023, 8:03:41 pm Jul 12, 2023, 2:03:32 pm	
	Jul 12, 2023, 2:03:31 am	
	location	
	Previous Next	

3. recovery.

n Ne	etApp BlueXP	Q Bluet/# Search Account V Workspace V Connector V Automation-te Accure-D8 Accessor	🗶	¢ 6	8
3	Restore "NTAP"	Restore Point and Location (2) Configuration (3) Review			
9					
•		Restore Point and Location Specify the restore point to which the database should to be restored.			
õ		specify the resore point to which the database should to be restored.			
۲		Restore Point			
••		jul 13, 2023, 8:03:40 am 🖤			
		Restore to original Restore to alternate location location			
					C
		Previous Next			

4. Define your **Restore Scope**, and **Recovery Scope**. All Logs mean a full recovery up to date including current logs.

n Ne	tApp BlueXP	Q Blued# Search Account * Workspace * Connector * 40 🔅 ?
8	Restore "NTAP"	Restore Point and Location 2 Configuration 3 Review
ଅ ଜ ଜ *	Restore Sc	All Cata Files Restore Control Files Control Files Restore O Database state will be changed if needed for restore and recovery.
		External Archive log locations mnt/log_location001 Open the database or the container database in READ-WRITE mode after recovery.
		Previous Next

5. Review and **Restore** to start database restore and recovery.

II Ne	etApp BlueXP		Q Bluett? Search Automation-te Workspace Y Connector Y Azure-D6 Azure-Connector Q	¢ 0	8
8	Restore "NTAP"	\odot	estore Point and Location 🥥 Configuration 🛛 🟮 Review		>
9					
•			Review		
Ģ		Backup Name	my_full_bkup_Hourly_NTAP_2023_07_13_12_03_07_43633_0		
0		Restore Scope	All Data Files		
••		Recovery Scope	All Logs		
		Force in Place Restore	No		
		Open Database or Contail Database in READ-WRITE Mode After Recovery	er Yes		
					-
			Previous Restore		C

6. From the **Job Monitoring** tab, we observed that it took 2 minutes to run a full database restore and recovery up to date.

	Backup and recover	y Volu	imes Restore	Applications	Virtual	Machines K	ubernetes	Job M	onitoring	Reports	-						
	0			at as so to													
	Job Monitoring > Job Name: Rest	ore for Oracle Data															
•			Job Name	Restore for (Jracle L				name my_ 868-9f279f8302		ip_Houri	y_NTAP_	2023_07	_13_12			
Ģ																	
۲			0			0			0			~					
•:			Other			ul 13 2023, 10:37:4	2 am		Jul 13 2023, 10:	39:15 am		🕑 Sue					
			Job Type		S	itart Time			End Time			Job Stat	tus				
		Sub-Jobs(6)													Collapse All 🔨		
		Job Name			≎ Job	ID	\$	Start Time		≎ En	d Time		🗘 Durat	ion	÷ 🔁		
		Resto	re for Oracle Database	NTAP using backu.	. 0	80882740-952d-4	acd-b	Jul 13 202	3, 10:37:42 am	Jul	13 2023, 10:	:39:15 am	2 Min	utes			
			Post Restore Cleanu	ip	Ø	0533d58b-7750-4	Oc1-a	Jul 13 202	3, 10:39:14 am	Jul	13 2023, 10:	:39:15 am	1 Seo	ond			
			Post Restore		Ø	64262431-041c-4	:21-8d	Jul 13 202	3, 10:38: <mark>4</mark> 8 am	Jul	13 2023, 10:	:39:14 am	26 Se	conds			
			Restore		C	918ad669-af04-41	7e-89	Jul 13 202	3, 10:38:24 am	Jul	13 2023, 10:	:38:48 am	24 Se	conds			
																	1

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.

NetApp	BlueXP							eXP Search	Automatic	n-te	Workspace Azure-DB	Ň	Connecto AzureCon	
Storage	Backup and recovery	Volumes	Restore App	olications	Virtual Machines	Kubernetes	Job Monitori	ing Report	's					
9 Health														
Protection														
Backup and rec		Cloud Native		Ŧ	Oracle		¥							
		_			_					Applicatio	n Protectio	n		
Governance		<u> </u>			88 3		88	0		2		<mark>9</mark> 1		
		Hosts			ORACLE			Clone		Protected		Unpro	tected	
		1033			ORACLE									
Mobility		1030			ONACLE									
		3 Databases			UNALL									
					UNACLE				Q Mar	nage Database	es 🔻	5	Settings	•
		3 Databases	~	Host Name			Policy Name			hage Database		: :	GH 2003/	1 7
		3 Databases Filter By +	•	Host Name 172.30.137.1			│ Policy Name ⊜ my_full_t		Pro			0.1	GH 2003/	1 🔻
		3 Databases Filter By + Name	•		142			e	Pro	otection State		ils	Settings] 🔻
		3 Databases Filter By + Name NTAP		172.30.137.1	142		Smy_full_t	e	Pro	otection State	View Deta On-Dema	≎ ils nd Backu	Settings	1 •
		3 Databases Filter By + Name NTAP db1	*	172.30.137.1 172.30.15.99	142		Smy_full_t	e	Pro	Protection State	us View Deta On-Dema	o ills nd Backu	Settings	

2. Select the **Restore Point** and check the **Restore to alternate location**.

III Ne	etApp BlueXP	Q BlueXP Search	Account ~ Automation-te	Workspace V Azure-DB	Connector Y AzureConnector	4 0 🌣	?	8
-	Restore "NTAP"	1 Restore Point and Location (2) Configuration	3 Review					×
Q								
•		Restore Point and Location Specify the restore point to which the database should to be	restand					
Ģ		Specify the restore point to which the database should to be	restored.					
۲		Restore Point	0					
••		Jul 13, 2023, 8:03:40 am	· ·					
		Restore to original location						
		Previous Next						0

3. In the next **Configuration** page, set alternate **Host**, new database **SID**, and **Oracle Home** as configured at alternate Azure VM.

n Ne	etApp BlueXP	Q Blueze Search Account V Workspace V Connector Acure-D8 Acure-C08
8	Restore "NTAP"	Restore Point and Location 2 Configuration 3 Review
9		
•		Configuration Specify the alternate host details on which the database will be restored and throughput.
6		speciry une ortennate nost octans on which the orazonase will be restored and dinoughput.
۲		Host SID 172.30.137.147 v NTAP1
0 ⁰ 0		Oracle Home Database Credentials Optional /u01/app/oracle/product/19.0.0/clone Image: Add Credential Image: Add Credential Maximum storage throughput (MIB/s) Optional Image: Add Credential Image: Enter throughput (1-4500) Image: Add Credential Image: Add Credential
		Previous Next Q

4. Review **General** page shows the details of cloned database such as SID, alternate host, data file locations, recovery scope etc.

🗖 Ne	tApp BlueXP			Q BlueXP Search	Account 💙 Automation-te	Workspace 🛩 Azure-DB	Connector ✓ AzureConnector	-	¢ (• •
	Restore "NTAP"	Ø	Restore Point and Location	Configuration	Review					×
9										
٠			Re	view						
¢		Ge	neral	Databas	e parameters					
Θ		Backup Name	my_full_bkup_Hourly_NTAF	P_2023_07_13_12_03_07_43633	_0					
••		SID	NTAP1							
		Host	172.30.137.147							
		Datafile locations	/u02_NTAP1							
		Control files	/u02_NTAP1/NTAP1/contro	ol/control01.ctl						
		Redo logs	RedoGroup = 2 TotalSize =	1024 Path = /u02_NTAP1/NTA 1024 Path = /u02_NTAP1/NTA 1024 Path = /u02_NTAP1/NTA	P1/redolog/redo02_01.ld	g				
		Recovery scope	Until cancel using selected	l backup's archive logs						
		Recovery Point	Jul 13, 2023, 8:03:40 am							
		Location	Alternate Location							
			Previous	Restore						C

5. Review **Database parameters** page shows the details of cloned database configuration as well as some database parameters setting.

n Ne	tApp BlueXP			Q BlueXP Search Account V Automation-te	Workspace V Azure-DB	Connector ✓ AzureConnector	40 🌣	?	8
Ð	Restore "NTAP"	\odot	Restore Point and Location	Configuration 3 Review					
,									
			Re	eview					
		Ger	eral	Database parameters					
		Database Credentials	None						
		Oracle home	/u01/app/oracle/product/	19.0.0/clone					
		Oracle OS user	oracle						
		Oracle group	oinstall						
		DB parameters	audit, file_dest = /u01/app audit_trail = DB open_cursors = 300 pga_aggregate_target_in_ processes = 320 remote_login_passwordfil sga_target_in_mb = 9216 undo_tablespace = UNDO	e = EXCLUSIVE					
			Previous	Restore					(

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.

Net	tApp BlueXP					Q Blu	eXP Search	Account Automation-te		orkspace cure-DB	*	Connector AzureConnector	~ 4 0	¢	?
•	🕝 Backup	and recovery	Volumes	Restore	Applications	Virtual Machines	Kubernetes	Job Moni	itoring F	leports					
,	Job Monitoring >	Job Name: Restore	Oracle Database NTAP	as NTAP1 on ho	ost 172.30.137.147 (using backup my_full	_bkup_Hourly_NTA	P_2023_07_13_1	12_03_07_4363	3_0					
,			Job Name: Re	store Oracl	e Database N	Job Id: 7a187d5a-7			using bacl	cup my	full_b	kup_Hourl			
			\bigcirc		0		0			~					
•:			Other Job Type		Jul 13 202 Start Time	3, 1:05:02 pm	Jul 13 2 End Tin	023, 1:13:15 pm 1e		-	Success				
		Sub-Jobs(6)											Collapse All	^	
		Job Name		;	Job ID	\$	Start Time	\$	End Time		‡	Duration	÷ 🔂		
		Restore C	Dracle Database NTAP as	NTAP1 on ho	🗇 7a187d5	a-7f7e-461a-83	Jul 13 2023, 1:05:0	12 pm	Jul 13 2023, 1	13:15 pm	l	8 Minutes			
		c	Collect the restore databa	ise job logs of	abc9342;	a-5777-4262-b	Jul 13 2023, 1:13:	4 pm	Jul 13 2023, 1	13:14 pm	2	0 Second			
		F	legister the restored data	abase metadata	D 15aefb90	0-b21b-418f-b0	Jul 13 2023, 1:12:	0 pm	Jul 13 2023. 1	12:30 pm	į	0 Second			
		F	lemove the temporary st	orage of the I	🗇 cc106fb9	9-7555-46c8-9c	Jul 13 2023, 1:12:	0 pm	Jul 13 2023, 1	13:14 pm	3	44 Seconds			
		1													

7. Validate the cloned database in BlueXP **Applications** page that showed the cloned database was immediately registered with BlueXP.

G	Backup and recovery Volume	es Restore Applications Virtual Ma	achines Kubernetes Job Monitori	ng Reports	
	Cloud Native	Oracle			
				Application Protect	ion
	<u> </u>	88 4	0 88	 2 	9 2
	Hosts	ORACLE	Clone	Protected	Unprotected
	4 Databases				
	Filter By +		Q	Manage Databases 🔻	Settings 🔻
	Name	A Host Name	Policy Name	Protection Status	0/1
	_	▲ Host Name 172.30.137.142	Policy Name Se my_full_bkup		• 1
	Name			Protection Status	
	Name NTAP	172.30.137.142		Protection Status	

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:
    SORACLE SID: SORACLE 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$database;
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 cloudbackup, 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

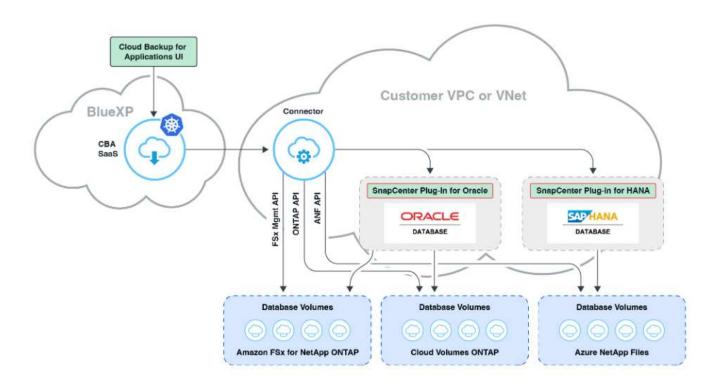
Audience

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.



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.

aws Services Q Search				[Alt+S]		🗘 🕜 Global y
🥘 Resource Groups & Tag Editor 🛛 🚾 FSx				· · · · · · · · · · · · · · · · · · ·		
Identity and Access Management (IAM)	Policies > snap					
Dashboard			Policy AR	N arn:aws:iam	::541696183547:policy/snapc	eenter 街
 Access management 			Descriptio	on Policy to gra	nt snapcenter service permis	sion to create connector in AWS.
User groups		- "				
Users	Permissions	Policy usage	Tags	Policy versions	Access Advisor	
Roles	Policy sum	mary {}JSON	Edit p	olicy		
Policies						
Identity providers	2 * {	"Version": "20	312-10-1	7.".		
	3*	"Statement":				
Account settings	4.*	{				
 Access reports 	5		ct": "All on": [Low",		
Access analyzer			iam:Creat	teRole",		
Access analyzer			iam:Delet			
Archive rules				plePolicy",		
Analyzers			iam:Creat	teInstanceProf teRolePolicy",	ile",	
10000000000000000000000000000000000000				pleToInstanceP		
Settings				veRoleFromInst		
Credential report				teInstanceProf		
credential report		10	iam:Pass	Role",		
Organization activity			iam:ListM			
Service control policies (SCPs)				ribeInstanceSt	atus",	
Service control policies (Ser 3)				nstances",	21	
				fyInstanceAttr		
Q Search IAM				teSecurityGrou teSecurityGrou		
				ribeSecurityGr		
				<pre>closecurityGrou </pre>		
AWS account ID:				prizeSecurityG		
541696183547				prizeSecurityG		
041030100047			ec2:Revol	ke5ecurityGrou	pIngress",	
		19	ec2:Creat	teNetworkInter	face",	
	9.0			aihoNotwonkInt		

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.

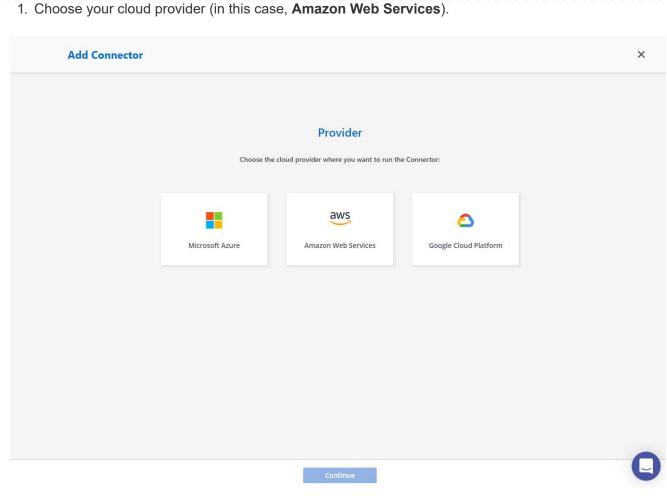
3. You also need the AWS VPC, subnet, security group, an AWS user account access key and secrets, an SSH key for ec2-user, and so on ready for connector provisioning.

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.

Manag	e Account: Automation-team	Overview	Members	Workspaces	BlueXP Connector	×
manag		Overview	Members	workspaces	SideXF Connector	~
	Manage the Blue	eXP connector Workspace	s			
		+ Ad	d New Workspace			
	Database		1			
	Database-2		1			
	sufians-k8		1			
	Workspace-1		1			
Click Add a Con	nector to launch the connector	provisioning wo	orkflow.			
App Cloud Manager		Account	Workspace	✓ Connector N/A	···	98
	Volumes Restore Applications Virtual Machines	Automation-team Kubernetes Job Monitoring	new-workspace	N/A		
(G) Backup & Restore	volumes restore Applications virtual machines	Kubernetes Job Monitoring				
	Backup & Restore			12 2,011	112.25 m to a barrier to a second to a se	D+
	Fully integrated data protection for O		2,017 team	nage (The folget state and the state of the	Q	Tena Maria
	and unstructured data across your ONTAP hybrid cloud en			Secret, Secret	lanen Jak I. an 21 100 laten 2013 spann 💮 en	· ·
	effective object storage. All you need to do is select the sou	irce the target and the	100	have, where, 1	Terrer, NW, J. Bay 31322, M.B.B.B. US Segment 2 in 1	ingen

-	Backup & Restore	President Insures States
9	Fully integrated data protection for ONTAP anywhere	12 2011 112.25 m 2014 Constraints Constraints <thconstraints< th=""> Constraints Constraits</thconstraints<>
۲	Cloud Backup dramatically reduces the complexity of backing up critical structured	2.011 Ansar Q Use strap Ministrating processing 1 Scale strap 4 ministration 1 Scale strap 7 ministration Scale strap Scale strap 6 Scale strap
	and unstrucutred data across your ONTAP hybrid cloud environments to cost-	 Marine 1 Marine 1 Marine 1 Marine 1 Marine 2 Marine 1 Marine 2 Marine 1 Marine 2 Marine 2<
0	effective object storage. All you need to do is select the source, the target and the protection policy and you're protected	ment Cart and and and and and and and and and and
ø	To start your Backup & Restore experience, please deploy our connector	61B 15 1.271 B 1 124 26.25 T8 5 19995 -
6	Add a Connector	
6		
0		
đ		
-		\$.)
(iii)	Simple & intuitive Hybrid Multicloud	Unmatched Efficiency
۲	No backup or cloud expertise required. Simply click Backup from On-premises or Cloud Volu the button above and follow the instructions ONTAP to AWS, Azure, GCP or StorageG	



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."

Add Connector -	AWS		×
	Deploying a Cont The Connector is a crucial component for the day It's used to connect Cloud Manager's services to y The Connector can then manage the resources and process Before you begin the deployment process, ensure that you have c will enable you to focus on the minimum requirem	-to-day use of Cloud Manager. our hybrid-cloud environments. es within your public cloud environment. ompleted the required preparations. This guid	ie
	Permissions Authentication Set up an IAM role with the Choose between two AWS equired permissions authentication methods: AW or assuming an IAM role	Networking Obtain details about the VPC a subnet in which the Connector reside	
	Skip to Deploymen	ıt	
	Previous	Continue	
. Enter your AWS ac	Previous	Continue Key and Secret Key.	C
. Enter your AWS ac Add Connect	count authentication with Access I		More Information X
	count authentication with Access I	Key and Secret Key.	
	count authentication with Access H or - AWS	Key and Secret Key. Nork (4) Security Group (5) Revi	
	or - AWS AWS Credentials C Details C Deta	Key and Secret Key. Nork (4) Security Group (5) Revi	
	or - AWS AWS Credentials AWS Credentials Context (Context) Context) Context) Context) Context) Con	Key and Secret Key.	
	or - AWS AWS Credentials AWS Credentials Context (Context) Context) Context) Context) Context) Con	Key and Secret Key.	
	or - AWS AWS Credentials AWS Credentials Control Contro Control Control Contr	Key and Secret Key.	

- Add Connector AWS More Information × AWS Credentials 2 Details (3) Network (4) Security Group (5) Review Details Connector Instance Name 0 0 Connector Role SnapCenterSvs Oreate Role O Select an existing Role Role Name Cloud-Manager-Operator-VZzSSP9-SnapCenter Add Tags to Connector Instance Om AWS Managed Encryption 0 Master Key: aws/ebs (default) Change Key Previous Next 1. Configure networking with the proper VPC, Subnet, and SSH Key Pair for connector access.
- 2. Name the connector instance and select Create Role under Details.

Add BlueXP Co	onnector - AWS			More Information	
	AWS Credentials O Details Setwork	4 Security Group	5 Review		
	Network				
	Connectivity Proxy	Configuration (Optional)			
	VPC HTTP F				
	vpc-0b522d5e982a50ceb - 172.30.15.0/25	mple: http://172.16.254.1:8080			
	Subnet Define	Credentials for this Proxy $$			
	172.30.15.0/25 priv-subnet-01 Vpload	d a root certificate 🗸			
	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.				
	Previous	Next			
Set the Security Gr	oup for the connector.				
	oup for the connector.			More Information	
		Security Group	(5) Review	More Information	
	onnector - AWS		3 Review	More Information	
	Onnector - AWS AWS Credentials Details Network Security Grou	up	(5) Review	More Information	
	ONNECTOR - AWS	up	3 Review	More Information	
	ONNECTOR - AWS	up		More Information	
Add BlueXP C	AWS Credentials O Details Network AWS Credentials O Details Network Security Grout The security group must allow inbound HTT Assign a security group: Create a new security group O Se	up P, HTTPS and SSH access.			
Add BlueXP C	AWS Credentials O Details O Network AWS Credentials O Details O Network Security Group The security group must allow inbound HTT Assign a security group: O Create a new security group O Se 1 Security Group	up P, HTTPS and SSH access. elect an existing security group		More Information	
Add BlueXP C	AWS Credentials O Details O Network AWS Credentials O Details O Network Security Group The security group must allow inbound HTT Assign a security group: O Create a new security group O Se 1 Security Group	up P, HTTPS and SSH access.			
Add BlueXP C	AWS Credentials O Details O Network Cecurity Group AWS Credentials O Details O Network Cecurity Group The security group must allow inbound HTT Assign a security group: O Create a new security group O Se Cecurity Group Security Group	up P, HTTPS and SSH access. elect an existing security group			
Add BlueXP C	AWS Credentials O Details O Network Cecurity Group AWS Credentials O Details O Network Cecurity Group The security group must allow inbound HTT Assign a security group: O Create a new security group O Se Cecurity Group Security Group	up P, HTTPS and SSH access. elect an existing security group			
Add BlueXP C	AWS Credentials O Details O Network Cecurity Group AWS Credentials O Details O Network Cecurity Group The security group must allow inbound HTT Assign a security group: O Create a new security group O Se Cecurity Group Security Group	up P, HTTPS and SSH access. elect an existing security group			
Add BlueXP C	AWS Credentials O Details O Network Cecurity Group AWS Credentials O Details O Network Cecurity Group The security group must allow inbound HTT Assign a security group: O Create a new security group O Se Cecurity Group Security Group	up P, HTTPS and SSH access. elect an existing security group			
Add BlueXP C	AWS Credentials O Details O Network Cecurity Group AWS Credentials O Details O Network Cecurity Group The security group must allow inbound HTT Assign a security group: O Create a new security group O Se Cecurity Group Security Group	up P, HTTPS and SSH access. elect an existing security group			
Add BlueXP C	AWS Credentials O Details O Network Cecurity Group AWS Credentials O Details O Network Cecurity Group The security group must allow inbound HTT Assign a security group: O Create a new security group O Se Cecurity Group Security Group	up P, HTTPS and SSH access. elect an existing security group			
Add BlueXP C	AWS Credentials O Details O Network Cecurity Group AWS Credentials O Details O Network Cecurity Group The security group must allow inbound HTT Assign a security group: O Create a new security group O Se Cecurity Group Security Group	up P, HTTPS and SSH access. elect an existing security group			
Add BlueXP C	AWS Credentials O Details O Network Cecurity Group AWS Credentials O Details O Network Cecurity Group The security group must allow inbound HTT Assign a security group: O Create a new security group O Se Cecurity Group Security Group	up P, HTTPS and SSH access. elect an existing security group			

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 Conne	ctor - AWS					More Information	×
	AWS Credentials	Oetails	Network	Security Group	5 Review		
			Review				
				Code for Terraf	orm Automation		
	BlueXP Connector Name	aws-snapctr-us-e	east				
	AWS Access Key	AKIAX4H43ZT56I	IWWR3TI				
	Region	us-east-1					
	VPC	vpc-0b522d5e98	2a50ceb - 172.30.15.	0/25			
	Subnet	172.30.15.0/25 p	priv-subnet-01				
	Key Pair	sufi_new					
	Public IP	Use subnet settin	ngs (Disable)				
	Ргоху	None					
	Security Group	default					
	Г	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.

Identity and Access X Management (IAM)	1A20 > Boles		
Q. Search LAM	Roles (106) and An IAM role is an identity you can inside that has specific premissions with insidestate that you that.	that are solid for short durations. Notes can be assumed by entities	Create role
Dashboard	Q. Sherit	C 1 2 2 4	5 5 > 6
Access management	Role name 🗢	Trusted entities	Last activity
their groups	Amazant Califord of aurochimized	AWS Server 1:12	
Users	AmazondSMRoleFormtancesQuickDetup	AWS Service ec2	156 days ago
Roles			
Policies	aws cottrobuver Administrativit vecubio Role	Account: 982617961887	
identity providers	aws-controlltower CariligReconterRide	AWS Service config	
Account settings	awa-controllower Forward/Instructification/fole	AWS Similor lambda	2 days ago
Access reports	awv controllower ReadOutyLeputer Pole	Amount 982/117661487	
Atuma analyter			
Authine rules	AWS QuickSetup StackSet Local AdministrationRole	AWS Service: doutformation	354 days 990
Avtatyzers	AWS QuickSetup BlackSet Const Ceesuble/Able	Account: 541091183547	354 daya aga
Sattings Credential report	AntiControl (wert recution	Account 292306980405	237 days aga
Organization activity	AWGResenved1530_AWGAdministratorAccess_3dbeb05ab00902fb	Identity Provider: arr/awy.autr.541696183547.saml-provider/AWSSSO_91b222138526441.00_WOT_DELETE	26 days ago
Service control policies (SCPs)	AmpleservedSSD_AmplogenUnionsFullAccess_36cRbs705667et53	Identity Provider, am awa laam. 541606183547 sami provider/AW8550, 91b22273b256441,00, NOT, DELETE	
	AWSReserved1100_AWSPowerfUterrAccess_S0905eadfoa40ed1	Identity Provider, am awstam, 541696183547 sami-provider/AWSSS0, 91b222188b2bf441, DO, NOT, DELETE	
Network consider	AWSReservedS00_AWSReadOrlyAccess_234340507478b114	Identity Provider: an aws sam: 541666183547 sami-provider/AWSSSO, 91b222738b256441_00_NOT_DELETE	
AWS Organizations 2	ANVARIANTING SO, SAA Dev Readonly, Infect 1 ad 010 eff 1 fe 7	Identity Provider: am aws sam: 541696183562 same provider AW0550, 91b222158b26441, 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.

12 Purchassions 13 not, revisions, and creates 14 Allow values feedwates to the SC2_L and bala, or others to be second. 15 Allow values feedwates to the SC2_L and bala, or others to be second. 15 Allow values feedwates to the SC2_L and bala, or others to be second. 15 Addition to the SC2_L and bala, or others to be second. 15 Addition to the SC2_L and bala, or others to be second. 15 Addition to the SC2_L and bala, or others to be second. 15 Addition to the SC2_L and bala, or others to be second. 15 Addition to the SC2_L and balan, or others to be second. 15 Addition to the SC2_L and bala, or others to be second. 16 Addition to the SC2_L and bala, or others to be second. 16 Addition to the SC2_L and bala, or others to be second. 17 Addition to the SC2_L and bala, or other to be second. 18 Addition to the SC2_L and the second. 19 Addition to the SC2_L and the second. 10 Not second. 10 Not second. 10 Not second. 10 Not second. 10 Not second. 10 Not second. <th>ect trusted entity</th> <th>Select trusted entity Info</th> <th></th> <th></th> <th></th>	ect trusted entity	Select trusted entity Info			
 MS service		Trusted entity type			
Allow users federated with SAML 2.0 from a coporate diversity for perform actions in this account. Create a custom trust policy to enable others to perform An AWS account Allow entities in other AVS accounts belonging to you or a 3rd party to perform actions in this account. Image: Create a custom trust policy to enable others to perform This account (\$416406183547) Image: Create a custom trust policy to enable others to perform actions in this account. On Another VIS account Belonging to you or a 3rd party to perform actions in this account. Image: Create a custom trust policy to enable others to perform On Another VIS account Belonging to you or a 3rd party to perform actions in this account. Image: Create a custom trust policy to enable others to perform On Another VIS account (\$416406183547) Image: Create a custom trust perform Image: Create a custom trust policy to enable others to perform Second ID Image: Create a custom trust policy to enable others Image: Create a custom trust policy to enable others Second ID is a 12 cogit number. Image: Create a custom trust policy to enable others Image: Create a custom trust policy to enable others Options Coping number. Image: Create a custom trust policy to enable others Image: Create a custom trust policy to enable others		Allow AWS services like EC2, Lambda, or others to	Allow entities in other AWS accounts belonging to you	Allows users federated by the specified external web identity provider to assume this role to perform actions	
Allow entities in other AWS accounts belonging to you or a 3rd party to perform actions in this account.		Allow users federated with SAML 2.0 from a corporate	Create a custom trust policy to enable others to perform		
Another AWS account Account ID United of the account that can use this noise S2013314444 Account ID is a 12-digit number. Options		Allow entities in other AWS accounts belonging to you or a 3rd party to p	erform actions in this account.		
Options		Another AWS account Account ID Ioenther of the account that can use this role 952013314444			
Require MFA		Options	sume this role)		

3. Filter permission policies by fsx and add **Permissions policies** to the role.

			Create policy
Q. Filter policies by property or policy name and press enter	ur.	4 matches	< 1 >
fsx X Clear filters			
■ Policy name □ ⁿ マ	Туре 🗢	Description	
E AmazonFSxReadOnlyAccess	AWS ma	Provides read only access to Amazon FSx.	
MazonFSxFullAccess	AWS ma	Provides full access to Amazon FSx and access to related AWS services.	
The AmazonFSxConsoleReadOnlyAccess	AWS ma	Provides read only access to Amazon FSx and access to related AWS services via the AWS Management Console.	
AmazonFSxConsoleFullAccess	AWS ma	Provides full access to Amazon FSx and access to related AWS services via the AWS Management Console.	
	'fsx' X Clear filters Policy name (z* ① ① ① ② ① ② ③ ④ ③ AmazonFSxFullAccess ○ ② ① ③ ③ ③ ③ ④ ③ AmazonFSxFullAccess	■ Policy name [2* ▼ Type ▼ ● ● AmazonFSxReadOnlyAccess AWS ma… ☑ ● ● AmazonFSxFullAccess AWS ma… ● ● AmazonFSxConsoleReadOnlyAccess AWS ma… ● ● AmazonFSxConsoleReadOnlyAccess AWS ma… ● ● AmazonFSxConsoleFullAccess AWS ma…	"fax" X Clear filters Policy name (2" V Type V Description Image: Strate Stra

4. In **Role details** page, name the role, add a description, then click **Create role**.

Step 1 Select trusted entity	Name, review, and create	
Step 2 Add permissions	Role details	
Step 3	Role name Enter a meaningful name to identify this role.	
Name, review, and create	fsxn_bluexp	
	Maximum 64 characters: Use alphanumeric and '+@' characters;	
	Description Add a short explanation for this role.	
	Grant permission for BlueXP access to FSxN in AWS.	
	A Maximum 1000 characters. Use alphanumeric and ++, characters.	
	Step 1: Select trusted entities	Edit
	<pre>1 - (() "Version": "2012-10-17", " "Statement": [</pre>	

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.

n Ne	tApp BlueXP			Q BlueXP Search	Account 💙 Automation-te	Workspace Y Database-2	Connector 💙 acao-aws-conn	40 🌣	?	8
	Credentials	Account credentials User credentials								
9					11000					
٠				e account-level credentials to de s in your cloud environment.	eploy					
6		5 Credentials				Add credentials				
0		aws shantanucrec Type: Assume F								
**		210811600188	nkarthik_kafka_nfs_role_FSxN							
		AWS Account ID	Assume Role							

6. Choose credential location as - Amazon Web Services - BlueXP.

	Q BluckP Search Account * Workspace * Connect Automation-te Database-2 acao-aw	tor 🎽 🦣 🏘 😗 😫
Add Credentials		×
8 9 0	Choose Credentials Location	
© **	Choose how to associate the credentials Connector	
	~	
	Net	c

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.

n Ne	tApp BlueXP	Q_BlueXP Search Account Workspace Connector account_ account_	le 9 9 4
	Add Credentials	Credentials Type 2 Define Credentials 3 Review	×
9		Define Amazon Web Services Credentials	
		Learn more about AWS authentication methods	
@ 0		When creating the IAM role, select Another AWS account and enter the account ID for Blue XP: 952013314444 d	
*		Credentials Name	
		Previous Next	Q

8. Review and Add.

n Net	etApp BlueXP		Q BlueXP Search Accor	onnector 👻 🔒 🏚 🌣 🥐
	Add Credentials	🚫 Crede	ntials Type 🥥 Define Credentials 3 Review	×
Q			Review	
٠				
ô		Credentials Type	AWS	
۲		Credentials Name	fsxn_bluexp	
4		Credential Storage	Cloud Manager	
		Role ARN	arn:aws:iam::541696183547:role/fsxn_bluexp	
		Pre	vious Add	

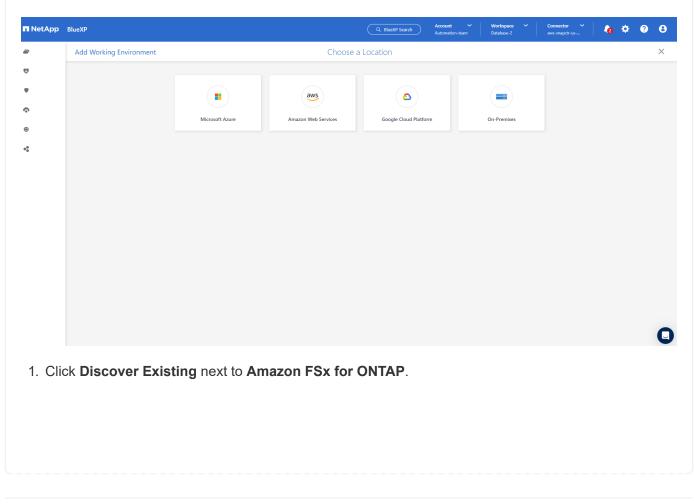
SnapCenter services setup

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.

		Q BlueXP Search	Account V Workspace Automation-team Database-2	Connector aws-snapctr-us	🏚 🌣 🥹 🖯
Canvas My Working Environments	My Opportunities				Go to Tabular View
+ Add Working Environment			C Enable Services 🚯	Working Environments	
Ŷ					
۲					
*					
	Amazon 53 5 Buckets AVS				
			-+		0

1. Choose Amazon Web Services as the location.



🗖 NetAp	op BlueXP				Q BlueXP Search	Account 🛩 Automation-team	Workspace V Database-2	Connector 🛩 aws-snapctr-us	6	¢	?	8
	Add Working Environment			Choose a	Location							×
Ø				aws								
•			Microsoft Azure	aws Amazon Web Services	Google Cloud Platform	On-Premises						
¢				Select	Туре							
۲		Cloud Volum	es ONTAP									
0 ⁰ 0		Single Node			Disc	cover Existing	udd new 🗸 🗸					
		Cloud Volum (High Availab			Disc	cover Existing	udd new 🗸 🗸 🗸					
		(FS%) Amazon FSx (High Availab			Disc	cover Existing	udd new 🗸 🗸					
		Kubernetes C Any	Cluster				Discover					
												0

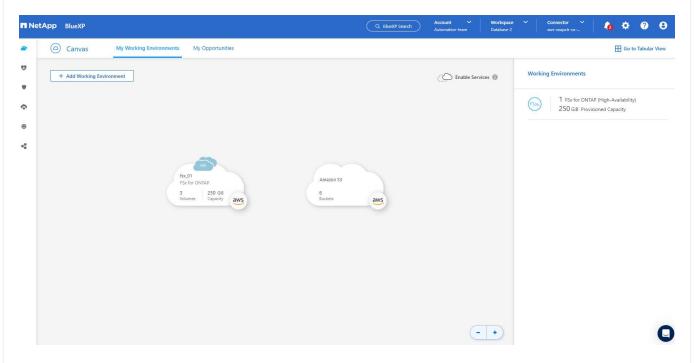
1. Select the **Credentials Name** that you have created in previous section to grant BlueXP with the permissions that it needs to manage FSx for ONTAP. If you have not added credentials, you can add it from the **Settings** menu at the top right corner of the BlueXP console.

n Net	App BlueXP	Q BlueXP Search Account V Workspace V Connector V Automation-Isam Database-2 ans-snapstr-us-	a 🌣 🥹 😆
	Add an Existing FSx for ONTAP	FSx for ONTAP Authentication	×
9			
٠		Select the credentials that provides BlueXP with the permissions that it needs to manage FSx for ONTAP.	
ବ		Credentials Name	
۲		DemoFSxNCMCredentials \sim	
4		To add a new set of credential, go to the Credentials Page.	
		Previous Next	0

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.

n Ne	etApp BlueXP						count Y	Workspace 💙 Database-2	Connector aws-snapctr-us	٠	0 B
	Add an Existing FS:	for ONTAP		Sele	ect FSx for ONT	AP					×
9											
٠			Choose an A	WS region and then s	select the working en	vironment that you w	ant to add				
6				Region us-east-1 US E	ast (N. Virginia)	~					
۲		1 FSx for ONTAP instance							Q		
** 0		Name	File System ID File System ID	VPC ID +	Subnet ID +	Management Addres	s ¢	Deployment modal	÷ Tags ÷		
		O fsx_01	fs- 02ad7bf3476b741df	vpc- 0b522d5e982a	subnet- 04f5fe7073ff5	management.fs- 02ad7bf3476b741df.ft	sx.us-east	Single Availability Zone	٩		
				Previou	us A	dd					0

1. The discovered Amazon FSx for ONTAP instance now appears in the working environment.



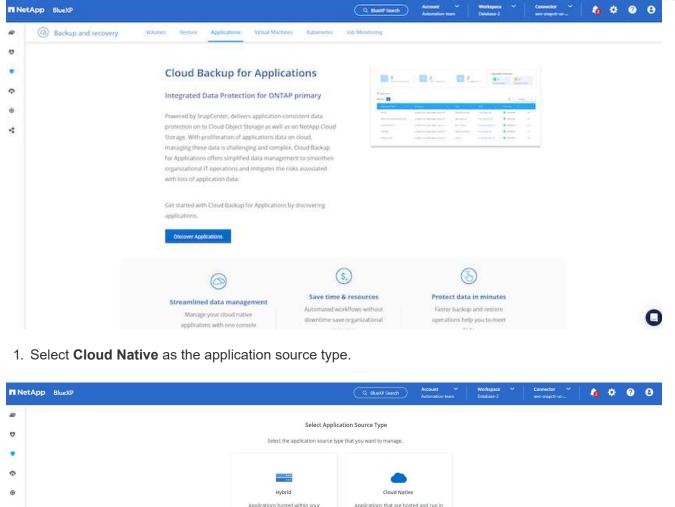
1. You can log into the FSx cluster with your fsxadmin account credentials.

n Ne	tApp BlueXP		Q. Bluez@ Search Account V Workspace V Connector V Actomation-team Database-2 avvs-snapttr-us Actomation-team	8
8	Isx_01	Overview Volumes	Timeline C (()
ø				
			Se the	
¢				
۲			One last store and you are ready to store	
-			One last step and you are ready to start Provide ONTAP Cluster Password	
			FLOWLE UNIAF LUSIEL PASSWORD	
			user name fszadmin	
			ONTAP Cluster Password	
			Save Back to Canvas	O

1. After you log into Amazon FSx for ONTAP, review your database storage information (such as database volumes).

Net	tApp Blu	еХР					Q BlueXP Searc	h Account Automation-team	Workspace	Connector aws-snapctr-u	ຼັ 📔 🙆	¢ 0	8
	🛞 fsx	_01 0	verview Volumes								Timeline	C	; (
			Volumes Summar	ry 😫	3 Volumes	250 GiB Provisioned Capacity	26 SSD	.03 GIB	O GIB Capacity Pool	licod			
					volumes	Provisioned capacity	330	osea	Capacity Poo	Used			
			3 Volumes						Q 🔚	Add Volume			
							A						
			ora_01_da	ita	ONL	INE Manage Volume	ora_01	logs	ON	LINE Manage Volume			
			INFO		CAPACITY		INFO		CAPACITY				
			Disk Type SVM Name	SSD svm_ora	Provisioned SSD Used	100 GiB 5.79 GiB	Disk Type SVM Name	SSD svm_ora	Provisioned SSD Used	100 GiB 1.14 GiB			
			Tiering Policy	Snapshot Only	Capacity Pool	0 GiB	Tiering Policy	Snapshot Only	Capacity Pool	0 GiB			
					Used				Used				
			ora_01_bit	ny	ONL	INE Manage Volume							
			INFO		CAPACITY								
			Disk Type	SSD	Provisioned	50 GiB							
			SVM Name	svm_ora	SSD Used	19.1 GiB							
			Tiering Policy	Snapshot Only	Capacity Pool Used	0 GiB							(

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**.



■ NetApp BlueXP		Q BlueXP Search Account ~ Automation-team	Workspace Y Database-2	Connector Y aws-snapctr-us	6 • (98
e 1	Select Applic	cation Source Type pe that you want to manage.				
•		•				
© **	Hybrid Applications hosted within your organization's infrastructure.	Cloud Native Applications that are hosted and run in the cloud using AWS, Azure, GCP, etc.,				
	Cancel	Next				0

1. Choose **Oracle** for the application type.

🗖 Ne	tApp BlueXP		Q BlueXP Search Automation-team	Connector Y aws-snapctr-us	lo K	F ?	θ
	Discover Applications						×
•							
¢		Select Appl	lication Type				
۲		ORACLE	SAPHANA				
•:		Oracle	SAP HANA				
		N	lext				0

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**.

n Net	tApp BlueXP				orkspace V Connector stabase-2 aws-snapctr-us	
	Backup and recovery	Volumes Restore	Applications Virtual Machines Kubernetes	Job Monitoring		
ø						
•			Add Ho			
Ģ						
۲			Host Installation Type O Manual 🚯	Using SSH 1		
			Host FQDN or IP	Connector 🚯		
			172.30.15.58	aws-snapctr-us-east 🖤		
			Username(Sudo) 🕦			
			ec2-user	Add SSH Private Key Optional (
			SSH Port	Plug-in Port		
			. 22	8145		
						0
			Cancel	Next		e

2. Paste in your ec2-user SSH key for the database EC2 host and click on **Validate** to proceed.

e 💙 Connector 🎽 👍 🏟 ?

3. You will be prompted for **Validating fingerprint** to proceed.

n Ne	tApp BlueXP					Q BlueXP Search	Account ~ Automation-team	Workspace Database-2	~	Connector 🛩 aws-snapctr-us	6	٠	?	θ
	Backup and recovery	Volumes Restore	Applications	Virtual Machines	Kubernetes	Job Monitoring								
ø														
				Provide the followi	Add Hos	t ost and discover applications								
6				Provide the following	ng detana to add in	stand discover applications								
۲			Host Insta	lation Type 🛛 🔿	Manual 🚯	Using SSH ()								
•														
			Host	Validating fingerpri	int				- 1					
			Usen	Algorithm		ssh-rsa								
			er?	Fingerprint		AAAAE2VjZHNhLXNoYT	ItbmlzdHAyNTYAAAAIb	ml						
			SSH I	By proceeding further	r, I confirm that the	above fingerprint for host is valid	d.							
			22											
			_				Procced	Cancel						
					Cancel	Next								0

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.

n Net	etApp BlueXP			Q BlueXP Search	Account V Workspace V Connector V Automation-team Database-2 aws-snaptir-us	?
8	Backup and recovery	Volumes Restore	Applications Virtual Machines Kubernet	les Job Monitoring		
69						
•			Oracle			
ô		Cloud Native	• Oracle	*		
0					Application Protection	
		1 Hosts	ORACLE	Clone	0 1 Protected Unprotected	
		1 Databases				
		Filter By 🕂		Q	Manage Databases V Settings V	
		-				
		Name	A Host Name	Policy Name	Protection Status	
		db1	172.30.15.58		😲 Unprotected 🚥	
					1-1 of 1 << < 1 > >>	
		: • : - I • • · · · - • • • •	Span Captar comilas	a for Oracla T	he payt three continue of this	
; co	mpletes the in	itial setup of a	SnapCenter service	S IUI UIACIE. I	he next three sections of this	

Oracle database backup

1. Click the three dots next to the database **Protection Status**, and then click **Polices** to view the default preloaded database protection policies that can be applied to protect your Oracle databases.

Image: Secure with the second w	pp BlueXP			Q BlueXP Search	Account V Workspace Automation-team Database-2	Connector aws-snapctr-us	´ 4	•
Application Protection Protected Protected Application Protection Protected Protected Policies Application Protection Protected Protected Policies About Hosts Policies About Hosts Policies About Hosts Policies About Hosts Policies About Hosts Policies Policies About Hosts Policies About Hosts Policies About Hosts Policies About Hosts	Backup and recovery	Volumes Restore	Applications Virtual Machines Ku	ubernetes Job Monitoring				
Application Protection Protected Application Protection © 0 0 1 Unprotected Unprotected Name								
Image: Detablases I		Cloud Native	▼ Oracle	~				
Hosts ORACLE Clone Protected Unprotected					Application Protection	i		
Name Name Name Policy Name Protection Status db1 172.30.15.58 ♥ Unprotected Hosts		1 Hosts	BB 1 ORACLE					
Name Name Policy Name Protection Status Policies db1 172:30.15:58 Unprotected Hosts								
Name Image: Host Name Policy Name Protection Status Polices db1 172.30.15.58 Unprotected Hosts		1 Databases						
Name Name Policy Name Protection Status db1 172:30.15:58 Unprotected Hosts		Filter By 🕂		1.0	Q Manage Databases 🔻			
Hosts		Name	∧ Host Name	Policy Name	Protection Status			
1-1of1 « < 1 > >>		db1	172.30.15.58		0 Unprotected	Hosts		
					1 - 1 of 1	«< < 1 > >>		

1. You can also create your own policy with a customized backup frequency and backup data-retention window.

Backup and recovery	Volumes Restore	Applications Virtual Machines	Kubernetes Job Monitoring			
Backup and recovery	volumes restore	Applications	Robernetes Job Monitoring			
Applications > Policies						
	Cloud Native	• Oracle	×			
	4 Policies		۹	Create Policy		
	Policy Name	A Backup Type	Schedules and Retention	1		
	Oracle Full Backup for Bronze	FullBackup	Daily: Repeats Every 1 Day, Keeps 14 copies Weekly: Repeats Every Fri, Keeps 4 Copies Monthly: Repeats Every 1st Day of Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep. Oct.	. Nov		
	Oracle Full Backup for Gold	FullBackup	Hourly: Repeats Every 6 Hrs, Keeps 16 copies Daily: Repeats Every 1 Day, Keeps 30 copies Weekly: Repeats Every 1 Hosps 4 copies Monthly: Repeats Every 1st Day of Jan, Feb. Mar, Apr. May, Jun. Jul. Aug. Sep. Oct.	•••		
	Oracle Full Backup for Silver	FullBackup	Hourly: Repeats Every 12 Hrs. Keeps 6 copies Daily: Repeats Every 1 Day, Keeps 14 copies Weekly: Repeats Every 14: Day of Jan. Feb. Mar, Apr. May, Jun. Jul. Aug. Sep. Oct Monthy: Repeats Every 14: Day of Jan. Feb. Mar, Apr. May, Jun. Jul. Aug. Sep. Oct	•••		
	my_full_bkup	FullBackup	Hourly: Repeats Every 1 Hr, Keeps 3 Days			
			1 - 4 of 4	<< < 1 > >>		
						1

1. When you are happy with the policy configuration, you can then assign your policy of choice to protect the database.

tApp BlueXP					Q BlueXP Search	Account Automati		Workspace Database-2	~	Connector 💙 aws-snapctr-us	- 60	٥	?
Backup and recovery	Volumes	Restore Applications	Virtual Machines	Kubernetes	Job Monitoring								
	Cloud Native		 Oracle 		×								
	_		_		_		Applicatio	on Protection	î				
	1 Hosts		CRACLE		Clone		0 🤡		<mark>り</mark> 1				
	Hosts		ORACLE		Clone		Protected		Unprote	cted			
	1 Databases												
	Filter By +					Q Ma	inage Databas	es ∣▼	Se	ttings 🔻			
	Name	∧ Host I	lame	\$	Policy Name	P	rotection Sta	tus	¢				
	db1	172.30	.15.58			2	🕖 Unprotecte	d		•			
								_	Details	1 > >>			
								Assig	n Policy				

1. Choose the policy to assign to the database.

Volumes Restore Applications Policies	Assigr	Job Monitoring POlicy packups of the database "db1" C Schedules Daily: Repeats Every 1 Day, Keeps 14 copies	
Policy Name	Assign a policy to start taking	c Schedules	
Policy Name	Assign a policy to start taking	c Schedules	
Policy Name	Assign a policy to start taking	c Schedules	
Policy Name			
Oracle Full Backup for Bronze	CullD adjust	Daily: Repeats Every 1 Day, Keeps 14 copies	
	ниваскир	Weekly: Repeats Every Fri. Keeps 4 copies Monthly: Repeats Every 1st Day of Jan, Feb. Mar, Apr. May, Jun, Jul, Aug. Sep. C	
Oracle Full Backup for Gold		Hourly: Repeats Every 6 Hrs. Keeps 16 copies Daily: Repeats Every 1 Day, Keeps 30 copies Weekly: Repeats Every F1r. Keeps 4 copies Monthly: Repeats Every 1st Day of Jan. Feb. Mar, Apr, May, Jun, Jul, Aug. Sep. C	
Oracle Full Backup for Silver	FullBackup	Hourly: Repeats Every 12 Hrs, Keeps 6 copies Daily: Repeats Every 10 Jay, Keeps 14 copies Weekly: Repeats Every F1: Keeps 4 copies Monthly: Repeats Every 1st Day of Jan, Feb, Mar, Apr, May, Jun, Jul, Aug. Sep, C	
my_full_bkup	FullBackup	Hourly: Repeats Every 1 Hr, Keeps 3 Days	
		1-4of4 << 1 > >>	
	Oracle Full Backup for Silver	Oracle Full Backup for Silver FullBackup	Oracle Full Backup for Gold FullBackup Weekly: Repeats Every 11, Keeps 4 copies Monthly: Repeats Every 11 Day, Keeps 4 copies Doracle Full Backup for Silver FullBackup Weekly: Repeats Every 11 Day, Keeps 4 copies Weekly: Repeats Every 11 St Day of Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, C my_full_bkup FullBackup Hourly: Repeats Every 11 Kr, Keeps 3 Days

1. After the policy is applied, the database protection status changed to **Protected** with a green check mark.

🗖 Ne	tApp BlueXP						Q BlueXP Search	Accou Autom	int 🗸	Workspace Database-2		Connector aws-snapctr-us-		6	٥	?	8
	Backup and recovery	Volumes	Restore	Applications	Virtual Machines	Kubernetes	Job Monitoring										
ø																	
•		Cloud Native			Oracle												
ê																	
۲		<u> </u>			88 1				Application	on Protectio	n 🚺 0						
•		Hosts			ORACLE		Clone		Protected		Unprote	ected					
		1 Databases						_					_				
		Filter By 🕂						۹ 🗖	Manage Databas	ses 🔻	S	ettings 🔻					
		Name		∧ Host Nam	ie i	0	Policy Name	$\langle \psi \rangle$	Protection Sta	tus	¢						
		db1		172.30.15.	58		Smy_full_bkup		Protected								
										1 - 1 of	1 <<	< 1 > >>	•				
																	C
																	-

1. The database backup runs on a predefined schedule. You can also run a one-off on-demand backup as shown below.

n Net	App BlueX	P			Q BlueXP Sea	rch Account ~		Vorkspace 🛩 🛛	Connecto aws-snapc		6	۰	?	8
	G Back	up and recovery	Volumes	Restore Applications	Virtual Machine	es Kubernetes Job	Monitorir	ng						
9														
•		Cloud Native		▼ Oracle										
@ 0								Application	Protectio					
•		Hosts				Clone		1 Protected		0 Unprotect	ed			
		1 Databases												
		Filter By +					۹ 🛛	Manage Databases	1 🔻	Setti	ings	•		
		Name	^	Host Name	0	Policy Name		Protection Statu	5	e				
		db1		172.30.15.58		Soracle Full Backup for	r Gold	Sector Protected						
									View Detai		1 >	>>		
									Assign Poli Un-assign					
									-11 035501					0
														U

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.

n NetApp	BlueXP		Q BlueXP Search	Account V Workspace Automation-team Database-2	Connector aws-snapctr-us	• • • • •
<u>م</u>	Backup and recovery Volum	es Restore Applications	Virtual Machines	Kubernetes Job Monitoring		
9	Applications > Database Details					
			Database I	Details		
۵	db1 Database Name	Service 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 +				Q Select Time	frame 🔻
	Backup Name	¢.	Backup Type	🔅 SCN 🔅 Backup Da	te ~	
	Oracle_Full_Backup_for_Gold_Weekly_	db1_2023_03_24_19_12_18_60900_1	Log	2589354 Mar 24, 20.	23, 3:12:34 pm Delete	
	Oracle_Full_Backup_for_Gold_Weekly_	db1_2023_03_24_19_11_51_51476_0	Data	2589306 Mar 24, 20	23, 3:12:18 pm •••	
	Oracle_Full_Backup_for_Gold_Hourly_0	db1_2023_03_24_18_10_31_71953_1	Log	2586621 Mar 24, 20.	23, 2:10:45 pm Delete	
	Oracle_Full_Backup_for_Gold_Hourly_0		Data	2586557 Mar 24, 20	23, 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.

NetApp	D BlueXP			Q. BlueXP Search	Account 🗸	Workspace Connector Database-2 aws-snapctr-us	la 🌣 🔞
(Backup and recov	very Volumes Restore	Applications Virtual Machines	Kubernetes Job Monitoring			
	Applications > Data	ibase Details					
				Database Details			
		db1 Database Name	Service Protected Protection	Oracle Full Back Policy Names	up for Gold	Database Type	
		172.30.15.58 Host Name	FSx Host Storage	Unreachable Database Versio	n	bKed8yv2T19BJ0V5QyqvA Agent ld	
		- Clones	- Parent Database				
		6 Backups					
		Filter By +				Q Select Timeframe 🔻]
		Backup Name	÷ 1.	Backup Type 🔅 🕴 SCN	0 Backup Date	→ 1	
		Oracle_Full_Backup_for_Gold_Hourly_db1	2023_03_24_18_10_31_71953_1	Log 2586621	Mar 24, 2023, 2:1	10:45 pm Delete	
		Oracle_Full_Backup_for_Gold_Hourly_db1	2023_03_24_18_10_03_70535_0	Data 2586557	Mar 24, 2023, 2:1	10:31 pm •••	
		Oracle_Full_Backup_for_Gold_Hourly_db1	2023_03_24_15_37_04_98851_1	Log 2580577	Mar 24, 2023, 11	:37:1 Restore e	
						ar so to the	

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.

	etApp BlueXP		Connector Y ws-snapctr-us	¢ ?	8
8	Restore "db1"	Restore Settings (2) Review			×
•		Restore Settings			
© •*		Restore Scope All Data Files Data Files Restore Control Files Control Files Restore 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. Database state will be changed if needed for restore and recovery. Recovery Scope All Logs Until System Change Number Date and Time No Recovery Archive Log Files Locations Imnt/log_location001 Open the database or the container database in READ-WRITE mode after recovery. 			
1.	Review and start data	Previous Next			0

III Ne	etApp BlueXP		Q BlueXP Search	Account 🛩 Automation-team	Workspace V Database-2	Connector 💙 aws-snapctr-us	6	¢ 6	8
-	Restore "db1"	Restore Se	ttings 2 Review						×
Q									
٠		F	Review						
6		Backup Name	Oracle_Full_Backup_for_Go	old Weekl					
۲		Баскор манте	y_db1_2023_03_24_19_11_5 0						
•:		Restore Scope	All Data Files						
		Recovery Scope	All Logs						
		Force In Place Restore	Yes						
		Open Database or Container Database in READ-WRITE Mode After Recovery	Yes						
		Previous	Restore						0

1. From the **Job Monitoring** tab, you can view the status of the restore job as well as any details while it is running.

	- Advanced Search	h & Filtering	Timeframe: Last 24 Hours				Last Updated March 24 2023, 15:25:33
		in a finaning					
ol		Туре	Resource Name	Status	🗘 Job Name	\$	Start Time
đ	[]] 1fdca0bd-a9c8-45aa			Success	Restore for Oracle Data	abase db1	Mar 24 2023, 3:16:28 pr
Ø	[]] f6f4fe2d-3040-497f			Success	Backup of db1 oracle c	latabase o	Mar 24 2023, 3:11:51 pr
Ø	[]] 5e3299f5-29db-4dcc			⊘ Success	Backup of db1 oracle c	latabase o	Mar 24 2023, 2:10:03 pr
đ	[]] 6da5e51e-1a79-4e7e			⊘ Success	Initialize FullBackup ba	ckup of po	Mar 24 2023, 2:10:01 pr

n N	letApp	BlueXP				BlueXP Search	Account Automation	√ team	Workspace V Database-2	*	Connector 🗸	6	\$?	8
3	6	Backup and recovery	Volumes Re	store	Applications Virtua	l Machines	Kubernetes	Job N	lonitoring						
9		Job Monitoring > Job Id: 1fdca0	bd-a9c8-45aa-9d7a-05	a07cb291	f4										
•						Job Det	ails								
6					Job ld: 1fdc	a0bd-a9c8-45aa	-9d7a-05a07cb29	1f4				Expand			
0		Sub-Jobs(6) Job Name		\$	Job ID	\$ Start Ti	me	¢	End Time	\$	Duration	¢	Ð		
•		Restore for Oracle D	atabase db1 using back	sup	1fdca0bd-a9c8-45aa-9	d 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			fb7b1171-9f6f-4228-9	e Mar 24	2023, 3:20:06 pm		Mar 24 2023, 3:23:19 pm		3 Minutes				
		Restore			Of4580d0-6598-458b-a	a7 Mar 24	2023, 3:17:49 pm		Mar 24 2023, 3:20:07 pm		2 Minutes				
															0

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.

NetAp	p BlueXP				Q BlueXP Search	Account ~ Automation-team	Workspace 💙 Database-2	Connector 🗸	6	1
- 17	Backup and recovery	y Volumes Restore	Applications Virtual Machin	nes Kubernetes	Job Monitoring					
	Applications > Databas	e Details								
				Datab	ase Details					
		db1 Database Name	⊘ Protected Protection		Oracle Full Back	up for Gold	Database Type			
		172.30.15.58	FSx		Unreachable		bKed8yv2T19BJ0V5Qy	/qvA		
		Host Name	Host Storage		Database Version	n	Agent Id			
		- Clones	- Parent Databa	se						
		2 Backups								
		Filter By					Q Selec	t Timeframe 🔻		
		Backup Name		0 Backup Type	e 🗘 SCN	0 Backup Date	· • 1			
		Oracle_Full_Backup_for_Gold_He	ourly_db1_2023_03_24_13_34_41_3049	1_1 Log	2575607	Mar 24, 2023, 9:34	4:55 am Delete			
		Oracle_Full_Backup_for_Gold_He	ourly_db1_2023_03_24_13_34_07_2674	8_0 Data	2575555	Mar 24, 2023, 9:34	4:41 am •••			
							Delete	< 1 > >>		
							Clone			

1. Select the **Basic** option if you don't need to change any cloned database parameters.

e Database of "db1"	Provide following details to create a clor Select Clone Options		Clone kup "Oracle_Full_Backup_for_Gold_Ho	urly_db1_2023_03_24_13_		
		ne from the database bac	kup "Oracle_Full_Backup_for_Gold_Ho	urly_db1_2023_03_24_13_		
			 Specification file (1) 			
	Clone Host 172.30.15.58 Clone Naming Scheme	•	Clone SID db1clone Oracle Home	•		
	Auto-generated Database Credentials	Optional	/u01/app/oracle/product/19.0.0/d ASM Credentials	Dptional		
	[Cancel	Next			C
		Add Credential				

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.

		Automation-team Database-2 aws-snapetr-us 🔽 🏧 🤝
Clone [Database of "db1"	Clone Details 2 Review
		Create Clone
		Provide following details to create a clone from the database backup "Oracle_Full_Backup_for_Gold_Weekly_db1_2023_03_24_19
		Select Clone Options O Basic 🕦 🐵 Specification file 🕦
		(j) Generate specification file to modify input parameters and use for clone.
		Specification File db1 3 24 2023 10 14 specision 1 Browse
		db1_3_24_2023_10_14_spec.json ⊥ Browse Clone Host Clone SID
		172.30.15.58 v
		Database Credentials Optional ASM Credentials Optional Add Credential Image: Credential description of the second descripticond description of the second description of the second d
		Cancel Next
Revie	w and launch th	ne job.
pp Blue	хр	Q BlueXP Search Account × Workspace × Connector × Automation team Database-2 ans-snaptfr-so-
Clone Data	abase of "db1"	Clone Details 2 Review
		Review
		General Database parameters
		Backup Name Oracle_Full_Backup_for_Gold_Hourly_dbl_2023_03_24_13_34_07_26748_0

1. Monitor the cloning job status from the **Job Monitoring** tab.

Clone Host

Datafile locations

Control files

Redo logs

Recovery scope

172.30.15.58

DATA_db1clone

+DATA_db1clone/db1clone/control/control01.ctl

RedoGroup = 1 TotalSize = 1024 Path = +DATA, db tclone/db tclone/redolog/redo01_01.log RedoGroup = 2 TotalSize = 1024 Path = +DATA, db tclone/db tclone/redolog/redo02_01.log RedoGroup = 3 TotalSize = 1024 Path = +DATA, db tclone/db tclone/redolog/redo03_01.log

Until cancel using selected backup's archive logs

Clone

0

Previous

Sub-Jobs(2) Job Name Glob Name clop Job ID cllp Start Time cllp End Time cllp Duration cllp Cloning Cracle Database db1 as db1done on h cllp cd30abaf-fbe2-4052-a6 Mar 24 2023, 130.36 pm
Job Details Job Id: cd30abaf-fbe2-4052-a6db-4bf965a8d29b Sub-Jobs(2) ■ Job Name \$ Job ID \$ Start Time \$ End Time \$ Duration \$ €
Job Details Job Job Id: cd90abaf-fibe2-4052-a6db-Abf965a8d29b Sub-Jobs(2) Job Name ↓ Job ID ↓ Start Time ↓ Duration
Sub-Jobs(2) Job Name \$ Job 1D \$ Start Time \$ End Time \$ Duration \$ €
Job Name Job ID Image: Start Time Image: End Time Image: Duration Image: Duration
Running pre scripts 🗇 51f152c1-853a-4ec6-a4f Mar 24 2023, 1:30:41 pm Mar 24 2023, 1:30:41 pm 0 Second
Validating clone request 🗇 193a6c44-2eb2-4c5e-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 Resou	irces			
ora.DATA.do	1			
	ONLINE	ONLINE	ip-172-30-15-58	STABLE
ora.DATA_DE	BICLONE, dg			
	ONLINE	ONLINE	ip-172-30-15-58	STABLE
ora.LISTENE	ER.1snr			
	ONLINE	ONLINE	ip-172-30-15-58	STABLE
ora.LOGS.dg	1			
	ONLINE	ONLINE	ip-172-30-15-58	STABLE
ra.LOGS_SC	0_2748138658	3.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
ora.cssd				
1	ONLINE	ONLINE	ip-172-30-15-58	STABLE
1 ora.db1.db			20 - High and South the States and Anna States	
1	ONLINE	ONLINE ONLINE	ip-172-30-15-58 ip-172-30-15-58	Open,HOME=/u01/app/
ora.db1.db	ONLINE		20 - High and South the States and Anna States	Open,HOME=/u01/app/ racle/product/19.0.
1 ora.db1.db 1 ora.db1clor 1	ONLINE ne.db ONLINE		20 - High and South the States and Anna States	Open,HOME=/u01/app/ racle/product/19.0.
1 pra.db1.db 1 pra.db1clor 1 pra.diskmor	ONLINE ne.db ONLINE	ONLINE	ip-172-30-15-58	Open,HOME=/u01/app/ racle/product/19.0. /db1,STABLE Open,HOME=/u01/app/ racle/product/19.0. /db1,STABLE
1 pra.db1.db 1 pra.db1clor 1 pra.diskmor 1	ONLINE ne.db ONLINE 1 OFFLINE	ONLINE	ip-172-30-15-58	Open,HOME=/u01/app/ racle/product/19.0. /db1,STABLE Open,HOME=/u01/app/ racle/product/19.0.
1 pra.db1.db 1 pra.db1clor 1 pra.diskmor 1 pra.driver.	ONLINE ne.db ONLINE oFFLINE	ONLINE ONLINE OFFLINE	ip-172-30-15-58 ip-172-30-15-58	Open,HOME=/u01/app/ racle/product/19.0. /db1,STABLE Open,HOME=/u01/app/ racle/product/19.0. /db1,STABLE STABLE
1 pra.db1.db 1 pra.db1clor 1 pra.diskmor 1 pra.driver. 1	ONLINE ne.db ONLINE 1 OFFLINE	ONLINE	ip-172-30-15-58	Open,HOME=/u01/app/ racle/product/19.0. /db1,STABLE Open,HOME=/u01/app/ racle/product/19.0. /db1,STABLE
1 pra.db1.db 1 pra.db1clor 1 pra.diskmor 1 pra.driver.	ONLINE ne.db ONLINE oFFLINE	ONLINE ONLINE OFFLINE ONLINE	ip-172-30-15-58 ip-172-30-15-58	Open,HOME=/u01/app/ racle/product/19.0. /db1,STABLE Open,HOME=/u01/app/ racle/product/19.0. /db1,STABLE STABLE

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=Cj0KCQiA54KfB hCKARIsAJzSrdqwQrghn6I71jiWzSeaT9Uh1-vY-VfhJixFxnv5rWwn2S7RqZOTQ0aAh7eEALw_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

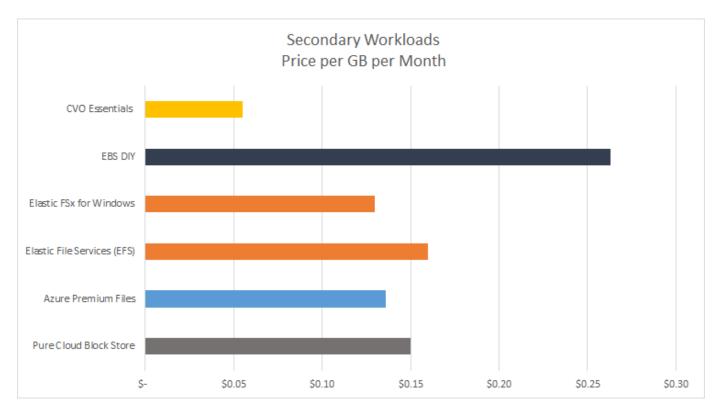
This solution provides NetApp field and customers with instructions and guidance for configuring, operating, and migrating databases to a hybrid cloud environment using the NetApp SnapCenter GUI-based tool and the NetApp storage service CVO in public clouds for the following use cases:

- 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 publiccloud 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:

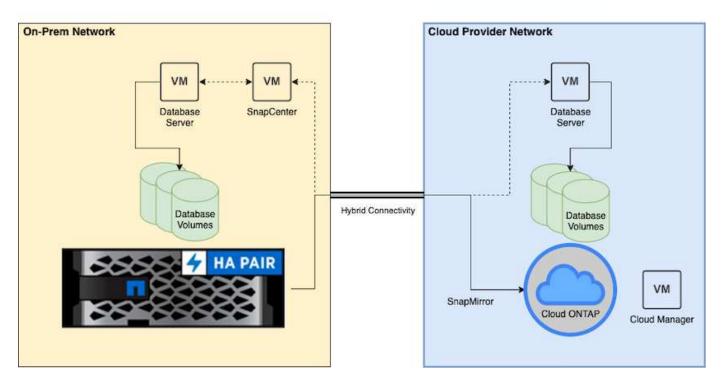
- Backup of an Oracle database across a Hybrid Cloud using SnapCenter
- SnapCenter- Clone DEV/TEST to AWS Cloud for an Oracle database

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 hybridcloud database workload environment.

SnapCenter installation and configuration

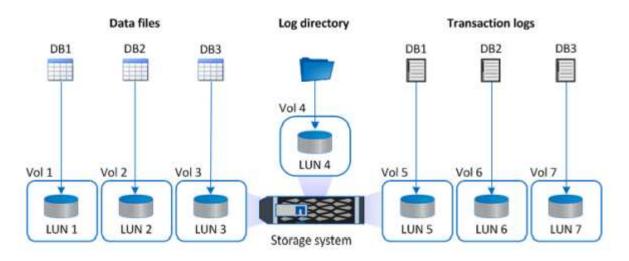
The NetApp SnapCenter tool is a Windows-based application that typically runs in a Windows domain environment, although workgroup deployment is also possible. It is based on a multitiered architecture that includes a centralized management server (the SnapCenter server) and a SnapCenter plug-in on the 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.

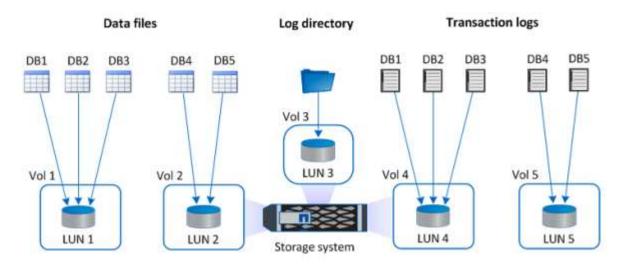
After careful pre-installation planning and consideration, click this SnapCenter installation workflow for details of SnapCenter installation and configuration.

On-premises database server storage configuration

Storage performance plays an important role in the overall performance of databases and applications. A welldesigned 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.

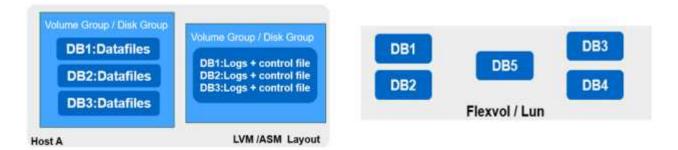


i.

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 particular 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 instruction to set up your Ansible control node for RedHat or CentOS machines: RedHat/CentOS Ansible Controller Setup.

Follow these instruction to set up your Ansible control node for Ubuntu or Debian machines: Ubuntu/Debian Ansible Controller Setup.

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.

	letApp Snap(Center®									•	s 0.	demo\administrator	SnapCenterAdmin	🖡 Sign Out
<		Global Sett	ings Polici	es Users and A	ccess	Roles	Credential	Software							
	Dashboard	Search b	y Name						Type	-				Add	
v	Resources		Name			11	Туре		Roles					Domain	
	Monitor		administrat	<u>or</u>			User		SnapCenterAdmin					demo	
			oradba				User		App Backup and Clone Admin					demo	
M	Reports		<u>sqldba</u>				User		App Backup and Clone Admin					demo	
A	Hosts														
ł.	Storage Systems														
	Settings														
A	Alerts														

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.
 - \circ 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.

п	NetApp SnapCenter® 🔹 🚱 Lemo\administrator SnapCenterAdmin 🗗 Sign Out											
<		Global Settings Policies Users and Access Roles Credential Software										
	Dashboard											
V	Resources	Global Settings										
٩	Monitor											
ííí	Reports	Hypervisor Settings 🚯						^				
Å	Hosts	VMs have ISCSI direct attached disks or NFS for all the hosts Update										
ł•	Storage Systems	Notification Server Settings 🟮						~				
靐	Settings	Configuration Settings 0						~				
▲	Alerts	Purge Jobs Settings 🚯						~				
		Domain Settings 0						~				
		CA Certificate Settings 🚯						~				

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.

	etApp	p SnapCenter®					1	•	-	8 -	⊥ demo\administrator	SnapCenterAdmin	🖡 Sign Out
>	Man	naged Hosts											
	Se	earch by Name		Add Host									
		Name	IL.	Host Type	Windows	•							
•		rhel2.demo.netapp.com		Host Name	sql-standby								
ŝ		sql1.demo.netapp.com		Credentials	Domain Admin	•	+						
A.				Select Plug-ins to Inst	tall SnapCenter Plug-ins Package 4.5 for Windows								
80 J					Microsoft Windows Microsoft SQL Server								
韢					Microsoft Exchange Server SAP HANA								
A					L) SAP HANA vrt, gMSA, Install Path, Custom Plug-Ins								
				Submit Cancel]								

4. After the plugin is installed on a Windows host, its Overall Status is shown as "Configure log directory."

	e tApp SnapC	Center	8					٠		8 -	👤 demo\adn	inistrator	SnapCent	terAdmin	🖡 Sign Out
		Mana	ged Hosts Disks Shares Initiator Groups is												
	Dashboard	Sea	rch by Name 🛛 🕅								+				More
2	Resources		Name	1E	Туре	System	Plug-in		Ve	rsion	Overall	štatus			
	Monitor		rhel2.demo.netapp.com		Linux	Stand-alone	UNIX, Oracle Database		4.5		Rur	ning			
			sql1.demo.netapp.com		Windows	Stand-alone	Microsoft Windows Server, Microsoft SQL Server		4.5		Rur	ning			
M	Reports		sql-standby.demo.netapp.com		Windows	Stand-alone	Microsoft Windows Server, Microsoft SQL Server		4.5		😑 Cor	figure log	directory		
A.	Hosts														
80 - S	Storage Systems														
÷=	Settings														
	Alerts														

5. Click the Host Name to open the SQL Server log directory configuration.

n Ne	tApp	SnapCenter®					• =	0 -	. demo\administrator	SnapCenterAdmin	🗊 Sign O	but
>	Mana	iged Hosts										×
	Sea	arch by Name		Host Details								
v		Name	11		sql-standby.demo.netapp.com	Alerts						
-		rhel2.demo.netapp.com sgl1.demo.netapp.com			10.221.2.56 Configure log directory	No Alerts						
<i>i</i> ii		sql-standby.demo.netapp.com		Host Type	Windows							
Δ.				System	Stand-alone							
֥				Credentials	Domain Admin 🖋							
華					SnapCenter Plug-ins package 4.5.0.6123 for Windows							
▲					Microsoft Windows Microsoft SQL Server <u>Remove Configure log directory</u>							
				More Options ; Po	rrt, gMSA, Install Path, Add Plug-Ins							
				Submit Cancel	Reset							

6. Click "Configure log directory" to open "Configure Plug-in for SQL Server."

Configure Plug-in for SQL Server											
Configure the log	backup directory for sql-standby.demo.netapp.com										
Configure host log	Configure host log directory										
Host log directory	dedicated disk directory path	Browse									
		Save	e								

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.

Configure Plug-in for SQL Server ×									
Configure the log backup directory for sql-standby.demo.netapp.com									
Configure host log directory									
Host log directory G:\	Browse								
Choose directory on NetApp Storage									
 sql-standby.demo.netapp.com G:\ System Volume Information 									
	Save	Close							



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.

NetApp Snap	Center@				♦ ≅ 0-	L demo\administrator	SnapCenterAdmin	🖡 Sign Out
<	Manag	d Hosts Disks Shares Initiator Groups iSCSI Session						
Dashboard	Searc	h by Name V				+ Add		More
Resources		Name IE	Туре	System	Plug-in	Version	Overall Status	
Monitor		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	
Reports		sql-standby.demo.netapp.com	Windows	Stand-alone	Microsoft Windows Server, Microsoft SQL Server	4.5	Running	
Hosts								
 Storage Systems 								
E Settings								
Alerts								

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.

NetApp Snap	Center®								٠	0 -	1 demo\administrator	SnapCenterAdmin	🗊 Sign Out		
<	Global Settin	ngs Policies	Users and Access	Roles C	redential	Software									
Dashboard	Search by t	Name	\supset				Type All	•					Remove		
Resources		Name			1E	Туре	Roles				Do	main			
Honitor		administrator				User	SnapCenterAdmin				de	mo			
		oradba				User	App Backup and Clone Admin				der	no			
Reports		sqldba				User	App Backup and Clone Admin				der	demo			
📥 Hosts															
Storage Systems															
Settings															
Alerts															

Assign Assets												
Asset	Type Host 💌	search										
	Asset Name	15										
	rhel2.demo.netapp.com											
	sql1.demo.netapp.com											
	sql-standby.demo.netapp.com											
		Save Close										

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.

						•	€.	L demo\administrator	SnapCenterAdmin	🖡 Sign Out	
											×
Add Host											
Host Type	Linux 👻										
Host Name	ora-standby										
Credentials	admin 👻	+	• 0								
Select Plug-ins to Insta	all SnapCenter Plug-ins Package 4.5 for Linux										
	 Oracle Database 										
	SAP HANA										
More Options : Por	rt, Install Path, Custom Plug-Ins										
Submit Cancel]										

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	8145 /opt/NetApp/snapcenter	6 6
	 Skip preinstall checks Add all hosts in the oracle RAC 	-
Custom Plug-ins	Choose a File Browse Upload	
	No plug-ins found.	*
	Save	ancel

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 de	termined 🕦	
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	
	Confirm and Submit	Close

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.

	Manage	d Hosts Disks Shares Initiator Groups iSCSI Session						
	-							
Dashboard	Searc	n by Name V						More
Resources		Name IE	Туре	System	Plug-in	Version	Overall Status	
Se Monitor		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	
Reports		sgl1.demo.netapp.com	Windows	Stand-alone	Microsoft Windows Server, Microsoft SQL Server	4.5	Running	
Hosts		sgl-standby.demo.netapp.com	Windows	Stand-alone	Microsoft Windows Server, Microsoft SQL Server	4.5	Running	
Storage Systems								
Settings								
A Alerts								

7. Assign the newly added host to the proper database management user ID (in our case, oradba).

II Ne	etApp SnapCenter®						8 -	L demo\administrator	SnapCenterAdmin	🕼 Sign O	Dut
>	Users and Access	Users/Groups	Details								×
	Search by Name	1	User Name oradba								
V	Name IE		Domain demo								
-	administrator		Roles App Backup and Clone Admin ×								
	oradba										
M	soldba	Assign A	ssets								
A		0							+ Assign	Unassign	
89 -			Asset Name Li	Т	pe			Asset Type			
÷2			10.0.0.1	D	ataOntap	Cluster		Storage Connect	ion	*	
A			192.168.0.101	D	ataOnta	Cluster		Storage Connect	ion		
-			admin					Credentials			
			Linux Admin					Credentials			
			Oracle Archive Log Backup					Policy			
			Oracle Full Online Backup					Policy			
			rhel2.demo.netaon.com					host		*	

Assi	gn Assets	×
Asset	t Type Host • search	
	Asset Name	ΨE
	ora-standby.demo.netapp.com	
	rhel2.demo.netapp.com	
	sql1.demo.netapp.com	
	sql-standby.demo.netapp.com	
	Save	e

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.

	etApp Snap(Center@	9						? ∙	👤 demo\oradba	App Backup and Clo	ne Admin	🖡 Sign Out
<		Oracle	Database 🕞										
	Dashboard	View	Database	Search databas	es V						Refre	🖈 sh Resources	New Resource Gro
9	Resources	- The	Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	0.			Last Backup	Overall S	itatus
•	Monitor		cdb2	Single Instance (Multitenant)	rhel2.demo.netapp.com							Not prot	ected
M	Reports												
h	Hosts												
	Hosts Storage Systems												
-													

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.

NetApp Snap						🖾 🚱 🕶 👤 demo\sqla	dba App Backup and Clone Admin 🗍 Sij	gn Out
<	Microson	t SQL Server 👻						-
Dashboard	View	Database 👻 search by n	ame V					esource Group
Resources	Til Ian	Name	Instance	Host	Last Backup	Overall Status	Туре	
Monitor		master	sql1	sql1.demo.netapp.com		Not available for backup	System database	
		model	sql1	sql1.demo.netapp.com		Not available for backup	System database	
Reports		msdb	sql1	sql1.demo.netapp.com		Not available for backup	System database	
📥 Hosts		tempdb	sql1	sql1.demo.netapp.com		Not available for backup	System database	
Storage Systems		tpcc	sql1	sql1.demo.netapp.com	09/14/2021 2:35:07 PM	Backup succeeded	User database	
\Xi Settings								
Alerts								

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.



II Ne	tApp SnapCenter®			٠	•	0-	👤 demo\sqldba	App Backup and Clone Admin	🗊 Sign Out
>	Microsoft SQL Server 🗨	Instance - Credentials							×
	search by name								Add Credential
0	Name	i The Microsoft SQL server or Windows credentials are neccessary to unlock the	ne selected instance. Click Refresh Resources to run a discovery with the associated	Auth.					. ×
-	sql-standby	Name	sql-standby						
<i>6</i> 1	sql1	Resource Group	None						
-		Policy	None						
•		Selectable	🖑 Not available for backup. DB is not on NetApp storage, auto-close is enabled	or in re	covery m	node.			
10 I.									
華									
A									

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.

	NetApp Snap(Center®					۰	-	••	👤 demo\sqidba	App Backup and Clone Admin	🗊 Sign Out
<		Microsoft S	QL Server 👻									
	Dashboard	View Ins	stance • search by name	V							Refresh Resources	New Resource Group
۲	Resources	Til Im	Name	Host	Resource Groups	Policies	Stat	e			Туре	
۲	Monitor		sql1	sql1.demo.netapp.com			Run	ning			Standalone (15.0.2000)	
a	Reports		sql-standby	sql-standby.demo.netapp.com			Run	ning			Standalone (15.0.2000)	
ň	Hosts											
20	Storage Systems											
÷	Settings											
▲	Alerts											

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.

 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 Manger. A default CVO deployment has intercluster LIFs configured automatically.

On-premises cluster:

ONTAP Sy	ystem Manager (Return to c	classic version)			Search actions, objects, and	d pages 🔍			? ‹›	-
DASHBOARD	Overview									
TORAGE V	IPspaces			+	Broadcast Domains	s				4
ETWORK ^	a anasas									
verview	Cluster	Broadcast Domains Cluster			Cluster	9000 MTU IP:	ipace: Cluster			
hernet Ports		Cluster			Default	1500 MTU IP	pace: Default			
Ports	Default	Storage VMs				on	Prem-01 e0a e0b e0c e0d e0e	eOf eOg eOh eOg-100 eOe-3	200 e0f-201	
ENTS & JOBS 🗸		svm_onPrem Broadcast Domains								
		Default								
verview										
ationships	Network Interface:	S							+ =	Filt
osts ^	Name	Status 🌻	Storage VM	IPspace	Address	Current Node	Current Port	Protocols	Туре	
AN Initiator Groups /Me Subsystem	onPrem-01_IC	0		Default	192.168.0.113	onPrem-01	e0b		Intercluster	
USTER ^	onPrem-01_mgmt1	0		Default	192.168.0.111	onPrem-01	e0c		Cluster/Node Mgmt	
verview	cluster_mgmt	0		Default	192.168.0.101	onPrem-01	e0a		Cluster/Node Mgmt	

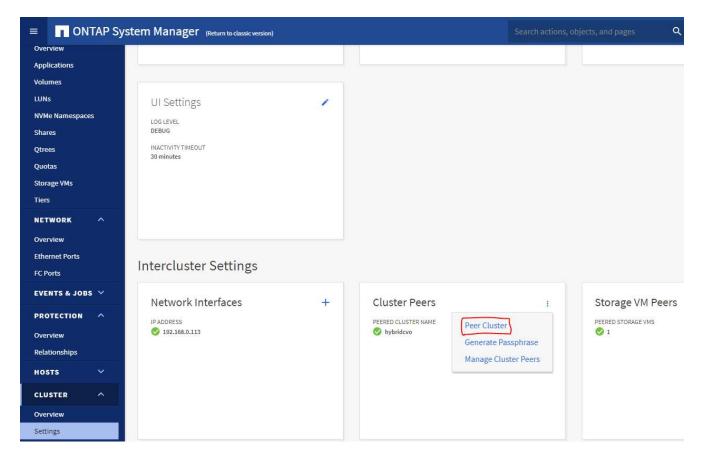
Target CVO cluster:

	TAP Sy	stem Manager				Search a	actions, objects, and pages	۹				? 🗘	-
DASHBOARD		Overview											
STORAGE		-											
NETWORK		IPspaces				+	Broadcast Doma	ains					
Overview Ethernet Ports		Cluster	Broadcast Domains Cluster				Cluster	9000 MTI	J IPspace: Cluster hybridcvo-01 e0b hybridcvo-02 e0b				
VENTS & JOBS	s ~ ~	Default	Storage VMs svm_hybridcvo				Default	9001 MT					
IOSTS			Broadcast Domains Default						hybridcvo-02 e0a				
LUSTER		Network Interfaces	3							+ Q Search	Land Land en		
										+	🛓 Download 🖙 Filter 🛛 🛛	Show / Hide Y	~
		Name ≑	Status	Storage VM	IPspace	A	uddress	Current Node	Current Port	Protocols		Through	
		Name 🗘	Status	Storage VM	IPspace Default		uddress	Current Node	Current Port	4			
		000000000 175		Storage VM		10				4	Туре	Through	;hput (I
		hybridcvo-02_mgmt1	۲	Storage VM	Default	10	0.221.2.104	hybridcvo-02	e0a	4	Type Cluster/Nade Mgmt	Through	;hput (1 0 0.02
		hybridcvo-02_mgmt1 inter_1	0	Storage VM	Default	10 10 10	0.221.2.104	hybridcvo-02 hybridcvo-01	e0a e0a	4	Type Cluster/Node Mgmt Intercluster,Cluster/Node Mgmt	Through	;hput (I 0

2. With the intercluster LIFs configured, cluster peering and volume replication can be set up by using dragand-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.

■ ONTAP S	ystem Manager (Return to classic version)		Search actions, objects, and pages
DASHBOARD	Volumes		
STORAGE ^	+ Add Telete Protect :	More	
Overview Applications	Name	rhel2_u03 All Volumes	
Volumes	onPrem_data rhel2_u01	Overview Snapshot Copies	Clone Hierarchy SnapMirror (Local or Remote)
NVMe Namespaces Shares	rhel2_u02		
Qtrees Quotas	rhel2_u03	status Online	Capacity
Storage VMs Tiers	8 sql1_data	style FlexVol	0% 10% 20% 30% 40% 50%
NETWORK ^	sql1_log	/rhel2_u03	SNAPSHOT CAPACITY
Overview Ethernet Ports	sql1_snapctr	storage vm svm_onPrem	0 Bytes Available 2.36 GB Used 2.36 GB Overflow
FC Ports	svm_onPrem_root	LOCAL TIER ONPrem_01_SSD_1	
PROTECTION V		SNAPSHOT POLICY default	Performance Hour Day Week
HOSTS ~		QUOTA Off TYPE	Latency
CLUSTER V		Read Write space reservation	1.5

5. Set the protection policy to Asynchronous. Select the destination cluster and storage SVM.

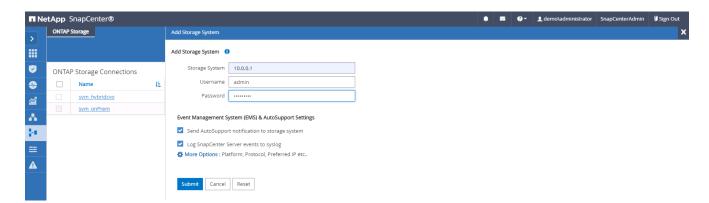
ONTAP System Manager (Return to classic version)		ctions, objects, and pages Q
BOARD	Protect Volumes	
	PROTECTION POLICY	
view	Asynchronous	×
ations	Source	
mes	Jource	
IS	CLUSTER	CLUSTER
e Namespaces	onPrem	hybridcvo
res	STORAGE VM	
ees	svm_onPrem	STORAGE VM
otas	SELECTED VOLUMES	svm_hybridcvo
rage VMs	rhel2_u03	 Destination Settings
s		2 matching labels
TWORK ^		VOLUME NAME
rerview		PREFIX
nernet Ports		vol_ <sourcevolume< td=""></sourcevolume<>
Ports		Override default storage s
ENTS & JOBS V		Overnoe default storage s
		Configuration Details
ROTECTION V		🔽 Initialize relationship 🛛 🧑
osts V		Enable FabricPool 👩
JSTER Y		
	Save Cancel	

6. Validate that the volume is synced between the source and target and that the replication relationship is healthy.

Volum	ies								
+ Add	🖥 Delete 🔍 Protect 🕴 !	More					₹F		
	Name	rhel2 U03 All Volumes							
	onPrem_data	_							
	rhel2_u01	Overview Snapshot Co	opies Clone Hierarchy	SnapMirror (Local or Remote)					
	rhel2_u02								
	1.10.00	Source	Destination	Protection Policy	Relationship Health	Relationship Status	Lag		
rhel2_u03		svm_onPrem:rhel2_u03	svm_hybridcvo:rhel2_u03_dr	MirrorAllSnapshots	🕑 Healthy	Mirrored	12 seconds		
	rhel2_u030923211942120311 8								

6. Add CVO database storage SVM to SnapCenter

- 1. Log into SnapCenter with a user ID with SnapCenterAdmin privileges.
- 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.

More Options		×
Platform	Cloud Volumes ON' 🗸 🗹 Secondary 🚺	
Protocol	HTTPS 🔻	
Port	443	
Timeout	60 seconds ()	
Preferred IP	0	
Save Cance	2	

4. Assign the storage systems to SnapCenter database management user IDs as shown in 3. SnapCenter host plugin installation.

NetApp	SnapCenter	8					9-	L demo\administrator	SnapCenterAdmin	🖡 Sign Oi
<	ONTA	P Storage								
Dashboard		ONTAP SVMs Search by Na	me						+ New	Deter
Resources	ONT	AP Storage Connections								
Monitor		Name	15	IP	Cluster Name	User Name	Platform	Control	er License	
Reports		svm hybridevo			10.0.0.1		CVO	0		
Hosts		svm_onPrem			192.168.0.101		CVO	~		
- Storage Sys	stems									
Settings										
A Alerts										

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.

	letApp Snap(Center®				0· 1	demo\oradba	Арр Вас	kup and Clo	ne Admin	🛿 Sign Out
<		Policies Credential									
	Dashboard	Oracle Database									
	Resources	Search by Name									Dolesa
-	Monitor	Name 15	Backup Type	Schedule Type	Replication		Ver	rification			
14	Reports	Oracle Archive Log Backup	LOG, ONLINE	Hourly	SnapMirror						
		Oracle Full Online Backup	FULL, ONLINE	Daily	SnapMirror						
A	Hosts										
÷.	Storage Systems										
	Settings										
▲	Alerts										

2. Click New to launch a new backup policy creation workflow or choose an existing policy for modification.

Modify Oracle [Database Backup	Policy	×				
1 Name	Provide a policy na	Provide a policy name					
2 Backup Type	Policy name	Oracle Full Online Backup	0				
3 Retention	Details	Backup all data and log files]				
4 Replication							
5 Script							
6 Verification							
7 Summary							
		Previous	Next				

3. Select the backup type and schedule frequency.

Modify Oracle	Database Backup Policy	×
1 Name	Select Oracle database backup options	-
2 Backup Type	Choose backup type	
3 Retention	Online backup	1
4 Replication	Otatafiles, control files, and archive logs	
5 Script	O Datafiles and control files	
6 Verification	○ Archive logs	
Summary	O Offline backup 🕕	
	◯ 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	
	O 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				×
1 Name	Retention settings				
2 Backup Type	Daily retention settings Data backup retention settings				
3 Retention	Total Snapshot copies to keep	7			
4 Replication	Keep Snapshot copies for	14	days		
5 Script	Archive Log backup retention settings O Total Snapshot copies to keep	7			
6 Verification	Keep Snapshot copies for	14	days		
7 Summary					
				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	×
1 Name	Select secondary replication options ()	
2 Backup Type	✓ Update SnapMirror after creating a local Snapshot copy.	
3 Retention	Update SnapVault after creating a local Snapshot copy.	
4 Replication 5 Script	Secondary policy label Daily The Daily Dai	
6 Verification		
7 Summary		
	Previous	Next

6. Specify any optional script to run before and after a backup run.

Modify Oracle [Database Backup	Policy				×			
1 Name	Name Specify optional scripts to run before and after performing a backup job								
2 Backup Type	Prescript full path	/var/opt/snapce	nter/spl/scripts/	Enter Prescript	path				
3 Retention	Prescript arguments								
4 Replication	Postscript full path Postscript arguments	/var/opt/snapce	nter/spl/scripts/	Enter Postscrip	t path				
5 Script	Script timeout	60 s	iecs						
6 Verification									
7 Summary									
					Previous	Next			

7. Run backup verification if desired.

Modify Oracle I	Database Backup	Policy			×					
1 Name	Select the options	to run backu	p verification							
2 Backup Type	Run Verifications for following backup schedules									
3 Retention	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.									
A Replication	🗌 Daily									
5 Script	Verification script	commands								
6 Verification	Script timeout	60	secs							
	Prescript full path	/var/opt/sna	pcenter/spl/scripts/	Enter Prescript path						
Summary	Prescript arguments	Choose optional arguments								
	Postscript full path	/var/opt/sna	pcenter/spl/scripts/	Enter Postscript path						
	Postscript arguments	Choose optional arguments								
				Previous	lext					

8. Summary.

Name	Summary	
Backup Type	Policy name	Oracle Full Online Backup
Retention	Details	Backup all data and log files
Retention	Backup type	Online backup
Replication	Schedule type	Daily
	RMAN catalog backup	Disabled
Script	Archive log pruning	None
Verification	On demand data backup retention	None
	On demand archive log backup retentio	n None
Summary	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

Create a database log backup policy for Oracle

- 1. Log into SnapCenter with a database management user ID, click Settings, and then click Polices.
- 2. Click New to launch a new backup policy creation workflow, or choose an existing policy for modification.

New Oracle Dat	tabase Backup Po	licy	×
1 Name	Provide a policy na	ame	
2 Backup Type	Policy name	Oracle Archive Log Backup	1
3 Retention	Details	Backup Oracle archive logs	
4 Replication			
5 Script			
6 Verification			
7 Summary			
		Previous	Next

3. Select the backup type and schedule frequency.

New Oracle Da	atabase Backup Policy	×
1 Name	Select Oracle database backup options	^
2 Backup Type	Choose backup type	1
3 Retention	Online backup	
4 Replication	O Datafiles, control files, and archive logs	
5 Script	O Datafiles and control files	
6 Verification	Archive logs	1
7 Summary	O Offline backup Mount	
	◯ Shutdown	
	Save state of PDBs 0	
	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.	
	O On demand	
	Hourly	
	O Daily	*
	Previous Next	

4. Set the log retention period.

New Oracle Da	atabase Backup Policy			>
1 Name	Retention settings ()			
2 Backup Type 3 Retention	Hourly retention settings Data backup retention settings ①	7		
4 Replication 5 Script	O Keep Snapshot copies for Archive Log backup retention settings	14	days	
6 Verification	O Total Snapshot copies to keep Keep Snapshot copies for	7 7 ≎	days	
7 Summary				
				Previous Next

5. Enable replication to a secondary location in the public cloud.

New Oracle Dat	tabase Backup Policy	×
1 Name	Select secondary replication options 🚯	
2 Backup Type	🗹 Update SnapMirror after creating a local Snapshot copy.	
3 Retention	Update SnapVault after creating a local Snapshot copy.	
Replication S Script	Secondary policy label Hourly Error retry count 3	
6 Verification		
7 Summary		
	Previous	Next

6. Specify any optional scripts to run before and after log backup.

New Oracle Dat	New Oracle Database Backup Policy						
Name Specify optional scripts to run before and after performing a backup job							
2 Backup Type	Prescript full path	/var/opt/snapcenter/spl/scripts/	Enter Prescript path				
3 Retention	Prescript arguments						
4 Replication	Postscript full path Postscript	/var/opt/snapcenter/spl/scripts/	Enter Postscript path				
5 Script	arguments						
6 Verification	Script timeout	60 secs					
7 Summary							
				Previous Next			

7. Specify any backup verification scripts.

New Oracle Dat	tabase Backup Po	licy		×					
1 Name	Select the options to run backup verification								
2 Backup Type	Run Verifications f	or following backup schedules							
3 Retention		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.							
A Replication	Verification script	commands							
5 Script	Script timeout	60 secs							
6 Verification	Prescript full path	/var/opt/snapcenter/spl/scripts/	Enter Prescript path						
7 Summary	Prescript arguments								
	Postscript full path	/var/opt/snapcenter/spl/scripts/	Enter Postscript path						
	Postscript arguments	Choose optional arguments							
	8								
			Previous	lext					

8. Summary.

Name	Summary	
Backup Type	Policy name	Oracle Archive Log Backup
3 Retention	Details	Backup Oracle archive logs
	Backup type	Online backup
Replication	Schedule type	Hourly
	RMAN catalog backup	Disabled
5 Script	Archive log pruning	None
6 Verification	On demand data backup retention	None
	On demand archive log backup retentio	n None
7 Summary	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

Create a full database backup policy for SQL

1. Log into SnapCenter with a database management user ID, click Settings, and then click Polices.

	letApp Snap	Center®			•	-	⊖- Lden	io\sqldba	App Backup and O	Ione Admin	🖡 Sign Out
<		Policies Credential									
	Dashboard	Microsoft SQL Server -									
10000	Resources	Search by Name					+ New				
	Monitor	Name	Backup Type	Schedule Type	Replication				cation		
		There is no match for your search or data is not availab	ble.								
	Reports										
	Hosts										
24	Storage Systems										
盐	Settings										
▲	Alerts										

2. Click New to launch a new backup policy creation workflow, or choose an existing policy for modification.

New SQL Serve	r Backup Policy		×
1 Name	Provide a policy na	ame	
2 Backup Type	Policy name	SQL Server Full Backup	1
3 Retention	Details	Backup all data and log files	
4 Replication			
5 Script			
6 Verification			
7 Summary			
		Previous	Next

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 Serve	r Backup Policy ×
1 Name	Select SQL server backup options
2 Backup Type	Choose backup type
3 Retention	Full backup and log backup
4 Replication	 Full backup Log backup
5 Script	Copy only backup
6 Verification	Maximum databases backed up per Snapshot copy: 100
7 Summary	Availability Group Settings
	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
	Previous Next

4. Set the backup retention period.

New SQL Serve	er Backup Policy	×
1 Name	Retention settings	
2 Backup Type	Retention settings for up-to-the-minute restore operation 🚯	
3 Retention	Keep log backups applicable to last 7 full backups	
(4) Replication	O Keep log backups applicable to last 14 days	
5 Script		
6 Verification	Full backup retention settings 🚯	
7 Summary	Total Snapshot copies to keep Keep Snapshot copies for 14 days	
	Previous	Next

5. Enable backup copy replication to a secondary location in cloud.

New SQL Serve	r Backup Policy	×
1 Name	Select secondary replication options 1	
2 Backup Type	✓ Update SnapMirror after creating a local Snapshot copy.	
3 Retention	Update SnapVault after creating a local Snapshot copy.	
Replication S Script 6 Verification 7 Summary	Secondary policy label Daily Error retry count 3	
	Previous Next	

6. Specify any optional scripts to run before or after a backup job.

New SQL Serve	r Backup Policy		×
1 Name	Specify optional se	cripts to run before performing a backup job	
2 Backup Type	Prescript full path		
3 Retention	Prescript arguments	Choose optional arguments	
4 Replication	-	cripts to run after performing a backup job	
	Postscript full path		
5 Script	Postscript arguments	Choose optional arguments	
6 Verification	Script timeout	60 secs	
7 Summary			
		Previous	Next

7. Specify the options to run backup verification.

New SQL Serve	r Backup Policy	×
1 Name	Select the options to run backup verification	Â
2 Backup Type	Run verifications for the following backup schedules	
3 Retention	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.	1
4 Replication	Daily	1
5 Script	Database consistency checks options	1
6 Verification	 Limit the integrity structure to physical structure of the database (PHYSICAL_ONLY) Suppress all information message (NO_INFOMSGS) 	1
7 Summary	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) 	1
	Log backup	1
	🗌 Verify log backup. 🚯	1
	Verification script settings	
	Script timeout 60 secs	
	Previous	

8. Summary.

New SQL Serve	r Backup Policy		×
1 Name	Summary		
2 Backup Type	Policy name	SQL Server Full Backup	
3 Retention	Details	Backup all data and log files	
3 Retention	Backup type	Full backup and log backup	
4 Replication	Availability group settings	Backup only on preferred backup replica	
	Schedule Type	Daily	
5 Script	UTM retention	Total backup copies to retain : 7	
6 Verification	Daily Full backup retention	Total backup copies to retain : 7	
	Replication	SnapMirror enabled , Secondary policy label: Daily , Error retry count: 3	
7 Summary	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 Finis	h

Create a database log backup policy for SQL.

1. Log into SnapCenter with a database management user ID, click Settings > Polices, and then New to launch a new policy creation workflow.

New SQL Serve	r Backup Policy		×
1 Name	Provide a policy na	ame	
2 Backup Type	Policy name	SQL Server Log Backup	1
3 Retention	Details	Backup SQL server log]
4 Replication			
5 Script			
6 Verification			
7 Summary			
		Previous	Next

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 Serve	er Backup Policy ×
1 Name	Select SQL server backup options
2 Backup Type	Choose backup type
3 Retention	Full backup and log backup
4 Replication	 Full backup Log backup
5 Script	Copy only backup
6 Verification	Maximum databases backed up per Snapshot copy: 100
7 Summary	Availability Group Settings
	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
	O Daily O Weekly
	○ Monthly
	Previous Next

3. SQL server data backup policy defines the log backup retention; accept the defaults here.

New SQL Serve	r Backup Policy ×
1 Name	Log backup retention settings
2 Backup Type	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
3 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.
(4) Replication	
5 Script	
6 Verification	
7 Summary	
	Previous Next

4. Enable log backup replication to secondary in the cloud.

New SQL Serve	r Backup Policy	×
1 Name	Select secondary replication options ()	
2 Backup Type	Update SnapMirror after creating a local Snapshot copy.	
3 Retention	Update SnapVault after creating a local Snapshot copy.	
4 Replication	Secondary policy label Hourly	0
6 Verification		
7 Summary		
		Previous Next

5. Specify any optional scripts to run before or after a backup job.

New SQL Serve	r Backup Policy		×
1 Name	Specify optional so	ripts to run before performing a backup job	
2 Backup Type	Prescript full path		
3 Retention	Prescript arguments	Choose optional arguments	
4 Replication	Specify optional so	ripts to run after performing a backup job	
	Postscript full path		
5 Script	Postscript arguments	Choose optional arguments	
6 Verification	Script timeout	60 secs	
7 Summary			
		Previous	Next

6. Summary.

New SQL Serve	r Backup Policy		×
1 Name	Summary		
2 Backup Type	Policy name	SQL Server Log Backup	
3 Retention	Details	Backup SQL server log	
3 Retention	Backup type	Log transaction backup	
4 Replication	Availability group settings	Backup only on preferred backup replica	
	Schedule Type	Hourly	
5 Script	Replication	SnapMirror enabled , Secondary policy label: Hourly , Error retry count: 3	
6 Verification	Backup prescript settings	undefined Prescript arguments:	
7 Summary	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 Finist	h

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.

	Center								l Clone Admin	
	Oracle	Database 🝷								
Dashboard		Database	- Search databases	V					Refresh Resources	New Resource Gr
Resources	19	Name	Oracle Database Type	Host/Cluster	Resource Group	Policies		Last Backup	Overall Stat	us
🛞 Monitor		cdb2	Single Instance (Multitenant)	rhel2.demo.netapp.com					Not protect	ed
Reports										
Hosts										
Hosts										

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.

ΠN	letApp Si	napCenter®								٠	2	0-	👤 demo\oradba	App Backup and Clone Admin	🖡 Sign Out
>		itabase 👻	New Resource Group												×
	Search	databases													
U	15.16	Name		2	3	4	5	6							
۲		cdb2	Name R	lesources	Policies	Verification	Notification	Summary							
aii			Provide a name	and tags for th	he resource g	iroup									
Å			Name	rhel2_cdb2					0						
ł۹.			Tags	orafullbkup					 0						
#			Use custom nam		shot copy										
▲			\$CustomText × rhel2_cdb2						 _						
			Backup settings												
			Exclude archive log destinations from backup					x 🔹 + 🟮							

3. Add database resources to the resource group.

n N	l etApp Sri	napCenter®		🛊 🖀 🔂 • 主 demo\oradba App Backup and Clone Admin 🕼 Sign Out
>		itabase 🔽	New Resource Group	×
	Search	databases		
U	IF P	Name	1 2 3 4 5 6 Name Resources Policies Verification Notification Summary	
٩		cdb2	Name Resources Policies Verification Nourication Summary	
.			Add resources to Resource Group	
A			Host	
<u>اور</u>			All Available Resources Selected Resources	
ŧ			(search available resources)	
A			cdb2 (rhel2.demo.netapp.com)	
			>	

4. Select a full backup policy created in section 7 from the drop-down list.

	n Ne	tApp SnapCenter®				٠	• =	♦ ≅ 0 -	🜲 🔤 💽 🔹 🛓 demo\oradba	🌲 📓 😨 🕈 L demo\oradba App Backup and Clone Admin	🌲 🔤 🥹 🔹 demo\oradba App Backup and Clone Admin 🖡 Sig	
	>	Oracle Database 👻	New Resource Group									
		Search databases										
Ī	9	lF ⊨ Name		4 5	6							
1	۵	cdb2	Name Resources Polici	ies Verification Notification	Summary							
1	ñĩ		Select one or more policies and confi	igure schedules								
	A		Oracle Full Online Backup	· + 0								
	-		Configure schedules for selected poli	icies								
1	E		Policy IL	Applied Schedules	Configure Schedules							
	A		Oracle Full Online Backup	None	+							
			Total 1									

5. Click the (+) sign to configure the desired backup schedule.

Daily										
Start date 09/10/2021 2:32 PM)						
Z Expires on	Expires on 12/31/2021 2:32 PM			1						
Repeat every	1	days		<		Dece	mber	2021		>
hepearerery				Su	Mo	Tu	We	Th	Fr	Sa
				28	29	30	1	2	3	4
				5	б	7	8	9	10	11
				12	13	14	15	16	17	18
				19	20	21	22	23	24	25
			_	26	27	28	29	30	31	1
. The sector	المغاجب وحاديات	محمد ما الم فا		2	3	4	5	6	7	8
zone.	edules are trij	ggered in tr	ie snap				Ø			
				i	_			-	-	
					Ca	ncel	100	ОК		

6. Click Load Locators to load the source and destination volume.

n Ne	tApp Sna	pCenter®						٠	9 -	👤 demo\oradba	App Backup and Clone Admin	🖡 Sign Out
>	Oracle Data	base 👻	New Resource Group									×
	Search da	itabases										
U	15.16	Name		-0		5	6					
•		cdb2	Name	Resources Policie	s Verification	Notification	Summary					
세 사			Load secondary verify backups o Secondary sto									
30 			Source Volume		Dest	ination Volume						
莘			svm_onPrem:rh	hel2_u02	SVI	m_hybridcvo:rhel2_u	02_dr •					
A			Policy 11	erification schedules Schedule Type	Applied Schedules		Configure Schedules					
			There is no mate	tch for your search or data is no	t available.							

7. Configure the SMTP server for email notification if desired.

ΠN	etApp Si	napCenter®									٠		0 -	👤 demo\oradba	App Backup and Clone Admin	🖡 Sign Out
>		atabase 👻	New Reso	ource Group												×
	Search	databases	A If you serve		notifications for schee	duled or on deman	d jobs, an SMTP serv	er must be configure	d. Continue to the S	ummary page to save	our information, an	d then go	to Settir	ngs>Global Settings>	Notification Server Settings to co	infigure the SMTP
U	17.19	Name		_	_	_	_									
۲		cdb2		0—				6	6							
a il				Name	Resources	Policies	Verification	Notification	Summary							
Å					ail settings 🕦											
÷٩.				Select the serv	vice accounts or peopl	e to notify regardin	g protection issues.									
橆				Email preferen	nce Never		-									
A				From	From email											
-				То	Email to											
				Subject	Notification											
				🗌 Attach job r	report											

8. Summary.

	etApp Sn	apCenter®				🏚 🔤 🥹 🔹 🔒	no\oradba App Backup and Clone Admin	🖡 Sign Out
>	Oracle Dat	abase 👻	New Resource Group					×
	Search d	latabases						
U	1E IN	Name	0-0-		6			
•		cdb2	Name Resources	Policies Verification Notification	Summary			
ណ៍			Resource group name	rhel2_cdb2				
Α.			Tags	orafullbkup				
24			Policy	Oracle Full Online Backup: Daily				
#			Plug-in Verification enabled for policy	SnapCenter Plug-in for Oracle Database None				
 A			Send email	No				
-								
	Total 1						Previor	us Finish

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.

NetApp Snape	letApp SnapCenter®								App Backup and Clone Admin	🖡 Sign Out
<	Oracle Database 👻	de Database 🖌								
Dashboard	View Resource Group	Search resource g	roup V							New Resource Group
Resources	Name	Resources	Tags	Policies				Last Backup	Overall Status	
Honitor	rhel2_cdb2	1	orafullbkup	Oracle Full Online Backup						
Reports										
📥 Hosts										
Storage Systems										
Alerts										

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.

n Ne	etApp SnapCenter®						٠	≥ (🗧 👤 demo\oradba	App Backup and Clone Admin	🖡 Sign Out
>	Oracle Database 👻	New Resource Group									×
	Search resource groups										
	Name	1 2 Name Resources	Policies Verification	5 Notification	5 Summary						
٠	rhel2_cdb2	Name Resources	Policies Verification	Notification	summary						
ай		Provide a name and tags for the	e resource group								
Α.		Name rhel2_cdb2_log				0					
80 -		Tags oralogbkup				0					
橆		Use custom name format for Snaps	hot copy								
A		\$CustomText × rhel2_cdb2_log									
		Backup settings									
		Exclude archive log destinations from backup		×	* + ()						

3. Add database resources to the resource group.

п	NetApp SnapCenter®		🌲 🔄 🚱 🐐 💄 demo\oradba App Backup and Clone Admin 🗿 Sign Out
>	Oracle Database 👻	New Resource Group	
	Search resource groups		
0			
•	rhel2_cdb2	Name Resources Policies Verification Notification Summary	
a i		Add resources to Resource Group	
٨		Host	
34		All Available Resources Selected Resources	
=		Available resources	
A		cdb2 (rhel2.demo.netapp.com)	
		>	
		ĸ	
	Total 1		Previous Next

4. Select a log backup policy created in section 7 from the drop-down list.

ΠN	etApp SnapCenter®		٠	0 •	👤 demo\oradba	App Backup and Clone Admin	🖡 Sign Out
>	Oracle Database 👻	New Resource Group					×
	Search resource groups						
	Name	1 2 3 4 5 6 Name Resources Policies Verification Notification Summary					
٠	rhel2_cdb2	Name Resources Policies Verification Notification Summary					
		Select one or more policies and configure schedules					
A		Oracle Archive Log Backup -					
3 0		Oracle Full Online Backup ✔ Oracle Archive Log Backup S					
÷		Policy IE Applied Schedules Configure Schedules					
▲		Oracle Archive Log Backup None +					
		Total 1					
	Total 1					Pre	vious Next

5. Click on the (+) sign to configure the desired backup schedule.

Add schedules for policy Oracle Archive Log Backup ×								
Hourly								
Start date	09/10/2021 3:00 PM							
Z Expires on	12/31/2021 3:00 PM							
Repeat every	1 hours 0 mins							
i The schedu zone.	les are triggered in the SnapCenter Server time							
	Cancel OK							

6. If backup verification is configured, it displays here.

	letApp SnapCenter®		٠	8 -	👤 demo\oradba	App Backup and Clone Admin	🖡 Sign Out
>	Oracle Database 👻	New Resource Group					×
	Search resource groups						
0	Name						
•	rhel2_cdb2	Name Resources Policies Verification Notification Summary					
A		Configure verification schedules					
34		Policy Li Schedule Type Applied Schedules Configure Schedules					
==		There is no match for your search or data is not available.					
 A							
-							
		Total 0					
	Total 1					Pre	vious Next

7. Configure an SMTP server for email notification if desired.

II N	etApp SnapCenter®	🌲 🗷 😌 🔹 🕹 demoloradba 🛛 App Backup and Clone Admin	l Sign Out
>	Oracle Database 👻	New Resource Group	×
	Search resource groups	If you want to send notifications for scheduled or on demand 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 configured.	ure the SMT
0	Name		
•	rhel2_cdb2		
1		Name Resources Policies Verification Notification Summary	
٨		Provide email settings 0	
} -1		Select the service accounts or people to notify regarding protection issues.	
÷		Email preference Never *	
A		From From email	
		To Email to	
		Subject Notification	
		Attach job report	
	Total 1	Previou	s Next

8. Summary.

	etApp SnapCenter®					• =	🕑 🔹 👤 demo\oradba	App Backup and Clone Admin	🖡 Sign Out
>	Oracle Database 👻	New Resource Group							>
	Search resource groups			-					
•	Name	0-2-		5					
٢	rhel2_cdb2	Name Resources	Policies Verification	Notification	Summary				
		Resource group name	rhel2_cdb2_log						
٨		Tags	oralogbkup						
24		Policy	Oracle Archive Log Backup: Ho						
≢		Plug-in	SnapCenter Plug-in for Oracle	e Database					
 A		Verification enabled for policy Send email	None						
	Total 1							Prev	ious Finish

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.

n Ne	etApp SnapCenter®				٠	⊠ 6	• l demo\sqldba	App Backup and Clone Admin	🖡 Sign Out
>	Microsoft SQL Server 👻	New Resource Group							×
	search by name								
U	Name		2 3 4 5 6						
٠	master	Name	Resources Policies Verification Notification Summary						
2 21	model	Provide a par	me and tags for the resource group						
*	msdb tempdb								
34	tpcc	Name	sql1_tpcc						
•• ==		Tags	sqlfullbkup name format for Snapshot copy	0					
		\$CustomTex							
A		sql1_tpcc							
	Total 5								evious Next

2. Select the database resources to be backed up.

n Ne	etApp SnapCenter®		🌲 📓 🚱 ד 👤 demo\sqldba App Backup and Clone Admin 🖡 Sign Out
>	Microsoft SQL Server 👻	New Resource Group	
	search by name		
•	Name		
•	master	Name Resources Policies Verification Notification Summary	
	model		
M	msdb	Add resources to Resource Group	
Δ.	tempdb	Host Resource Type SQL Server Instance	
34	tpcc	All • Databases • sql1 •	
-		Available Resources Selected Resources	
莘		search available resources	
A		Auto select all the resources from the same storage volume 0	
		tpcc (sql1)	
		>	
		K	
	Total 5		Previous Next

3. Select a full SQL backup policy created in section 7.

ΠN	etApp SnapCenter®		🌲 🦉 🕈 💄 demo\sqldba App Backup and Clone Admin 🛛 🖡 Sign Out
>	Microsoft SQL Server	New Resource Group	
	search by name		
	Name		
•	master	Name Resources Policies Verification Notification Summary	
۲	model	Select one or more policies and configure schedules	
	msdb	SQL Server Full Backup	
<u>*</u>	tempdb tpcc	✓ SQL Server Full Backup	
<u>اور</u>		SQL Server Log Backup S	
##		Policy Li Applied Schedules Configure Schedules	
▲		SQL Server Full Backup None +	
		Total 1	
		Use Microsoft SQL Server scheduler 0	
	Total 5		Previous Next

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								
Z Expires on	12/31/2021 6:20 PM								
Repeat every	1 days								
i The schedu zone.	les are triggered in the SnapCenter Server time								
	Cancel OK	I							

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.

n Ne	tApp SnapCenter®		٠	0 -	👤 demo\sqldba	App Backup and Clone Admin	🖡 Sign Out
>	Microsoft SQL Server 👻	New Resource Group					×
	search by name						
	Name						
	master	Name Resources Policies Verification Notification Summary					
	model	Select the verification servers					
	msdb						
*	tempdb	Verification server Select one or more servers					
20	tpcc	Load secondary locators to					
₩ •		Verify backups on secondary Load locators Secondary storage location: SnapVault or SnapMirror Source Volume Destination Volume sym_onPremsql1_data sym_hybridcoosql1_data_dr sym_onPremsql1_log sym_hybridcoosql1_data_dr Configure verification schedules sym_hybridcoosql1_log_dr Policy is Schedule Type Applied Schedules There is no match for your search or data is not available.					
	Total 5					Prev	lious Next

6. Configure the SMTP server for email notification if desired.

	etApp SnapCenter®	🌲 😆 😔 🔹 1 demolsqlidba 🛛 App Backup and Clone Admin 🛛 🖉 Sign O	ut
>	Microsoft SQL Server	New Resource Group	×
	search by name	If you want to send notifications for scheduled or on demand 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 server.	SMT
0	Name		
	master		
	model	Name Resources Policies Verification Notification Summary	
a ii	msdb		
Δ.	tempdb	Provide email settings 👔	
24 C	tpcc	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	
	Total 5	Previous N	ĸt

7. Summary.

II N	etApp SnapCenter®					٠	8	8 -	👤 demo\sqldba	App Backup and Clone Admin	🖡 Sign Out
>	Microsoft SQL Server	New Resource Group									2
	search by name										
•	Name		-0	5							
۲	There is no match for your search or data is not available.	Name Resources	Policies Verification	Notification	Summary						
ай		Resource group name	sql1_tpcc								
A.		Tags	sqlfullbkup								
34		Policy	SQL Server Full Backup: Daily								
- =		Plug-in	SnapCenter Plug-in for Micro	osoft SQL Server							
		Verification Server Verification enabled for policy	None								
A		Send email	None								
	Resources are not found. Click Refresh Resources to discover databases in the database view or create new resource group										
	on the discovered databases from the resource view.									Prev	ious Finish

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.

II Ne	tApp SnapCenter®			• •	s 0-	👤 demo\sqldba	App Backup and Clone Admin	🖡 Sign Out
>	Microsoft SQL Server	New Resource Group						×
	search by name							
U	Name	1 2 3 4 5 6 Name Resources Policies Verification Notification Summary						
	sql1_tpcc	Name Resources Policies Venication Notrication Summary						
a i		Provide a name and tags for the resource group						
A		Name sql1_tp:cc_log	0					
89 -		Tags sqllogbkup	0					
1 25		Use custom name format for Snapshot copy ScustomText						
		sql1_tpc_log						
	Total 1							vious Next

2. Select the database resources to be backed up.

	NetApp SnapCenter®		٠	 Ø ▼ demo\sqldba	App Backup and Clone Admin	🛿 Sign Out
>	Microsoft SQL Server 👻	New Resource Group				×
	search by name					
•	Name					
۲	sql1_tpcc	Name Resources Policies Verification Notification Summary				
a i		Add resources to Resource Group				
А		Host Resource Type SQL Server Instance				
20		All • Databases • sql1 •				
=		Available Resources Selected Resources				
▲		Auto select all the resources from the same storage volume tpcc (sql1)				
		>				
		"				
	Total 1				Pre	vious Next

3. Select a SQL log backup policy created in section 7.

n Ne	tApp SnapCenter®		● 🜌 🚱 - L demo\sqldba App Backup and Clone Admin I Sign Out
>	Microsoft SQL Server	New Resource Group	×
	search by name		
0	Name		
٠	sql1_tpcc	Name Resources Policies Verification Notification Summary	
a il		Select one or more policies and configure schedules	
A		SQL Server Log Backup 🔹 🕂 🕕	
} 4		SQL Server Full Backup SQL Server Log Backup S	
韢		Policy IL Applied Schedules Configure Schedules	
▲		SQL Server Log Backup None +	
		Total 1 Use Microsoft SQL Server scheduler	
	Total 1		Previous Next

4. Add exact timing for the backup as well as the frequency.

III Ne	tApp SnapCenter®		٠	8	0 -	⊥ demo\sqldba	App Backup and Clone Admin	🗊 Sign (Dut
>	Microsoft SQL Server 🚽	New Resource Group							×
	search by name								
Ø	Name								
•	sql1_tpcc	Name Resources Policies Verification Notification Summary							
※i ト ÷ ▲		Select one or more policies and configure schedules SQL Server Log Backup • • • Configure schedules for selected policies Policy Is Applied Schedules Configure Schedules SQL Server Log Backup Hourly: Repeat every 1 hours • • • Total 1 Use Microsoft SQL Server scheduler • • •							
	Total 1						Pre	vious N	lext

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.

	Net	:App SnapCenter®		۰	0-	👤 demo\sqldba	App Backup and Clone Admin	🖡 Sign Out
>		Microsoft SQL Server	New Resource Group					×
		search by name						
		Name						
2		sql1_tpcc	Name Resources Policies Verification Notification Summary					
1			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					
4	•		Source Volume Destination Volume					
			svm_onPrem:sql1_data svm_hybridcvo:sql1_data_dr 👻					
			svm_onPrem:sql1_log svm_hybridcvo:sql1_log_dr					
			Configure verification schedules					
			Policy Ii Schedule Type Applied Schedules Configure Schedules					
			There is no match for your search or data is not available.					
								-
		Total 1					Prev	rious Next

6. Configure the SMTP server for email notification if desired.

п	NetApp SnapCenter®	🌲 📓 🚱 🔹 🛔 demolsajibba 🛛 App Backup and Clone Admin 🕷	Sign Out
>	Microsoft SQL Server 👻	New Resource Group	×
	search by name	🔺 If you want to send notifications for scheduled or on demand 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 configured.	ure the SMTF
0	Name		
	sql1_tpcc		
1		Name Resources Policies Verification Notification Summary	
٨		Provide email settings 🕚	
÷.		Select the service accounts or people to notify regarding protection issues.	
幸		Email preference Never -	
A		From From email	
_		To Email to	
		Subject Notification	
		Attach job report	
	Total 1	Previous	Next

7. Summary.

	letApp SnapCenter®				• =	. demo\sqldba	App Backup and Clone Admin	🖡 Sign Out
>	Microsoft SQL Server	New Resource Group						2
	search by name							
	Name	0-0-	-3	-6				
•	sql1_tpcc	Name Resources	Policies Verification Notification	Summary				
*		Resource group name Tags	sql1_tpcc_log sqllogbkup					
		Policy	SQL Server Log Backup: Hourly					
-		Plug-in	SnapCenter Plug-In for Microsoft SQL Server					
華		Verification Server	None					
▲		Verification enabled for policy Send email	None					
		Serio eman	NU					
	Total 1						Prev	rious Finish

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.

NetApp Snap	App SnapCenter®						App Backup an	d Clone Admin	🗊 Sign Out
2	Jobs	Schedules							
Dashboard	searc	h by name	<u>-</u>						
Resources	Jobs - I	Filter							
Monitor	ID	Status	Name	Start date		End	i date Ov	mer	
窳 Reports	532	~	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/14/2021 8:35:01 PM 🛱		09/14/2021 8:37:10 P	M 🛱 de	no\sqldba	
Hosts	528	~	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/14/2021 7:35:01 PM 🛱		09/14/2021 7:37:09 P	M 🛱 de	no\sqldba	
	524	~	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/14/2021 6:35:01 PM 🛱		09/14/2021 6:37:08 P	M 🛱 de	no\sqldba	
 Storage Systems 	521	~	Backup of Resource Group 'sql1_tpcc' with policy 'SQL Server Full Backup'	09/14/2021 6:25:01 PM 🛱		09/14/2021 6:27:14 P	M 🛱 dei	mo\sqldba	
Settings	517	~	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/14/2021 5:35:01 PM 🛱		09/14/2021 5:37:09 P	M 🛱 der	no\sqldba	
Alerts	513	*	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/14/2021 4:35:01 PM 📛		09/14/2021 4:37:08 P	M 🛱 de	no\sqldba	
	509	~	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/14/2021 3:35:01 PM 🛱		09/14/2021 3:37:10 P	M 🛱 de	mo\sqldba	
	503	~	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/14/2021 2:35:01 PM 🛱		09/14/2021 2:37:09 P	M 🛱 de	no\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.

I Ne	tApp Snap(Center®					• ≊ 9- 1	demo\oradba	App Backup and Clone Ac	lmin 🌗 Sigr	h Out
	Oracle Databas	se 🔽	cdb2 Topology								
	Search datab	bases							and the second second second second second second second second second second second second second second second	otect R	H lefresh
	17 1	Name	Manage Copies								
		cdb2	197 Backups								
	層	cdb2dev	197 Backups 3 Clones				Summary Card				
		cdb2dr	Mirror copies					394 Backu			
	a	cdb2dr2	Local copies						ta Backups		
	1	cdb2test						300 LO	g Backups		
			Primary Backup(s)					O I	rie 1 @ me Cons Restore Mour		
			Backup Name	Count	Туре	17 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	

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



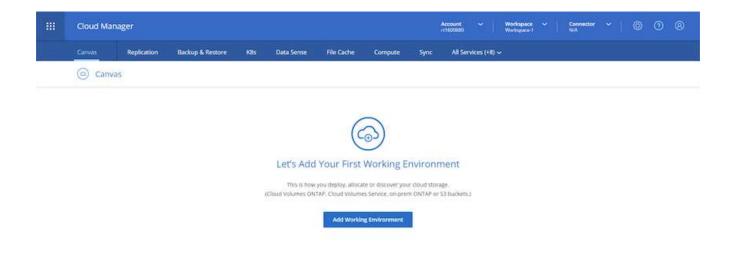
There are many methods for deploying Cloud Manager and Cloud Volumes ONTAP; this method is the simplest but requires the most permissions. If this method is not appropriate for your AWS environment, please consult the NetApp Cloud Documentation.

Deploy the Cloud Manager connector

1. Navigate to NetApp Cloud Central and log in or sign up.

Log In to NetApp Cloud Central	
Don't have an account yet? Sign Up	
rt1600680@demo.netapp.com	
LOGIN	

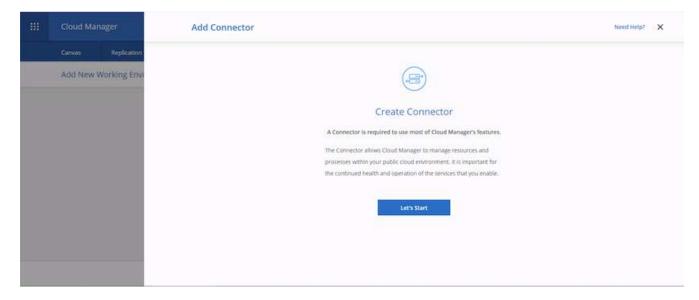
2. After you log in, you should be taken to the Canvas.



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.

***	Cloud Ma	nager							Account ~		kspace ~	0 1	onector X	@ 0	0 0
	Canvas	Replication	Backup & Restore	K8s	Data Sense	File Cache	Compute	Sync	All Services	(+8) ~					
	Add New	Working Enviro	nment												×
						aws	0		-						
				Wicrosh An	Ana Ana	zan Web Services	Google Cloud	Nation	On Prem	ives					
			Cho	iose Type											
				G	0)	6	a	1	\bigcirc						
				-	mes ONTAP		HES ONTAP HA	ci	oud Volumes Ser	vice					
					Node		ontaking -		High Availability						
						1		<u></u>							
							Next/								

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.

-	Cloud Manager	Add Connector				Need Help?	×
	Canvas Replication						
	Add New Working Envi						
				Provider			
			Choose the	cloud provider where you want to run th	e Connector:		
				0			
				aws	0		
			Microsoft Azure	Amazon Web Services	Google Cloud Platform		
				Continue			

6. Enter your secret key and access key. Make sure that your user has the correct permissions outlined on the NetApp policies page.

	Cloud Manager	Add Connector	Need Help? X
	Canvas Replication	🕗 Get Ready 👩 AWS Credentials 🗿 Details 🕣 Network 🗿 Security Group 🚳 Review	
_	Add New Working Envi	AWS Credentials	
		AWS Access Key AWS Access Key Is required AWS Secret Key Region Region Us-east-1 US East (M. Virginia)	
		Previous Next	

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.

 Cloud Manager	Add Connector	Need Help7 🗙
Canvan Replication	⊘ Get Ready ⊘ AWS Credentials 👩 Details 🗿 Network 🕥 Security Group 🕥 Review	
Add New Working Envi	Details	
	Connector Instance Name O Connector Role O	
	awscloudmanager Create Role Create Role Select an existing Role	
	Role Name	
	Add Tags to Connector Instance Cloud-Manager-Operator-IBNt24)	
	Previous Next	

- 8. Give the networking information needed to deploy the connector. Verify that outbound internet access is enabled by:
 - a. Giving the connector a public IP address
 - b. Giving the connector a proxy to work through
 - c. Giving the connector a route to the public internet through an Internet Gateway

-	Cloud Manager	Add Connector	Need Help? X
	Canves Replication	🧭 Get Ready 🔗 AWS Credentials 🔗 Details 💽 Network 🔇 Security Group 🚯 Review	
	Add New Working Env	Connectivity Proxy Configuration (Optional) VPC HTTP Proxy vpc 083fcd/19975dfb6e-10.221.0.0/16 Example: impuff 22.16.264.1 mml Subnet Define Credentials for this Proxy. ~ 10.221.4.0/24 publicSN_us-exist-1a_rt1600 Upload a root certificate ~ Key Pair Image:	
		Previous Next	

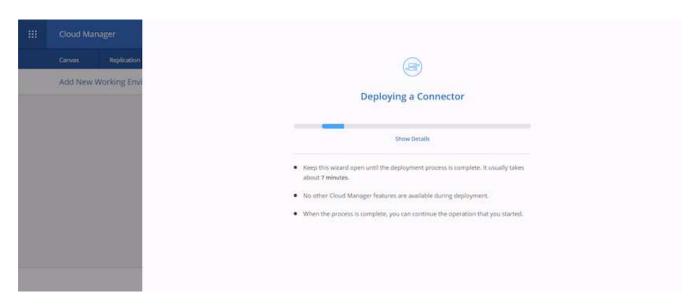
9. Provide communication with the connector via SSH, HTTP, and HTTPs by either providing a security group or creating a new security group. I have enabled access to the connector from my IP address only.

 Cloud Manager	Add Connector	Need Help? X
Canvas (Replication)	🧭 Get Ready ⊘ AWS Credentials 🕝 Details ⊘ Network 👩 Security Group 👩 Review	
Add New Working Envi	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 440) SSH (Port 22)	
	Source Type Source Type Source Type	
	My IP v My IP v	
	Source (CDR) Source (CDR) Source (CDR)	
	216.240.31.345/32 216.240.31.345/32 216.240.31.345/32	
	Previous Next	

10. Review the information on the summary page and click Add to deploy the connector.

-	Cloud Manager	Add Connector	Need Help? X
	Canvas. Replication	🧭 Get Ready 📿 AWS Credentials 🕝 Details ⊘ Network ⊘ Security Group 👩 Review	
	Add New Working Envi	Code for Terraform Automation	
		Connector Name awscloudmanager	
		Region us-east-1	
		VPC vpc-083fcbd79f75dfb6e - 10.221.0,0/16	
		Subnet 10.221.4.0/24 publicSN_us-east-1a_rt1600680	
		Key Pair rt1600680	
		Public IP Enable	
		Proxy None	
		Security Group HTTP: 216;240;31;145/32; HTTPS: 216;240;31;145/32; S5H: 216;240;31;145/32	
		Previous Add	

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.

Cloud Manager	x
Canvas Replication	Connector Successfully Created The Connector win created successfully.

Deploy Cloud Volumes ONTAP

1. Select AWS and the type of deployment based on your requirements.

 Cloud Ma	nager							ccount 1500580		Workspace Workspace-1	Connector evolvormena.	@ (0 0
Canvas	Replication	Backup & Restore	K8s	Data Sense	Ele Cache	Compute	Sync	All Servi	ces (+8) 🗸				
Add New	Working Enviro	onment											×
					aws	0			3				
			Microsft An	And And	uan Web Services	Google Clouif	Sachum	Do P	remites				
		c	ioose Type										
			6	6		6		\bigcirc					
				mes ONTAP	and the second second	HES ONTAP HA	Clos	ad Volumes	Service				
				Node		reliability		High Availab	7.				
					<u></u>	0							
						vest :							

2. If no subscription has been assigned and you wish to purchase with PAYGO, choose Edit Credentials.

	Cloud Ma	nager					\$	Account Verkapace Verkapace Verkapace 1	Connector avschadmana.	0	0
	Canvas	Replication	Backup & Restore	K8s Data Sense	File Cache	Compute	Sync	All Services (+8) ~			
	Create a No	ew Working Environ	oment		Details an	d Credential	S				
	Previous	Step .	Instance Profile Credential Nam			Ne outpurpettee in tarketplace Subscrip		Edit Credentials			
			Defails Working Enviro	nment Name (Cluster Nam	ej	Creder User N					
			Up to 40 char			admi Passwo	n				
			Add Tags	Optional field Up	to floor tags		nd n Password				
					6	social					
Coud M	langer3.9.9 Built	0 Aug 10, 2021 04,133	5 eri lutti								

3. Choose Add Subscription.

	Cloud Man	ager						Account	۰ (Workspace Workspace 1	•	Consector	1 4	0	8
	Canvas	Replication	Backup & Restore K8s	Data Sense	File Cache	Compute	Sync	All Ser	rvices (+8) ~						
	Create a Nev	w Working Enviro	ament		Details an	d Credentia									
	↑Prisiolis 3	Rep .	Instance Profile Credential Name	Edit Cre	edentials & Ac	ld Subscripti	an .		504	Determin					
			Details Working Environment A	Credentia	Subscription to			•							
			O Add Tags	9 m	place Subscription subscription		nderdaar								
000		Aug 16, 227, AAAA	a and a set		Apply	Ca	scel								

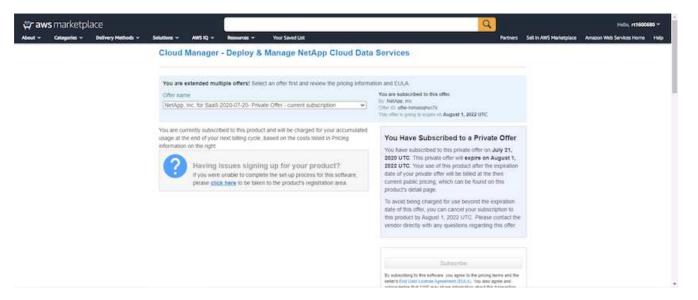
4. Choose the type of contract that you wish to subscribe to. I chose Pay-as-you-go.

 Cloud Ma	nager							Account	* •	forkspace forkspace (Connector availantment	© (0 0
Canvas	Replication	Backup & Restore	Kās	Data Sense	File Cache Cebentibilis for A	Compute	Sync	All Service	s (+8) 🛩					
Greate a He	w Working Enviror	strant:		Select a s details an Pay Pay	ubscription option nd then subscribe. y-Per-TiB - Annual y for Cloud Volum th an annual, upfri	and click Continue Contract es ONTAP	The AWS M	arketplace enal lay-as-you-go lay for Cloud V m hourly rate.	lalumes O					
 lange 115 Juli	1 - Seg 16, 207, 4411.2	Lan dik		(2) CO	t steps: IS Marketslace bscribe and then cliu wit Manager re your subscription		Marketplace		nth your Al	NS credentij	iks.			

5. You are redirected to AWS; choose Continue to Subscribe.

💝 aws marketplace	ver maansen waar				<u>a</u>	Hallo, r1600660 + etplace Amazon Web Services Home Heb
About + Gategoriet + Delivery Habo	■ NetApp	Answares Voia Savel Lint Cloud Manager - Deplo Data Services lot by: NetApp, Inc. tart here to deploy and manage Cloud loud Backup and Cloud Volumes Servi Shoe more	Volumes ONTAP, Cloud To	ning, Cloud Data Sense,	Parnes Sell in ANS Hart Continue to Soldisolbe Save to ligt	etplana Annazon Web Services Horne Help
	Product Over	Pricing	0srgę	Support	Reviews	
	NetApp Cloud Manager is th and operating NetApp's Clu - Cloud Volumes ONTAP - FI - Cloud Backup - Increment archiving CVD and On-Prem - Cloud Testing - Treing Infr - Cloud Testing - Treing Infr - Cloud Testing - Treing Infr - Cloud Manager also manager Cloud Manager eases the da environment inclusing code	e management and automation platfo al Data Services including: le and block storage fur enterprise wor d block-level Backup & Restore capabili	kloads tries for protecting and or AFF ar cloud storage act of your active		environments NetApp based storage and lity zones or to and from trators to audit and track	

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.

T NetApp		Falatic View	Fill -	22
111 Eable's view	Subscription Assignment ×			
NC: Feature Spotlight	Your subscription to Eloud Manager / Cloud Volumes ONTAP from AWS Marketplace was created successfully!			
10 Products (Name your subsiciption 0 demo.netapp.com-cloud-volumes-ontap-38695345			
×.	NetApp Cloud Central Account			
ET an	We've assigned your subsiciption to all of your NetApp Cloud Central accounts. You can choose to unassign specific accounts.			
terrises Status 13	Silve .			
Comart Us				

8. When successful, a check mark page appears. Navigate back to your Cloud Manager tab.

# NetApp		Fabric View	Full -	8
EE Fairs: View				
NGE: Former Spartlight	•			
Statistica i				
X tum :				
ATT AN				
Services Status (2	Your subscription demo.netapp.com-cloud-volumes-ontap-386953e5 saved successfully			
Contract Us				

9. The subscription now appears in Cloud Central. Click Apply to continue.

Create a New Working Environment	Edit Credentials & Add Subscription		
	Associate Subscription to Credentials		
	You subscribed successfully?	_	
	Credentals		
	Instance Profile Account ID: 322944748816		
	Subscription		
	 demo.netapp.com-cloud-volumes-ontap-386953e5 	•	
	Add Subscription		
	77 11		
	Apply Cancel	h i	
Charl Hengel 348 (Sell 8: Aug 18:201) 341000 are UTC			

- 10. Enter the working environment details such as:
 - a. Cluster name
 - b. Cluster password
 - c. AWS tags (Optional)

 Cloud Ma	nager						Account ~ Workspace ~ Con r11000880 * Workspace 1 avo	metor - 🗐 🧿 🤇	8
Canvas	Replication	Backup & Restore	K8s Data Sense	File Cache	Compute	Sync	All Services (+8) ~		
Create a N	ew Working Enviro	nment		Details an	d Credentia	Is			
↑ Previou	s Step	Instance Profile Credential Nam			demo.netapp.com- Marketplace Subscri		Edit Credentials		
		Details			Crede				
		Working Enviro hybridawscvo	nment Name (Cluster Nam	r)	User Madm				
					Passw	ord			
		🕒 Add Tags	Optional Held Up	to finur tags					
						m Password			
					ontinue	***			
anana 163 Kulot	5 Aug 18, 2021 04,130	A weak little							

11. Choose which additional services you would like to deploy. To discover more about these services, visit the NetApp Cloud Homepage.

ш	Cloud Ma	nager							Account r1500880	~ *	forkspac forkspace	•	Connector averalizationaria.	۲	0 (0
	Canvas	Replication	Backup & Restore	K8s D	ata Sense	File Cache	Compute	Sync	All Service	es (+8) ~						
	Create a N	ew Working Enviror	ament			Ser	vices									
	t Previou	s Step														
			💿 Dat	a Sense & Corr	ipliance						•	~				
			(i) Ван	kup to Cloud							•	~				
			(iii) Mo	nitoring							•	~				
						1 E										
		8 Aug 18, 2021 (A.13.2)				Co	ntinue									

12. Choose whether to deploy in multiple availability zones (reguires three subnets, each in a different AZ), or a single availability zone. I chose multiple AZs.

-	Cloud Ma	nager						Account m1600580		Workspace Workspace 1		Connector ~	@ @	0
	Canvas	Replication	Backup & Restore	K8s Da	ta Sense - File	Cache Con	pute Sync	All Servi	(ces (+8) ~	i.				
	Create a Ne	ew Working Enviro	nment.		HA D	eployment	Models							
	† Previous	Step	OP Enables set	aximum protectio	n against AZ failure lility zones. partner goes offlini	5.	Single availa group, spre	inst failures v bility zone, Hi id across disti	A nodes ar inct under	gle AZ e in a placem ying hardwan r goes offline.	e.			

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.

Cloud Manager							Account ~	Workspace	1 - 1	Connector	¥ 1	@ @	0 0	
Canvas	Replication	Backup & Restore	K8s	Data Sense	File Cache	Compute	Sync	All Services (*	R) 🛩					
Create a Ne	w Working Environ	iment			Region	n & VPC								
1 Previous	Step	AWS Region			VPC			Security group						
		US East N. Virginia)	*	vpc-083fcbd79f75d 10.221.0.0/16	fb6e -	×	Use a generated s	scurity group	*				
		Node 1:			Node 2:			Media	on					
		Availability Zone			Availability Zone			Availability Zone						
		us-east-1a			us-east-1b		•	us-east-1c		•				
		Subnet			Subnet			Subnet						
		10.221.1.0/24		*	10.221.2.0/24		*	10.221.3.0/24		+				
					Cor	stinue								
	Create a Ne		Create a New Working Environment Previous Step AWS Region US East 1 N. Virgen Figure Node 1: Avuilability Zone Us-east-1a Subnet	Creato a New Working Environment Previous Step AWS Region US East N. Virginia Node 1: Node 1: Vauitability Zone US-east-1a Subnet	Creato a New Working Environment Previous Step WWS Region US Edst 1 N. Virgenia Node 1: Availability Zone Ist-east-1a Subnet	Creato a New Working Environment Previous Step VPC VPC VPC VPC VPC VPC VPC VPC VPC VP	Create a New Working Environment Region & VPC Previous Step WVS Region US East 1 N. Virginia VVC VPC VPC VPC VPC VPC VPC VPC VPC VP	Create a New Working Environment Region & VPC Previous Step WYS Region US East 1 N. Virginia VPC Vpc Vpc Vpc Vpc Vpc Vpc Vpc Vpc Vpc Vpc	Create a New Working Environment Region & VPC Security group 1 Previous Step AWS Region VPC Security group US East 1 N. Virginia VPC Security group Vis East 1 N. Virginia VPC VPC Submet VPC VPC VPC Submet VPC VPC VPC Vis East 1 N. Virginia VPC VPC VPC Vis East 1 N. Virginia VPC VPC VPC Vis East 1 N. Virginia VPC VPC <td>Create a New Working Environment Region & VPC Security group 1 Previous Step AWS Region VPC Security group US East 1 N. Virginia VPC Security group US East 1 N. Virginia VPC Security group Vector VPC Security group Us East 1 N. Virginia VPC Security group Vector VPC VPC Vector VPC Security group Vector VPC VPC Vector <</td> <td>Create a New Working Environment Region & VPC Security group * Previous Step AWS Region VPC Security group US East 1 N. Virgenia * VPC Security group VPC VPC Security group Use a generated security group VPC VPC VPC Security group VPC VPC VPC Security group VPC VPC VPC VPC VPC VPC VPC Security group VPC VPC VPC VPC VPC VPC VPC Security group VPC VPC VPC Security group VPC VPC VPC Security group VPC VPC VPC VPC VPC</td> <td>Create a New Working Environment Region & VPC Security group * Previous Step AWS Region VPC Security group US East N. Virginia * VPC Security group VPC vpc0053ftbdd79f75dfb6e- * Us a generated security group VPC vpc0053ftbdd79f75dfb6e- * Us a generated security group VPC vpc014billity Zone * * Subnet * * * Subnet * * * 10.2211.0/24 * * *</td> <td>Create a New Working Environment Region & VPC Security group * Previous Step AWS Region VPC Security group US East N. Virgria * VPC Security group VPC VPC Security group * VPC VPC VPC VPC * VPC VPC VPC VPC * * VPC VPC VPC VPC * * * VPC VPC VPC VPC * * * VPC VPC VPC * * * * * * * * * *</td> <td>Create a New Working Environment Region & VPC Security group * Previous Step AWS Region VPC Security group US East N. Virgria * VPC Security group VPC vpc/dd3ftcbd79ff5dtb6e* * Use a generated security group VPC vpc/dd3ftcbd79ff5dtb6e* * Use a generated security group VPC vpc/dd3ftcbd79ff5dtb6e* * * VPC vpc/dd10e* * * * VPC vpc/dd10e* * * * VPC vpc/dd10e* * * * VPC v</td>	Create a New Working Environment Region & VPC Security group 1 Previous Step AWS Region VPC Security group US East 1 N. Virginia VPC Security group US East 1 N. Virginia VPC Security group Vector VPC Security group Us East 1 N. Virginia VPC Security group Vector VPC VPC Vector VPC Security group Vector VPC VPC Vector <	Create a New Working Environment Region & VPC Security group * Previous Step AWS Region VPC Security group US East 1 N. Virgenia * VPC Security group VPC VPC Security group Use a generated security group VPC VPC VPC Security group VPC VPC VPC Security group VPC VPC VPC VPC VPC VPC VPC Security group VPC VPC VPC VPC VPC VPC VPC Security group VPC VPC VPC Security group VPC VPC VPC Security group VPC VPC VPC VPC VPC	Create a New Working Environment Region & VPC Security group * Previous Step AWS Region VPC Security group US East N. Virginia * VPC Security group VPC vpc0053ftbdd79f75dfb6e- * Us a generated security group VPC vpc0053ftbdd79f75dfb6e- * Us a generated security group VPC vpc014billity Zone * * Subnet * * * Subnet * * * 10.2211.0/24 * * *	Create a New Working Environment Region & VPC Security group * Previous Step AWS Region VPC Security group US East N. Virgria * VPC Security group VPC VPC Security group * VPC VPC VPC VPC * VPC VPC VPC VPC * * VPC VPC VPC VPC * * * VPC VPC VPC VPC * * * VPC VPC VPC * * * * * * * * * *	Create a New Working Environment Region & VPC Security group * Previous Step AWS Region VPC Security group US East N. Virgria * VPC Security group VPC vpc/dd3ftcbd79ff5dtb6e* * Use a generated security group VPC vpc/dd3ftcbd79ff5dtb6e* * Use a generated security group VPC vpc/dd3ftcbd79ff5dtb6e* * * VPC vpc/dd10e* * * * VPC vpc/dd10e* * * * VPC vpc/dd10e* * * * VPC v

14. Choose the connection methods for the nodes as well as the mediator.

	Cloud Ma	inager							Account		Workspace Workspace	; ~ ;	Connector	 ۲	0	8
	Canvas	Replication	Backup & Restore	K8s	Data Sense	File Cache	Compute	Sync	All Ser	vices (+B)	~					
	Create a N	ew Working Enviro	nment		Conr	ectivity & S	SH Authen	tication								
	t Previou	s Step	Note Note	les.			<u>.</u>	Mediator								
			SSH Auther	tication Meth	od		Security Gr	oup								
			Password	s.,		•	Use a ger	verated secur	uth &uanti		•					
							Key Pair Na	me								
							/1160068	0)(
							internet Co	nnection Me	thod							
							Public IP	address			•					
						Cor	tinue									
CoutM	lanager 3.5.9 Suice	8 Aug 10, 2021 (AU) (3	S en UTC													



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.

 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 Documenation.

 Cloud Mar	ager			Account rt1618549		Workspace Workspace-1	Connector 🗸	Ĺ,	0						
Canvas	Replication	Backup & Restore	K8s	Data Sense	File Cache	Compute	Sync	All Services (+8) 🗸							
Create a Nev	w Working Enviror	iment				Floa	ting IPs								
↑ Previous !	Step		Floating IP	HA nodes if failu	Irres occur. To acce fy IP addresses th Floating IP add 10.222.0.200 Floating IP add 10.222.0.201 Floating IP add 10.222.0.202	ess the data from at are outside of ress for cluster m ress 1 for NFS and ress 2 for NFS and ress 2 for SVM mar	outside the Vi the CIDR block nanagement d CIFS data d CIFS data	CIFS data access. These ft PC, you can set up an AW ks for all VPCs in the selec	S transit gatew	ay.	tween				
						Cc	ontinue								

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 Ma	nager						Account ~	Workspace ~	Connector ~	® (0 0
Canvas	Replication	Backup & Restore	×85	Data Sense	Ble Ci	iche Compute	Sync All Services (+	8) ~			
Create a N	ew Working Environ	iment			R	loute Tables					
1 Previou	s Step						bles client access to the Cloud e route table cannot access th				
					AB	Stonal information 🗐					
		[Name			Main	iD	Associate with Subnet	t Tags			
		[2] private	_rt_rt1606680		No	rtb-08b4cb88/65c826a5	3 Subnets	1 Tags			
		2 public	rt_rt1600680		Ves	rtb-0e46720d0da10c59	1 Subnets	1 Tags			
		2 Route Table) The main is	ute table is the defa	ult for the	VPC					
					1	Continue					
						Continue					

3. Choose whether to enable AWS managed encryption or AWS KMS to encrypt the ONTAP root, boot, and data disks.

-	Cloud Ma	nager		Account m600680		Workspace Workspace 1	Connector ~		0						
	Canvas	Replication	Backup & Restore	K8s	Data Sense	File Cache	Compute	Sync	All Ser	vices (+8) •	×				
	Create a N	ew Working Enviro	oment			Data E	ncryption								
	1 Previous	s Step		ist	3 AWS Managed IS is responsible for andled by AWS key fault Master Key: a	data encryption a management sen		erabons. Key	nianageme	nt					
		5 Aug 18 2011 08 115				Co	ntinue								

4. Choose your licensing model. If you don't know which to choose, contact your NetApp representative.

-	Cloud Ma	nager							Account +1600680		Workspace Workspace-1	Connector	na.~	٢	0	
	Canvas	Replication	Backup & Restore	K8s	Data Sense	Ble Cache	Compute	Sync	All Sen	vices (+8) •						
	Create a N	ew Working Enviro	nment	Clou	d Volumes (ONTAP Cha	rging Meth	ods & N	ISS Acco	ount						
	1 Previous	i Step	Cloud Volumes ONTAP Learn more about our che © Pay-As-You-	Go by the h	ods.		Lean To rr shou Dorr finist	n more abou gister this C ild add Netaj t have a Net h deploying t	loud Volume pp Support S App Support this system.A	pport Site (is ONTAP to file Accourt Site accourt site accourt	NSS) accounts i support you ii 17Select go to					
			Freemium (Up to 500Gl	8)			Add Netapp	Support Site	Account						
						co	ntinue									
Churth	lauge 153 files	0 Aug 18, 2021 04:130	23 mm 10715													

5. Select which configuration best suits your use case. This is related to the sizing considerations covered in the prerequisites page.

***	Cloud Ma	nager							Account ~	Workspace Workspace 1		nector ~ doutmena	0	0	
	Canvas	Replication	Backup & Restore	Kās	Data Sense	Ele Cache	Compute	Sync	All Services (+8)	¢.					
	Create a Ne	ew Working Enviro	oment			Preconfigu	red Packag	95							
	† Previous	Step	Select a proc	configured	Cloud Volumes O Preconfi	NTAP system the guined settings ca	at best matches y in be modified at	your needs, a later timo	or create your own o	configuration.		Ch	ange Config	guration	
			mall workloads TB of storage	D	atabase ant applic production wor Up to 10TR of at	kloads		Cost effectiv			erformance p workloads to 36818 of too				
Contra	langer159 Suit	0 Aug 16, 2021 (64, 15.3	S en UTC			Co	atinue								

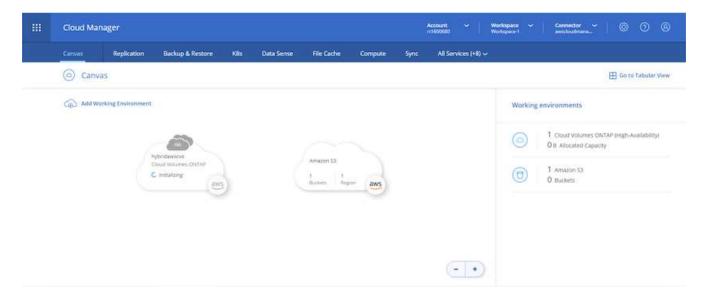
6. Optionally, create a volume. This is not required, because the next steps use SnapMirror, which creates the volumes for us.

**	Cloud Ma	nager							Account	• [Workspace Workspace 1		Connector ~ avectingtrung.	0	0	0
	Canvas	Replication	Backup & Restore	KBs	Data Sense	File Cache	Compute	Sync	All Servi	ces (+8) 🛩						
	Create a N	ew Working Environ	iment			Create	Volume									
	1 Previou	i Step	Details & Pro	tection			Protoco	d.								
			Volume Name:		Sue	:G81 @	145	5	CIFS		(5C5)					
			1		V	dume size	Access Con	anal .								
			Snapshot Policy:					export policy	<i>8</i>			2				
			default													
			ID Default Policy				Custom ex	port policy				0				
							10.221.0	0/16								
							Advanced	options				~				
						< Continue:	Sk	ip	Ĩ							
Charles M.	mana 160 Robert	9 Aug 18, 2027 64 53 35	and letter													

7. 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.

		er							4000000 11600680		Workspace 7		Connector avacloutmen		۲	0	
-0	anvas -	Replication	Backup & Restore	Kās	Data Sense	File Cache	Compute	Sync	All Service	rs (+8) 🛩	1						
	Create a New W	orking Environn	nent			Review (& Approve										
,	Previous Step	hybridawsc	VO ast-1 HA									5	iow API reque	st			
			d that in order to activate	support. I m	est first vegister (Doud Volumes ONT	AP with NetApp. Me	ne informatio	n F								
		V Lunderstan	d that Cloud Manager wil	allocate the	appropriate AW5	resources to comp	ly with my above re	quicements	More informat	ien it							
		Overview	Networking	5	torage									-			
		Storage System:	Cloud v	vilumes ONTA	на		HA Deploym	innt Model:	Mu	tiple Aval	ability Zones						
		License Type:	Cloud V	Diumes ONTA	f Standard		Encryption		AW	5 Mariage	d()						
		Capacity Limit	1078				Customer M	anter Key		urebic							

8. 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.



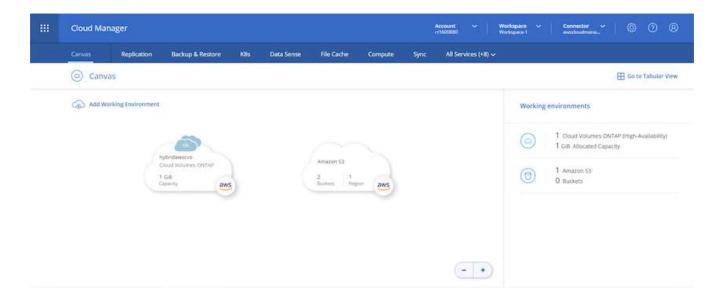
9. You can monitor the progress by navigating to the Timeline.

	Cloud M	anager						ccount ~	Workspace ~ Workspace-1	Connector ~	• •	8
	Canvas	Replication	Backup & Restore	Kās	Data Sense	File Cache Compute	Sync:	All Services (+)	n~			
		Resources Canvas Review	CHO, CVS, ANF & On Premises	*	6	Digital Wallet View & Manage Digital Wallet	#	*	Timeline Vew Activity & Servits			
		Services.	tion plication	*	6	Backup & Restore Data Protection for CVD and On-Preme	1	۲	KEs Cloud Native Development			
		Data Se Data Go	ense vernance & Compilante	*	0	Compliance Privacy & Compliance Controls	H.		Tiering Lift and DON'T shift	<i>\$</i>		
		(h) Monito	ring Optimize and Secure	2	(1)	File Cache Consolidate your Data into the Cloud	1	\bigcirc	Compute Optimize your clinic spend			
Margin /// South	ranagan metapip com	Sync Automa	ted Data Synchronization	×		SnapCenter Application Data Management	1	*	Active IQ Digital Advisor	1		

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.

-	Cloud Man	Cloud Manager				Account ~ Workspace ~ Commetter ~ 🚱 🕥 (
	Canvas	Replication Backup & Resto	re K8s Data Sense File	Cache Compute	s Sync	All Services (+8) ~			
	(Timel	ne							
		〒 Filters:	Action Agent (1) Resource	e Uier	Status Rea	et			
								Q ±	
		Time .	Action :	Service :	O Agent :	• Resource :	User :	Status :	
		> Aug 182021, 94232 pm	Check Connectivity	Coud Manager	avocloudman	hybridawscvo	Full Name	Success	
		Aug 18 2021, 9:42:00 pm	Create Aws Ha Working Environment	Cloud Manager	ampcloudma	hybridawucvo	full Name) Pending 🖷	
		Ave 18 2021 (1000-30 pm	Describe Operation Status					• Success	
			and a spectrum balance						
		1.11.12.2021, 10:00.20 mm	Details An Owned the state of the					A.Contan H.	

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.

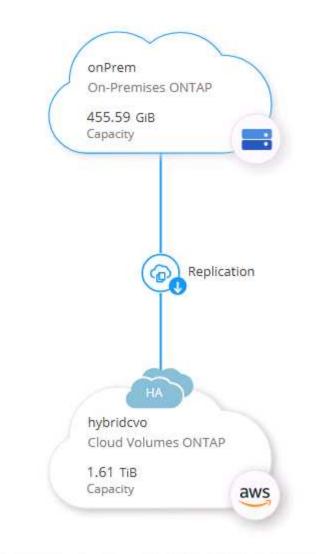


Configure SnapMirror from on-premises to cloud

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 > Menu > Replicate.



Select Enable.

SERVIC	ES		
6	Replication	Enable	

Or Options.

onPrem • On		×
DETAILS		
On-Premises ONTAP		
SERVICES		
Replication	1 Replication Target	(

Replicate.

onPrem • On DETAILS		(1) (1	
DETAILS			
On-Premise	s ONTAP		
SERVICES			
Re Re	e plication On	1 Replication Target	(
Re Re Ba		1 Replication Target	(

2. If you did not drag and drop, choose the destination cluster to replicate to.

Replicate Data	
From: onPrem	
To: select the Working Environment to which you wan	t to replicate data
Replication Target	
hybridcvo (Cloud Volumes ONTAP)	
Start Replication Wizard	Cancel

3. Choose the volume that you'd like to replicate. We replicated the data and all log volumes.

Replication S	etup				Source Volu	ime Selection					
rhel2_u03			ONLINE	e rhel2_u0309	923211942120311	18	ONLINE	sql1_data			ONLINE
NFO Storage VM Name Fiering Policy /olume Type	svm_onPrem None RW	CAPACITY 100 GB Allocated	7.29 GB Disk Used	INFO Storage VM Name Tiering Policy Volume Type	svm_onPrem None RW	CAPACITY 100 GB Allocated	35.83 MB Disk Used	INFO Storage VM Name Tiering Policy Volume Type	svm_onPrem None RW	CAPACITY 53.37 GB Allocated	■ 45.09 GB Disk Used
sql1_log			ONLINE	sql1_snapct	r	l concern	ONLINE				
NFO Storage VM Name Fiering Policy /olume Type	svm_onPrem None RW	CAPACITY 21.35 GB Allocated	18.16 GB Disk Used	INFO Storage VM Name Tiering Policy Volume Type	svm_onPrem None RW	24.87 GB Allocated	21.23 GB Disk Used				

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.

Replication Setup	Previous Step Destination Disk Type Image: Control of the previous state of the previous			
↑ Previous Step		General Purpose SSD - Dynamic		
	Enabled Disabled	ioning must be enabled on volumes created i		
		Continue		

- Cloud Manager 3.9.10 Build: 2 Sep 12, 2021 06:47:41 am UTC
- 5. Select the destination volume name: we chose [source_volume_name]_dr.

Destination Volume Name

Destination \	/olume Name	
---------------	-------------	--

sql1_data_dr

Destination Aggregate

Automatically select the best aggregate

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.

	Sec. 18	and the second se
 Limited to: 	100	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.

	Replicati	on Policy
	Default Policies	Additional Policies
Mirror Typically used for disaster recovery		Mirror and Backup (1 month retention) Configures disaster recovery and long-term retention of backups on the same destination volume
More info		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.

One-time copy	10min	12-hourly	5min	6-hourly
No schedule	Every hour Minutes: 0th, 10th, 20th, 3	Every day Hours: 12 AM and 12 PM Minutes: 15th minute	Every hour Minutes: 0th, 5th, 10th, 15t	Every day Hours: 12 AM, 6 AM, 12 PM Minutes: 15th minute
	8hour	daily	hourly	monthly
	Every day Hours: 2 AM, 10 AM and 6 Minutes: 15th minute	Every day Hours: 12 AM Minutes: 10th minute	Every hour Minutes: 5th minute	 Every month Days: 2nd Hours: 12 AM Minutes: 20th minute
		ī		
	pg-15-minutely	pg-6-hourly	pg-daily	pg-daily-set2
	C Fueny hour	C Even day	C Even day	O Evenuday

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.

Replication Setup			Review & Ap	oprove			
↑ Previous Step	Source onPrem onSrem	Destination (im) hybridcvo I → sql1_data_copy	Review your selection and start I understand that Cloud Manu- More information > Source Volume Allocated Size: Source Volume Used Size: Source Thin Provisioning: Destination Volume Allocated Siz Destination Volume Disk Type: Capacity Tiering:	ager will allocate the appropri 53.37 G8 45.09 G8 Yes	ate AWS resources to comply with m Destination Thin Provisioning: Destination Aggregate: Destination Storage VM: Max Transfer Rate: SnapMirror Policy: Replication Schedule:	y above requirements. Yes aggr1 (Automatically s svm_hybridcvo 100 MB/s Mirror daily	
			Go				

- 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.

Volume R	elationships	153.32 GIB Replicated Capacity	0 Currently Transl	erring	7 Healthy	⊗ 0 _{Failed}	
7 Volume Relationships							Q C
Health Status 💠	Source Volume +	Target Volume +	Total Transfer Time 💠	Status 🔤	Mirror State 🕴	Last Successful Transfer	•
Ø	rhel2_u01 onPrem	rhel2_u01_dr hybridcvo	43 minutes 43 seconds	idle	snapmirrored	Sep 30. 2021. 12:12:50 Af 19.73 MiB	
\odot	rhel2_u02 onPrem	rhel2_u02_dr hybridcvo	1 hour 37 minutes 59 seconds	idle	snapmirrored	Sep 30, 2021, 2:37:08 PM 239.78 MiB	-
\odot	rhel2_u03 onPrem	rhel2_u03_dr hybridcvo	16 hours 1 minute 9 seconds	idle	snapmirrored	Sep 30, 2021, 4:07:14 PM 225.37 KIB	
\odot	sql1_data onPrem	sql1_data_dr hybridcvo	1 hour 6 minutes 50 seconds	idle	snapmirrored	Sep 30, 2021, 12:12:28 Af 24.56 KiB	

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.

	Oracle Database	-					
Dashboard	View Database	Search datab	ases V			Ref	resh Resources New Resource
Resources	🌔 Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Overall Status
Monitor	cdb2	Single Instance (Multitenant)	rhel2.demo.netapp.com	rhel2_cdb2 rhel2_cdb2_log	Oracle Archive Log Backup Oracle Full Online Backup	09/17/2021 3:00:09 PM	Backup succeeded
Reports							
Reports							
Hosts							

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.

	racle Database 👻	cdb2	Topology								
	Search databases										rotect
17	LF 🕅 Name cdb2		age Copies	184 Backups 0 Clones Mirror copies					368 Backu 16 Da	ita Backups g Backups	
		D 1									
		(see	nary Backup(s) rich 꼬						Craing Res	tina Optie Ressore Moo	
		Sec			Count	Туре	l∓ End Date	Verified			
		Back	rch V		Count	Type Log	IF End Date 09/17/2021 3:00:10 PM ☐	Verified Not Applicable			SCN
		Backer	rrch 🛛 🐨		Count 1			Not	Citalog Reve Mounted	RMAN Cataloged	SCN 598200
		Back rhel2	rch 文 up Name _cdb2_log_09-17-2021_15.00.01.1317_1		Count 1 1	Log	09/17/2021 3:00:10 PM 🛱	Not Applicable Not	Mounted False	RMAN Cataloged	5982003
		Back rheiz rheiz	rch T		Count 1 1 1 1	Log	09/17/2021 3:00:10 PM 🛱	Not Applicable Not Applicable	Cristig Reca Mounted False False	RMAN Cataloged Not Cataloged Not Cataloged	SCN 5982003 5980623 5980584
		Back rheiz rheiz	rch p Name cdb2_log_09-17-2021_15.00.01.1317_1 .cdb2_09-17-2021_14.35.01.4997_1 .cdb2_09-17-2021_14.35.01.4997_0		Count 1 1 1 1 1 1	Log Log Data	09/17/2021 3:00:10 PM	Not Applicable Not Applicable Unverified Not	Consign Rena Mounted False False False	RMAN Cataloged Not Cataloged Not Cataloged Not Cataloged Not Cataloged Not Cataloged	SCN 598200 598062 598058 598058
		Back rheiz rheiz rheiz rheiz	rch rch V up Name		Count 1 1 1 1 1 1 1 1	Log Log Data Log	09/17/2021 3:00:10 PM	Not Applicable Not Applicable Unverified Not Applicable Not	Catalog Reca Mounted False False False False	Cons Press Noo RMAN Cataloged Not Cataloged Not Cataloged Not Cataloged Not Cataloged Not Cataloged	

3. Toggled to the mirrored backups view by clicking mirrored backups. The secondary mirror backup(s) is then displayed.

Oracle D	Database 👻	cdb2 Topology							
									7
Search	h databases							and the second second second second second second second second second second second second second second second	Protect
112 100	Name cdb2	Manage Copies 184 Backups 0 Clones Local copies Secondary Mirror Backup(s)					368 Backu 16 Da	ata Backups g Backups	
		search 7	Count	Ture	15 Fiel Date	Vestfied			
		Search Image: Search Badup Name rhel2_cdb2_log_09-17-2021_15.00.01.1317_1	Count 1	Type Log	I₹ End Date 09/17/2021 3:00:10 PM 🛱	Verified Not Applicable			SCN
		Backup Name	Count 1			Not	Mounted	RMAN Cataloged	SCN 5982003
		Backup Name mel2_cdb2_log_09-17-2021_15.00.01.1317_1	Count 1 1	Log	09/17/2021 3:00:10 PM 🛱	Not Applicable Not	Mounted False	RMAN Cataloged	SCN 5982003 5980629
		Backup Name mel2_cdb2_log_09-17-2021_15.00.01.1317_1 mel2_cdb2_09-17-2021_14.35.01.4997_1	Count 1 1 1 1 1	Log Log	09/17/2021 3:00:10 PM 🛱	Not Applicable Not Applicable	False	RMAN Cataloged Not Cataloged Not Cataloged	SCN 5982003 5980629 5980588
		Backup Name rhel2_cdb2_log_09-17-2021_15.00.01.1317_1 rhel2_cdb2_o9-17-2021_14.35.01.4997_1 rhel2_cdb2_c09-17-2021_14.35.01.4997_0	Count 1 1 1 1 1	Log Log Data	09/17/2021 3:00:10 PM	Not Applicable Not Applicable Unverified Not	Mounted False False False	RMAN Cataloged Not Cataloged Not Cataloged Not Cataloged Not Cataloged Not Cataloged	Othersure SCN 5982003 5980629 5980588 5980588
		Backup Name mel2_cdb2_log_09-17-2021_15.00.01.1317_1 mel2_cdb2_09-17-2021_14.35.01.4997_1 mel2_cdb2_09-17-2021_14.35.01.4997_0 mel2_cdb2_log_09-17-2021_14.00.01.1042_1	Count 1 1 1 1 1 1 1	Log Log Data Log	09/17/2021 3:00:10 PM 🖨 09/17/2021 2:35:21 PM 🖨 09/17/2021 2:35:12 PM 🖨 09/17/2021 2:30:10 PM 着	Not Applicable Not Applicable Unverified Not Applicable Not	Mounted False False False False	RMAN Cataloged Not Cataloged Not Cataloged Not Cataloged Not Cataloged Not Cataloged	

4. 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 backup	DS			×
Choose the host to mount the backup	ora-standby.demo.netapp.com]		
Mount path :	/var/opt/snapcenter/sco/backup_mount/rhel2_cdb2_09-1	7-2021_14.35.01.4997_1/cdb2		
Secondary storag	ge location : Snap Vault / Snap Mirror			
Source Volume		Destination Volume		
svm_onPrem:rhel	2_u03	svm_hybridcvo:rhel2_u03_dr 🔹		
			Mount	Cancel

-	tApp Snap	oCenter®					I ⊠ 0- 1	demo\oradba	App Backup and Clone A	dmin 🛛 🖡 Sign C
	Oracle Datab	ase 👻	cdb2 Topology							
	Search data	abases								vtoct Refr
	17 19	Name	Manage Copies							
•		cdb2	184 Backups							
	1	Cdb2dev							ary Card	
			Mirror copies					368 Backup		
			Local copies						ta Backups	
									g Backups	
								1 Clone		
			Secondary Mirror Backup(s)							
			Secondary Mirror Backup(s)					Cost	ing Cone Resource More	a Usmeane
				Count	Туре	l₹ End Date	Verified			
			(search)	Count	Type Log	I F End Date 09/17/2021 4:00:10 PM □	Verified Not Applicable	Cara	leg Cone Restore Mou	
			(search) Backup Name	Count 1			Not	Mounted	RMAN Cataloged	n Uhmount SCN
			Search V Backup Name rhel2_cdb2_Jog_09-17-3021_16.00.01.2156_1	Count 1	Log	09/17/2021 4:00:10 PM 🛱	Not Applicable Not	Mounted False	RMAN Cataloged	5985272
			Search V Backup Name	1	Log	09/17/2021 4:00:10 PM	Not Applicable Not Applicable Not	Mounted False False	RMAN Cataloged Not Cataloged Not Cataloged	SCN 5985272 5982003



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.

5. Highlight the full database backup copy to be cloned, and then click the clone button to start the DB clone Workflow.

cdb2 Topology							-
						and the second second second second second second second second second second second second second second second	otect Refre
search T					• It Catalog Renam	e Cone Restore Mour	a 🖬
Backup Name	Count	Туре	l∓ End Date	Verified	Mounted	RMAN Cataloged	SCN
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	Ť	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	True	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	Ť	Log	09/17/2021 2:00:10 PM 📋	Not Applicable	False	Not Cataloged	5978388

6. Choose a proper clone DB SID for a complete container database or CDB clone.

Clone from cdb	02				×
1 Name	Complete Database	Clone			
2 Locations	Clone SID	cdb2test			
3 Credentials	Exclude PDBs	Type to find PDBs			
4 PreOps	○ PDB Clone				
5 PostOps	Secondary storage locati	on : Snap Vault / Snap Mirror			
6 Notification	⊙ Data				
7 Summary	Source Volume		Destination Volume		
Junnary	svm_onPrem:rhel2_u02		svm_hybridcvo:rhel2_u02_dr	•	
	⊙ Logs				
	Source Volume		Destination Volume		
	svm_onPrem:rhel2_u03		svm_hybridcvo:rhel2_u03_dr	•	
				Previous	Next

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 cdb	02					×
1 Name	Select the host to o	reate a clone				
2 Locations	Clone host	ora-standby.demo.i	netapp.com	1	•	
3 Credentials	⊙ Datafile locations	0				
4 PreOps	/u02_cdb2test			Reset		
5 PostOps						
6 Notification	⊙ Control files ()					× * +
7 Summary		test/control/control01 test/control/control02				× + Reset
	Dede la ca					
	⊙ Redo logs () Group		Size	Unit	Number of files	
	RedoGroup 1	×	200	MB	1	+
	/u02_cdb2tes	/cdb2test/redolog/re	do03.log			X TReset
	RedoGroup 2	×	200	MB	1	+
						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 cdb	02		×
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/19800/cdb2	
6 Notification	Oracle OS User	oracle	
7 Summary	Oracle OS Group	oinstall	
		Previous Next	
		T EVIOUS NEAL	

9. Specify the scripts to run before clone operation. More importantly, the database instance parameter can be adjusted or defined here.

Clone from cdb	02					×
1 Name	Specify scripts to r	un before clone op	eration	0		
2 Locations	Prescript full path	/var/opt/snapcente	er/spl/scr	ipts/ Enter Prescript path		
3 Credentials	Arguments					
4 PreOps	Script timeout	60 secs				
	O Database Paramet	ter settings				
5 PostOps	processes			320	× •	
6 Notification	remote_login_passv	vordfile		EXCLUSIVE	× +	
	sga_target			4311744512	× Reset	
7 Summary	undo_tablespace			UNDOTBS1	×	
					Previous	xt

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.

NAME NAME Recover Database Outdations Until Cancel Octedentials Date and Time Date time format: MMX/DDVYYY hh:mm:ss Until SCN (System Change Number) Specify external archive log locations © Nommary Create new DBD Create new DBD Create tempfile for temppirary tablespace Create tempfile for temppirary tablespace Enter SQL queries to apply when clone is created Enter scripts to run after clone operation Create tempfile for temppirary tablespace Enter scripts to run after clone operation Create tempfile for temppirary tablespace Create tempfile for temporary tablespace Enter scripts to run after clone operation Create tempfile for temporary tablespace Create tempfil	Clone from cd	b2	×
Outril Cancel ③ Credentials ③ PreOps ⑤ PostOps ⑥ Notification ⑦ Notification ⑦ Summary ③ Create new DBID ① ④ Create new DBID ① ④ Create tempfile for temporary tablespace ① ③ Create tempfile for temporary tablespace ① ④ Create tempfile for temporary tablespace ① ④ Enter SQL queries to apply when clone is created ③ Enter scripts to run after clone operation ① Previous Next	1 Name	Recover Database	
Concernance () Date and Time	2 Locations	O Until Cancel	
PreOps Outli SCN (System Change Number) 5980629 Specify external archive log locations System Change Number) Specify external archive log locations System Change Number) Specify external archive log locations O Specify external archive log locations O Var/opt/snapcenter/sco/backup_mount/rhel2_cdb2_09-17-2021_14.35.01.4997_1/cdb2/1/orareco/CDB2/archivelog/ /var/opt/snapcenter/sco/backup_mount/rhel2_cdb2_09-17-2021_14.35.01.4997_1/cdb2/1/orareco/CDB2/archivelog/ /var/opt/snapcenter/sco/backup_mount/rhel2_cdb2_09-17-2021_14.35.01.4997_11	3 Credentials	🔿 Date and Time 💼 🕚	
S PostOps 6 Notification 7 Summary 9 Create new DBID 9 Create new DBID 9 Create new DBID 9 Create tempfile for temporary tablespace 9 Create tempfile for temporary tablespace 9 Enter SQL queries to apply when clone is created 9 Enter SqL queries to apply when clone is created 9 Enter SqL queries to apply when clone is created 9 Enter SqL queries to apply when clone is created 9 Enter Scripts to run after clone operation 9 Enter Scripts to run after clone operation 9 Enter SqL queries to apply when clone is created 9 Enter Scripts to run after clone operation 9 Enter SqL queries to apply of apply query (apply query q	A PreOps	Date-time format: MM/DD/YYYY hh:mm:ss	
6 Notification 7 Summary Ø Create new DBID Ø Create new DBID Ø Create tempfile for temporary tablespace Ø Enter SQL queries to apply when clone is created Ø Enter scripts to run after clone operation Ø excedera-standay/mp	U	Until SCN (System Change Number) 5980629	
• Construction of the first of	5 PostOps	Specify external archive log locations 💽 💿 🚯	
<pre> recte@restandby.tmp rectomestandby.tmp rectomestan</pre>	6 Notification	/var/opt/snapcenter/sco/backup_mount/rhel2_cdb2_09-17-2021_14.35.01.4997_1/cdb2/1/orareco/CDB2/archivelog/	
<pre> recte@restandby.tmp rectomestandby.tmp rectomestan</pre>	Cummani		
Create tempfile for temporary tablespace () Enter SQL queries to apply when clone is created Enter scripts to run after clone operation () Previous Next Previous Next	Journmary		
Create tempfile for temporary tablespace () Enter SQL queries to apply when clone is created Enter scripts to run after clone operation () Previous Next Previous Next			
Create tempfile for temporary tablespace () Enter SQL queries to apply when clone is created Enter scripts to run after clone operation () Previous Next Previous Next			
Create tempfile for temporary tablespace () Enter SQL queries to apply when clone is created Enter scripts to run after clone operation () Previous Next Previous Next			
Oracle@ora-standby/tmp1 is /var/opt/snapcenter/sco/backup_mount/rhe12_cdb2_09-17-2021_04.35.01.4997_1/cdb2/1/orareco/CDB2/archivelog/ 021_09_26_2021_09_30_2021_09_01_2021_09_01_2021_09_06_2021_09_07_2021_09_11_2021_09_13_2021_09_15_2021_09_17 2021_09_26_2021_09_31_2021_09_01_2021_09_06_2021_09_06_2021_09_10_2021_09_11_2021_09_15_2021_09_17 2021_09_26_2021_09_31_2021_09_01_2021_09_06_2021_09_06_2021_09_10_2021_09_11_2021_09_15_2021_09_17 2021_08_27_2021_08_31_2021_09_01_2021_09_06_2021_09_06_2021_09_10_2021_09_11_2021_09_15_2021_09_17 2021_08_27_2021_08_31_2021_09_01_2021_09_06_2021_09_06_2021_09_10_2021_09_11_2021_09_15_2021_09_17 2021_08_27_2021_08_31_2021_09_01_2021_09_06_2021_09_06_2021_09_10_2021_09_11_2021_09_15_2021_09_17 2021_08_27_2021_08_31_2021_09_01_2021_09_06_2021_09_07_2021_09_10_2021_09_11_2021_09_15_2021_09_17 2021_08_27_2021_08_31_2021_09_01_2021_09_06_2021_09_07_2021_09_10_2021_09_11_2021_09_15_2021_09_17 2021_08_27_2021_08_31_2021_09_01_2021_09_06_2021_09_07_2021_09_10_2021_09_11_2021_09_15_2021_09_17 2021_08_27_2021_08_31_2021_09_01_2021_09_06_2021_09_07_2021_09_10_2021_09_11_2021_09_15_2021_09_17 2021_08_27_2021_08_31_2021_09_01_2021_09_06_2021_09_07_2021_09_10_2021_09_11_2021_09_15_2021_09_17 2021_08_27_2021_08_31_2021_09_01_2021_09_06_2021_09_07_2021_09_10_2021_09_11_2021_09_15_2021_09_17 2021_08_27_2021_08_31_2021_09_01_2021_09_06_2021_09_07_2021_09_10_2021_09_11_2021_09_15_2021_09_17 2021_08_27_2021_08_31_2021_09_01_2021_09_06_2021_09_07_2021_09_11_2021_09_11_2021_09_15_2021_09_17 2021_08_27_2021_08_31_2021_09_10_2021_09_00_2021_09_10_2021_09_10_2021_09_11_2021_09_15_2021_09_17 2021_08_27_2021_08_31_2021_09_00_2021_09_00_2021_09_10_2021_09_11_2021_09_11_2021_09_11_2021_09_10_2021_09_10_2021_09_10_2021_09_10_2021_09_10_2021_09_10_2021_09_10_2021_09_10_2021_09_10_2021_09_10_2021_09_10_2021_09_10_2021_09_10_2021_09_10_2021_09_10_200_100_100_00_100_00_00_00_00_00_00_00_0			
Previous Next previous // Nex			
		① Enter scripts to run after clone operation ①	
		Provious	
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/ 021_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_17 021_08_27 2021_08_28 2021_08_31 2021_09_02 2021_09_04 2021_08_06 2021_09_08 2021_09_10 2021_09_12 2021_08_14		Previous Next	
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/ 021_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_17 021_08_27 2021_08_28 2021_08_31 2021_09_02 2021_09_04 2021_08_06 2021_09_08 2021_09_10 2021_09_12 2021_08_14			
	oracle@ora-standby tmm	ol\$ 1s /var/opt/snapcenter/sco/backup mount/rhe12 cdb2 09-17-2021 14.35.01.4997 1/cdb2/1/orareco/CDB2/archivelog/	×
	oraciegora-standby tmp	10	

11. Configure the SMTP server for email notification if desired.

Clone from cdb	02		×
1 Name	Provide email sett	ings 🕕	
2 Locations	Email preference	Never 👻	
3 Credentials	From	From email	
4 PreOps	То	Email to	
5 PostOps	Subject	Notification	
6 Notification			
7 Summary			
		r Clone jobs, an SMTP server must be configured. Continue to the Summary page to save your gs>Global Settings>Notification Server Settings to configure the SMTP server.	×
		Previous	Next

12. Clone summary.

Clone from cdb	52		×
1 Name	Summary		^
2 Locations	Clone from backup	rhel2_cdb2_09-17-2021_14.35.01.4997_0	
3 Credentials	Clone SID	cdb2test	
3 Credentials	Clone server	ora-standby.demo.netapp.com	
4 PreOps	Exclude PDBs	none	
•	Oracle home	/u01/app/oracle/product/19800/cdb2	
5 PostOps	Oracle OS user	oracle	
6 Notification	Oracle OS group	oinstall	
	Datafile mountpaths	/u02_cdb2test	
7 Summary	Control files	/u02_cdb2test/cdb2test/control/control01.ctl /u02_cdb2test/cdb2test/control/control02.ctl	1
	Redo groups	RedoGroup =1 TotalSize =200 Path =/u02_cdb2test/cdb2test/redolog/redo03.log RedoGroup =2 TotalSize =200 Path =/u02_cdb2test/cdb2test/redolog/redo02.log RedoGroup =3 TotalSize =200 Path =/u02_cdb2test/cdb2test/redolog/redo01.log	l
	Recovery scope	Until SCN 5980629	
	Prescript full path	none	
	Prescript arguments		
	Postscript full path	none	
	Postscript arguments		*
		Previous Finish	

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.

P oracle@ora-standby:/tmp				- 🗆 ×
[oracle@ora-standby tmp]\$ export ORACLE SI] [oracle@ora-standby tmp]\$ export ORACLE MO [oracle@ora-standby tmp]\$ export PATH=\$PAT] [oracle@ora-standby tmp]\$ sqlplus / as syst	ME=/u01/app/oracle/product H:\$ORACLE_HOME/bin	:/19800/cdb2		
SQL*Plus: Release 19.0.0.0.0 - Production (Version 19.3.0.0.0	on Fri Sep 17 17:49:29 202	21		
Copyright (c) 1982, 2019, Oracle. All right	nts reserved.			
Connected to: Oracle Database 19c Enterprise Edition Rel Version 19.3.0.0.0	ease 19.0.0.0.0 - Producti			
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			
3 CDB2_PDB1 1 4 CDB2_PDB2 1	READ ONLY NO READ WRITE NO READ WRITE NO READ WRITE NO			
SQL>				

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.

	Microsof	t SQL Server 👻					
Dashboard	View	Database - search	by name				Refresh Resources New Resou
Resources	1E M	Name	Instance	Host	Last Backup	Overall Status	Туре
Monitor		master	sql1	sql1.demo.netapp.com		Not available for backup	System database
2		model	sql1	sql1.demo.netapp.com		Not available for backup	System database
Reports		msdb	sql1	sql1.demo.netapp.com		Not available for backup	System database
Hosts		tempdb	sql1	sql1.demo.netapp.com		Not available for backup	System database
Storage Systems		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
E Settings		model	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database
Alerts		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.

Microsoft SQL Server 🚽	tpcc (sql1) Topology					
search by name					Cione Lifecycle Protect	i Details
Name	Manage Copies					
master	7 Backups				C	
model	7 Backups 0 Clones				Summary Card	
msdb	Local copies Mirror copies				14 Backups 0 Clones	
tempdb	Local copies				U ciones	
tpcc						
1	Drimony Packun(c)					
	Primary Backup(s)					
		Count	Туре	1F	End Date	
	(search 💟	Count 1	Type Full backup	1F	End Date 09/16/2021 6:25:05 PM 🛱	
	(search 🛛 🔍) Backup Name			IF		Verified
	Search Image: Control of the search Backup Name sql1_tprcc_06-16-2021_18.25.01.4024	1	Full backup	IF	09/16/2021 6:25:05 PM 🛱	Verified Unverified
	Search V Backup Name sql1_tpcc_06-16-2021_18.25.01.4024 sql1_tpcc_06-15-2021_18.25.01.4604	1	Full backup Full backup	17	09/16/2021 6:25:05 PM	Pename Cone Resolu- Verified Unverified Unverified
	Search V Backup Name sql1_tprc_06>16-2021_18.25.01.4024 sql1_tprc_06>15-2021_18.25.01.4604 sql1_tprc_09-14-2021_18.25.01.5238	1 1 1	Full backup Full backup Full backup	17	09/16/2021 6:25:05 PM 🛱 09/15/2021 6:25:06 PM 🛱 09/14/2021 6:25:05 PM 🛱	Verified Unverified Unverified Unverified
	Search V Backup Name sql1_tprc_09-16-2021_18.25.01.4024 sql1_tprc_09-15-2021_18.25.01.4024 sql1_tprc_09-15-2021_18.25.01.4024 sql1_tprc_09-14-2021_18.25.01.5233 sql1_tprc_09-13-2021_18.25.01.4550	1 1 1 1	Full backup Full backup Full backup Full backup	1F	09/16/2021 6:25:05 PM 🛱 09/15/2021 6:25:06 PM 🛱 09/14/2021 6:25:05 PM 🛱 09/13/2021 6:25:05 PM 🛱	Verified Unverified Unverified Unverified Unverified

 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.

RELA	pp SnapCenter®				• = •-	L demo\sqldba App Backup an	d Clone Admin 🛛 🖡 Sign
	Microsoft SQL Server 👻	tpcc (sql1) Topology					
	search by name					Cione Lifecycle Protect	i Details R
r	Name	Manage Copies					
1	master	7 Backups					
1	model	7 Backups 0 Clones				Summary Card	
	msdb	0 Clones Mirror copies				14 Backups	
f	tempdb	Local copies				0 Clones	
1	tpcc						
1							
		Secondary Mirror Backup(s)					
7			Count	Туре	17	End Date	
7		(search T)	Count 1	Type Full backup	1F	End Date 09/16/2021 6:25:05 PM 🛱	
T		search T			17		Core Renor Verified
-		Search T Backup Name sql1_spcc_09-16-2021_18.25.01.4024	1	Full backup	IF	09/16/2021 6:25:05 PM 🛱	Verified Unverified
		Search T Backup Name sql1_tpcc_09-16-2021_18.25.01.4024 sql1_tpcc_09-15-2021_18.25.01.4604	1	Full backup Full backup	1F	09/16/2021 6:25:05 PM 🛱 09/15/2021 6:25:06 PM 🛱	Cone Reson Verified Unverified Unverified
		Search V Backup Name sql1_spcc_09-16-2021_18.25:01.4024 sql1_spcc_09-15-2021_18.25:01.4664 sql1_spcc_09-11.2021_18.25:01.5233	1	Full backup Full backup Full backup	15.	09/16/2021 6:25:05 PM 🛱 09/15/2021 6:25:06 PM 🛱 09/14/2021 6:25:05 PM 🛱	Verified Unverified Unverified Unverified
		search V Backup Name sql1_spcc_0916-2021_18.25.01.4024 sql1_spcc_09515-2021_18.25.01.4024 sql1_spcc_09514-2021_18.25.01.5233 sql1_spcc_09514-2021_18.25.01.4500 sql1_spcc_09514-2021_18.25.01.4500	1 1 1 1	Full backup Full backup Full backup Full backup		09/16/2021 6:25:05 PM 09/15/2021 6:25:06 PM 09/14/2021 6:25:05 PM 09/14/2021 6:25:05 PM 09/13/2021 6:25:05 PM	Core Resort

4. Choose a backup copy, and then click the Clone button to launch the Clone from Backup workflow.

ΠN	atApp SnapCenter®						• ≅ 0··	👤 demo\sqldba 🛛 App E	Backup and	Clone Admin	🖡 Sign Out
	Microsoft SQL Server	tpcc (sql1) Topology									×
	search by name								Protect	1 Details	Refresh
	Name							Cione Lifecycle	Protect	Details	Refresh
•	master	Manage Copies	7 Backups								
	model	0 Clones	1 Clone					Summar	y Card		
	msdb	Local copies	Mirror copies					14 Backups 1 Clone			
4	tempdb tpcc										
b.	master										
華	model	Secondary Mirror Backup(s)									
▲	msdb	(search)								m	4
	tempdb									Clone	Restore
	tpcc_clone	Backup Name sql1_tpcc_09-19-2021_18.25.01.4134			Count	Type Full backup	17	09/19/2021 6:25:05	d Date	Verified Unverifier	4
		sql1_tpcc_09-18-2021_18.25.01.3963			1	Full backup		09/18/2021 6:25:05		Unverifie	
		sql1_tpcc_09-17-2021_18-25-01-4218			1	Full backup		09/17/2021 6:25:05		Unverifie	
		sql1_tpcc_09-16-2021_18.25.01.4024			1	Full backup		09/16/2021 6:25:05		Unverifier	
		sql1_tpcc_09-15-2021_18.25.01.4604			1	Full backup		09/15/2021 6:25:06		Unverifier	
		sql1_tpcc_09-14-2021_18.25.01.5233 sql1_tpcc_09-13-2021_18.25.01.4500			1	Full backup Full backup		09/14/2021 6:25:05		Unverifier Unverifier	
C	lone from back	up									×
(Clone Options	Clone settings									
(2 Logs	Clone server	Choose		-	9					
(3 Script	Clone instance	Nothing selected		-	•					
(4 Notification	Clone name	tpcc								
(5 Summary	Choose mount opt	ion								
		Auto assign mou	nt point 🚺								
		 Auto assign volu 	me mount point under path	full file path		0					
		Secondary storage	location : Snap Vault / Sr	nap Mirror							
		Source Volume		Destination V	/olume						
		svm_onPrem:sql1_c	ata	svm_hybrid	lcvo:sql1	1_data_dr		•			
		svm_onPrem:sql1_b	og	svm_hybrid	lcvo:sql′	1_log_dr		•			
								Prev		Ne	¢t

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 back	up				×
1 Clone Options	Clone settings				
2 Logs	Clone server	sql-standby.demo.netapp.	com •	• 1	
3 Script	Clone instance	sql-standby	-	• •	
4 Notification	Clone name	tpcc_clone			
5 Summary		unt point 1 Ime mount point under path		0	
	Source Volume	e location : Snap Vault / Sr	Destination Vol	ume	
	svm_onPrem:sql1_u	data	svm_hybridcv	ro:sql1_data_dr 🔹	
	svm_onPrem:sql1_l	og	svm_hybridcv	o:sql1_log_dr •	
				F	Previous Next

6. Determine a recovery point either by a log backup time or by a specific date and time.

Clone from back	up		×
1 Clone Options	Choose logs		
2 Logs	○ All log backups		
3 Script	• -) 8	7/2021 6:25:10 PM 🔹	
4 Notification	By specific date until None	9/17/2021 6:25:05 PM	
5 Summary	O None		
			Previous Next

7. Specify optional scripts to run before and after the cloning operation.

Clone from back	up		×
1 Clone Options	Specify optional so	cripts to run before and after performing a clone from backup job	
2 Logs	Prescript full path		
3 Script	Prescript arguments	Choose optional arguments	
4 Notification	Postscript full path		
	Postscript arguments	Choose optional arguments	
5 Summary	Script timeout	60 secs	
		Previous	ext

8. Configure an SMTP server if email notification is desired.

Clone from back	up		×
1 Clone Options	Provide email setti	ngs 🕦	
2 Logs	Email preference	Never •	
3 Script	From	From email	
	То	Email to	
4 Notification	Subject	Notification	
5 Summary	🗌 Attach Job Report		
		one jobs, an SMTP server must be configured. Continue to the Summary page to save your Global Settings>Notification Server Settings to configure the SMTP server.	
		Previous	Next

9. Clone Summary.

Clone from back	up		×
1 Clone Options	Summary		
2 Logs	Clone server	sql-standby.demo.netapp.com	
3 Script	Clone instance	sql-standby	
3 Schpt	Clone name	tpcc_dev	
4 Notification	Mount option	Auto assign volume mount point under custom path	
	Prescript full path	None	
5 Summary	Prescript arguments		
	Postscript full path	None	
	Postscript arguments		
	Send email	No	
		Previous	Finish

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.

NetApp Snap	Center®	0		•	a 6	→ L demo\sqldba	Арр Вас	kup and Clone Admin	🖡 Sign Out
	Jobs	Schedules	Events Logs					a 2	•
Dashboard	searc	h by name	Y						
Resources	Jobs -	Filter							
Monitor	ID	Status	Name	Start date		En	d date	Owner	
Reports	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	4	Discover resources for all hosts	09/16/2021 7:56:49 PM 🛱		09/16/2021 7:56:54	PM 🛱	demo\sqldba	
🐴 Hosts	761	4	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	
Storage Systems	760	A	Discover resources for all hosts	09/16/2021 7:19:05 PM 🛱		09/16/2021 7:19:09	PM 🛱	demo\sqldba	
E Settings	759	A	Discover resources for all hosts	09/16/2021 7:18:43 PM 🛱		09/16/2021 7:18:48	PM 🛱	demo\sqldba	
	756	A	Discover resources for all hosts	09/16/2021 6:59:51 PM 🛱		09/16/2021 6:59:56	PM 🛱	demo\sqldba	
Alerts	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 🛱	Demo\administrato	r .
	745	4	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.

B oracle@hel2-	×
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);	
l 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 Oz.12.13.00000 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.

■ NetApp Snap	Center®				٠	≅ 6	• 1 demo\oradba	App Backup and Clone Admin	🖡 Sign Out
	Oracle Database 👻								
Dashboard	View Resource Group	Search resource g	tuont A						New Resource Gro
Resources	Name	Resources	Tags	Policies			Last Bac	kup Overall Status	
	rhel2_cdb2	1	orafullbkup	Oracle Full Online Backup			09/17/2021 2:38:16 PM	Completed	
	rhel2_cdb2_log	1	oralogbkup	Oracle Archive Log Backup			09/17/2021 6:02:13 PM	Completed	
Reports									
🐴 Hosts									
Storage Systems									
🛱 Settings									
Alerts									

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.

n Ne	etApp SnapCenter®									8 -	👤 demo\oradba	App Backup and C	lone Admin	🖡 Sign Ou
>	Oracle Database 👻	rhel2_cdb	02_log Details											
	Search resource groups	search									Modify Resource	Sroup Back up Now	Maintenance	Delete
0	Name	Resource	Name	Туре	Host									
٩	rhel2_cdb2	cdb2		Oracle Database	rhel2.den	no.netapp.c	om							
2 2 2 2	rhel2_cdb2_log													
*														
34														
 A														
4														
	1													
_														
	Backup										×			
	Bachap													
	Create a backı	JD f	or the select	ted resource a	roup)								
	ereate a bacit	чр i	or the belief	curesource a	- o ap						_			
	Resource Group		rhel2_cdb2_l	0.7										
L '	resource droup		THEIZ_CODZ_I	og										
		[
	Policy		Oracle Archiv	e Log Backup		-	0							
1	oney	l		- ·										
											- 1			
					Г									
						C	ancel	Ba	acku	n				
										P				

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.

i

NetApp	SnapCenter®						٠	≅ 0 • .	👤 demo\oradba	App Backup and Clone Ad	min 🖡 Sign (Out
Oracle I	Database 👻		cdb2 Topology									×
	h databases									Database Settings	Rest Ref	fresh
0 .	Name	11	Manage Copies									
•	cdb2		185 Backups						Summ	nary Card		
2 1	cdb2de cdb2te		0 Clones Mirror copies						370 Backu			
*	cuozte	51	Local copies							ta Backups		
24									354 Lo 2 Clone	g Backups s		
=												-
 A			Secondary Mirror Backup(s)									_
44			search 7						Cat	log Cone Restore Mount	A Unmount	
			Backup Name	Count	Туре	47 E	nd Date V	/erified	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			pt/snapcenter/sco/backup_mount/rhel2_cdb2_log_09- tion : Snap Vault / Snap Mirror	17-202	1_18.2	0.04.1177_1/cc	ib2					
	e Volume	-		Destina	tion Vo	olume						
svm_0	onPrem:rh	nel2_u03	[svm_h	ybrido	vo:rhel2_u03_	dr		•			
									N	Aount C	ancel	

5. Select the last full database backup and click Clone to initiate the clone workflow.

	Oracle Database	-	cdb2 Topology							
	Search databa	ises								votect R
	(N	Name 11	Manage Copies 185 Backups							
1		cdb2	185 Backups 2 Clones					Summ	hary Card	
	1	cdb2dev	0 Clones Mirror copies					370 Backu	ips	
	Tin I	cdb2test	Local copies						ita Backups	
								354 Lo 2 Clone	g Backups	
								2 cione	<i>.</i>	
			Secondary Mirror Backup(s)							
			(search)					Cata		
			Backup Name	Count	Туре	↓₹ 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
									Not Cataloged	
			rhel2_cdb2_log_09-17-2021_15.00.01.1317_1	Ť	Log	09/17/2021 3:00:10 PM 🛱	Not Applicable	False	Not Cataloged	5982003
			rhel2_cdb2_log_09-17-2021_15:00:01.1317_1 rhel2_cdb2_09-17-2021_14:35:01.4997_1	1	Log Log	09/17/2021 3:00:10 PM		False	Not Cataloged	5982003 5980629

6. Select a unique clone DB ID on the host.

Clone from cdb	52				×
1 Name	Complete Database (Clone			
2 Locations	Clone SID	<u>cdb2dr</u>			
3 Credentials	Exclude PDBs	Type to find PDBs			
4 PreOps	O PDB Clone				
5 PostOps	Secondary storage location	on : Snap Vault / Snap Mirror			
6 Notification	⊙ Data				
7 Summary	Source Volume		Destination Volume		
	svm_onPrem:rhel2_u02		svm_hybridcvo:rhel2_u02_dr	•	
	\odot Logs				
	Source Volume		Destination Volume		
	svm_onPrem:rhel2_u03		svm_hybridcvo:rhel2_u03_dr	•	
				Previous	Next

7. Provision a log volume and mount it to the target DR server for the Oracle flash recovery area and online logs.

	System Mar	nager			Search actions, objects, and pages	s Q
DASHBOARD	Volum	nes				
STORAGE ^	+ Add	: More				
Overview		Name	Storage VM	Status	Capacity	
Applications Volumes	~	ora_standby_u01	svm_hybridcvo	🕑 Online	12.3 GB used 17.7 GB available	31.6 GB
LUNS	~	rhel2_u01_dr	svm_hybridcvo	o Add Vol	ume	×
Shares			1.1.1			
Qtrees Quotas	~	rhel2_u02_dr	svm_hybridcvo	⊘ o NAME ora_stanc	lby μ03	
Storage VMs	~	rhel2_u02_dr09172116081193 60	svm_hybridcvo	•		
Tiers		rhel2_u02_dr09172117035348	svm_hybridcvo			
NETWORK V		63	- /	20	GB 🗸	
EVENTS & JOBS 🗸 🗸	~	rhel2_u03_dr	svm_hybridcvo	⊘ ○	Ontions Conset	Save
PROTECTION V	~		svm_hybridcvo	Ø 0	Options Cancel	Save
ноѕтѕ 🗸		75				

P ec2-user@ora-standby:/tmp						3222	
ec2-user@ora-standby tmp]\$ sudo mkdir /u03_odb ec2-user@ora-standby tmp]\$ chown oracle:olmsta hown: changing ownership of /u03_odb2dr: Ope ec2-user@ora-standby tmp]\$ sudo chown oracle:o ec2-user@ora-standby tmp]\$ sudo chown oracle:	ll <mark>/u03 cd</mark> ration not install /u	permi 03_cdk	o2dr	1103	/u03_cdb2dr		
ec2-user@ora-standby tmp1\$ df -h		_					
filesystem	Size	Used	Avail	Use%	Mounted on		
levtmpfs	7.6G		7.6G	0%	/dev		
mpfs	7.6G		7.6G	0%	/dev/shm		
mpfs	7.6G	17M			/run		
mpfs	7.6G		7.6G	0%	/sys/fs/cgroup		
dev/nvme0n1p2	10G	9.0G		90%			
0.221.1.6:/ora_standby_u01		13G	18G	42%	/u01		
mpfs	1.6G		1.6G	0.%	/run/user/1000		
0.221.1.6:/Sc28182452-3fa8-448c-9e4a-c5a9e465f	353 100G	3.1G	97G	4%	/u02_cdb2dev		
mpfs	1.6G		1.6G	0.%	/run/user/54321		
0.221.1.6:/Sc39c05df8-4b00-4b3a-853c-9d6d338e5	df7 100G		97G	4%	/u02_cdb2test		
0.221.1.6:/Sccf886a5c-3273-475e-ad97-472b2a8dc	cee 100G	3.8G			/var/opt/snapcenter/sco/backup mount/rhe12 cdb2 log 09-17-2021 18.20.04.1177 1/c	db2/1	
0.221.1.6:/ora standby u03		320K	20G		/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 cd	lb2
1 Name	Select the host to create a clone
2 Locations	Clone host ora-standby.demo.netapp.com
3 Credentials	⊙ Datafile locations ①
(4) PreOps	/u02_cdb2dr Reset
5 PostOps	
6 Notification	○ Control files ● /u02_cdb2dr/cohtrol/control01.ctl ×
(7) Summary	/u03_cdb2dr/cdb2dr/control/control02.ctl
	⊙ Redo logs ()
	Group Size Unit Number of files
	▲ RedoGroup 1 × 200 MB 1 + +
	/u03_cdb2dr/cdb2dr/redolog/redo03.log
	RedoGroup 2 X 200 MB 1 +
	Previous Next

9. Select the credentials for the clone. Fill in the details of the Oracle home configuration on the target server.

Clone from cdb	o2			x
1 Name	Database Credentials for	the clone		
2 Locations	Credential name for sys user	None	+ 0	
3 Credentials	Database port	1521]	
4 PreOps	Oracle Home Settings)		
5 PostOps	Oracle Home	/u01/app/oracle/product/19800/cdb2		
6 Notification	Oracle OS User	oracle		
7 Summary	Oracle OS Group	oinstall		
			Previous	ĸt

10. Specify the scripts to run before cloning. Database parameters can be adjusted if needed.

Clone from cdb	02								×
1 Name	Specify scripts to r	un before d	clone o	peration	0				
2 Locations	Prescript full path	/var/opt/s	napcen	ter/spl/scr	ipts/ Enter Prescript path]
3 Credentials	Arguments								
4 PreOps	Script timeout	60	secs						
	🖸 Database Paramet	ter settings							
5 PostOps	audit_file_dest				/u01/app/oracle/admin/cdb2dr/adump	×	^		
6 Notification	audit_trail				DB	×		+	
	open_cursors				300	×		Reset	
7 Summary	pga_aggregate_targ	et			1432354816	×	•		
							Previo	us N	lext

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 cd	lb2	×
1 Name	Recover Database	
2 Locations	Until Cancel	
3 Credentials	🔿 Date and Time 🗂 🗂	
4 PreOps	Date-time format: MM/DD/YYYY hh:mm:ss	
	O Until SCN (System Change Number)	
5 PostOps	Specify external archive log locations 💿 💿 🕕	
6 Notification	/var/opt/snapcenter/sco/backup_mount/rhel2_cdb2_log_09-17-2021_18.20.04.1177_1/cdb2/1/orareco/CDB2/archivelog/	1
7 Summary		
	Create new DBID 1	
	Create tempfile for temporary tablespace 1	
	② Enter SQL queries to apply when clone is created	
	 Enter scripts to run after clone operation 	

12. Configure the SMTP server for email notification if needed.

Clone from cdb	02		×	
1 Name	Provide email settings 1			
2 Locations	Email preference	Never 👻		
3 Credentials	From	From email		
4 PreOps	То	Email to		
	Subject	Notification		
5 PostOps				
6 Notification				
7 Summary				
		r Clone jobs, an SMTP server must be configured. Continue to the Summary page to save your is>Global Settings>Notification Server Settings to configure the SMTP server.	*	
		Previous	Next	

13. DR clone summary.

Clone from cdb	52		×
1 Name	Summary		*
2 Locations	Clone from backup	rhel2_cdb2_09-17-2021_14.35.01.4997_0	
Conductively	Clone SID	cdb2dr	
3 Credentials	Clone server	ora-standby.demo.netapp.com	
4 PreOps	Exclude PDBs	none	
	Oracle home	/u01/app/oracle/product/19800/cdb2	
5 PostOps	Oracle OS user	oracle	
6 Notification	Oracle OS group	oinstall	
	Datafile mountpaths	/u02_cdb2dr	
7 Summary	Control files	/u02_cdb2dr/cdb2dr/control/control01.ctl /u03_cdb2dr/cdb2dr/control/control02.ctl	
	Redo groups	RedoGroup =1 TotalSize =200 Path =/u03_cdb2dr/cdb2dr/redolog/redo03.log RedoGroup =2 TotalSize =200 Path =/u03_cdb2dr/cdb2dr/redolog/redo02.log RedoGroup =3 TotalSize =200 Path =/u03_cdb2dr/cdb2dr/redolog/redo01.log	
	Recovery scope	Until Cancel	
	Prescript full path	none	
	Prescript arguments		
	Postscript full path	none	
	Postscript arguments		-
		Previous Finis	sh

14. Cloned DBs are registered with SnapCenter immediately after clone completion and are then available for backup protection.

			atabase 🗾						
	Dashboard	View	Database	Search databases	V				esh Resources New Resource Gro
9	Resources	IF IM	Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Overall Status
0	Monitor		cdb2	Single Instance (Multitenant)	rhel2.demo.netapp.com	rhel2_cdb2 rhel2_cdb2_log	Oracle Archive Log Backup Oracle Full Online Backup	09/17/2021 7:00:10 PM 🛱	Backup succeeded
M	Reports	-	cdb2dev	Single Instance (Multitenant)	ora-standby.demo.netapp.com				Not protected
•	Hosts		cdb2dr	Single Instance (Multitenant)	ora-standby.demo.netapp.com				Not protected
	Storage Systems		cdb2test	Single Instance (Multitenant)	ora-standby.demo.netapp.com				Not protected
2	Settings								
A	Alerts								

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.

P oracle@ora-standby//u01/app/oracle/product/19800/cdb2/dbs	-	Х
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_pdbl open;		
Fluggable database altered.		
SQL> alter session set container=cdb2_pdbl;		
Session altered.		
SQL> select * from pdbadmin.dr_test;		
ID		
EVENT		
 DT		
l testing DB clone for DR and roll forward DB to last available log 17-5EP-21 02.12.13.000000 PM		
SQL>		~

2. Configure the flash recovery area.



- 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	-	- 🗆	\times
C:\Users\administrator.DEMO>sqlcmd 1> select host_name() 2> go			î
SQL1			
(1 rows affected) 1> use tpcc 2> go Changed database context to 'tpcc'. 1> insert into snap_sync values ('test snap mirror D 2> go	R for SQL', getdate())		
(1 rows affected) 1> select * from snap_sync 2> go event d ²	ŧ		
test snap mirror DR for SQL 20	021-09-20 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.

II Ne	tApp SnapCenter®				٠		0 -	1 demo	sqldba App	Backup and C	lone Admin	🗊 Sign Out
>	Microsoft SQL Server 🚽	sql1_tpcc_logDetails										×
	search by name	search			Modify	Resource Gro		Lup Now	Cone Lifecycle	Maintenance	i Edit/View Details	Delete
0	Name	Resource Name	Туре	Host								
	sql1_tpcc	tpcc (sql1)	SQL Database	sql1.demo.netapp.com								
a il	sql1_tpcc_log											
• }•												

3. Manually run a log backup to flush the last transaction to be replicated to secondary storage in the public cloud.

Backup		×					
Create a backup for the selected resource group							
Resource Group	sql1_tpcc_log						
Policy	SQL Server Log Backup 🔹 🚺						
		_					
	Cancel Backup						

4. Select the last full SQL Server backup for the clone.

n Ne	tApp SnapCenter®				٠	-	🕽 🔹 💄 demo\sqldba	App Backup ar	nd Clone Admin	🖡 Sign Ou
>	Microsoft SQL Server 🚽	tpcc (sql1) Topology								
	search by name						Cone Life		i Details	Refresh
9	Name	Manage Copies								
2	master	7 Backups						C		
	model	7 Backups 2 Clones						nmary Card		
M	msdb	0 Clones Mirror copies						ckups		
	tempdb	Local copies					2 Cl	ones		
	tpcc									
	master									
2	model	Secondary Mirror Backup(s)								
	msdb								6	
	tempdb	(search)								Restore
	tpcc_clone 1	Backup Name	Count	Туре		F		End Date	Verified	
	tpcc_dev	sql1_tpcc_09-19-2021_18:25.01.4134	1	Full backup			09/19/2021 6	:25:05 PM 🛱	Unverifie	ed
		sql1_tpcc_09-18-2021_18.25.01.3963	1	Full backup			09/18/2021 6	:25:05 PM 🛱	Unverifie	ed
		sql1_tpcc_09-17-2021_18.25.01.4218	1	Full backup			09/17/2021 6	:25:05 PM 🛱	Unverifie	ed

5. 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.

Clone from back	up				x
1 Clone Options	Clone settings				
2 Logs	Clone server	sql-standby.demo.netapp.c	om 🔻	• •	
3 Script	Clone instance	sql-standby	•	6	
4 Notification	Clone name	tpcc_dr			
5 Summary	Choose mount opt	ion			
	Auto assign mou	int point 🕦			
	 Auto assign volu 	me mount point under path	full file path	•	
	Secondary storage	location : Snap Vault / Sn	ap Mirror		
	Source Volume		Destination Volu	ume	
	svm_onPrem:sql1_c	lata	svm_hybridcv	o:sql1_data_dr 🔹	
	svm_onPrem:sql1_l	og	svm_hybridcv	o:sql1_log_dr 🔹	
					Previous Next

6. Select all log backups to be applied.

Clone from back	up		>	ĸ
1 Clone Options	Choose logs			
2 Logs	All log backups			
3 Script 4 Notification 5 Summary	 By log backups until By specific date until None 	9/19/2021 6:25:10 PM 09/19/2021 6:25:05 PM		
			Previous Next	

7. Specify any optional scripts to run before or after cloning.

Clone from back	up		×
1 Clone Options	Specify optional so	ripts to run before and after performing a clone from backup job	
 2 Logs 3 Script 4 Notification 5 Summary 	Prescript full path Prescript arguments Postscript full path Postscript arguments Script timeout	Choose optional arguments Choose optional arguments 60 secs	
		Previous	ext

8. Specify an SMTP server if email notification is desired.

Clone from back	up		×
1 Clone Options	Provide email setti	ngs 🚯	
2 Logs	Email preference	Never •	
3 Script	From	From email	
	То	Email to	
4 Notification	Subject	Notification	
5 Summary	🗌 Attach Job Report		
		one jobs, an SMTP server must be configured. Continue to the Summary page to save your Global Settings>Notification Server Settings to configure the SMTP server.	×
		Previous	Next

9. DR clone summary. Cloned databases are immediately registered with SnapCenter and available for backup protection.

Clone from back	up		×
1 Clone Options	Summary		
2 Logs	Clone server	sql-standby.demo.netapp.com	
2 Carriet	Clone instance	sql-standby	
3 Script	Clone name	tpcc_dr	
4 Notification	Mount option	Auto Mount	
0	Prescript full path	None	
5 Summary	Prescript arguments		
	Postscript full path	None	
	Postscript arguments		
	Send email	No	
			During
			Previous Finish

	Microsoft	t SQL Server 👻					
Dashboard	View	Database • search b	y name V				Refresh Resources New Resource (
Resources	15.00	Name	Instance	Host	Last Backup	Overall Status	Туре
Monitor		master	sql1	sql1.demo.netapp.com		Not available for backup	System database
		model	sql1	sql1.demo.netapp.com		Not available for backup	System database
Reports		msdb	sql1	sql1.demo.netapp.com		Not available for backup	System database
Hosts		tempdb	sql1	sql1.demo.netapp.com		Not available for backup	System database
Storage Systems		tpcc	sql1	sql1.demo.netapp.com	09/22/2021 5:35:08 PM 🛱	Backup failed, Schedules on hold	User database
		master	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database
E Settings		model	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database
Alerts		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
	- E	tpcc_clone	sql-standby	sql-standby.demo.netapp.com		Not protected	User database
	a	tpcc_dev	sql-standby	sql-standby.demo.netapp.com		Not protected	User database
	- E	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 Snap	Center®			•	🔤 🕐 🗜 demo\sqldba App B	ackup and Clone Admin 🛛 🖡 Sigr
	Jobs	Schedules				
Dashboard	search	by name	<u>_</u> 7 9			
Resources	Jobs - F	ilter				
Monitor	ID	Status	Name	Start date	End date	Owner
Reports	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\sqldba
	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\sqldba
Hosts	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\sqldba
Storage Systems	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\sqldba
Ξ Settings	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\sqldba
≃ octuingo	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\sgldba

2. Validate that last transaction has been replicated and recovered with all log file clones and recovery.

Administrator: Command Prompt - sqlcmd - SQLCMD		- 0	×
C:\Users\administrator.DEMO>sqlcmd 1> select host_name() 2> go			^
SQL-STANDBY			
(1 rows affected) 3> ges tpcc_dr 2> gg Changed database context to 'tpcc_dr'. 1> select * from snap_sync 2> go wvent			
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 inquires.

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-premisses 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 }}'

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
```

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-premisses to AWS cloud.
- A database solution architect who is interested in Oracle database migration from on-premisses 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-premisses 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.localdomain4
::1 localhost localhost.localdomain localhost6
localhost6.localdomain6
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 prerequsites.

```
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
```

 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

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" {
                                   = var.ami
    ami
                                   = var.instance_type
    instance_type
    subnet id
                                   = var.subnet id
    key name
                                   = var.ssh key name
    root_block_device {
                                   = "qp3"
      volume type
      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_szie.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

Open HTML parser AWR mod	lule.
← C A (≜ https://app.atroposs.com/#/awr-module	四 田 谷 公 GI印
Aatroposs	AWR Assessment The solution is built using cutting-edge technologies WebAssembly(WASM), Javascript and Python. Click to learn more. As ATROPOSS runs in a modern browser there is no software or additional libraries installation needed. Since ATROPOSS runs on the local browser, the hosting platform won't have access to the user data. How does it work?
Meet your expert!	SQL Files
MODULES	Feedback
AWR Module	If you have any improvements for the module, feel free to contact us or much better please contribute to enhance the module. Thanks in advance for your interest and we hope to hear from you.
Database Assessment	SQL analysis source ® Top 20 SQL (first 30 chars) O Complete List of SQL (full sql)
App Analyse Module	Clustering of DB workload DB SKU templates: Standard, D2_v5 (2 vCPUs, 8 GiB)
SQL Module	Number of DB clusters: 4
Powerapp Module	Output file format ⊠ csv □.json □.xisx
Netapp ANF Module	Upload AWR files
DEV-MODULES	Upload files
Development Module	Note: Runtime data will be removed once the user closes either the tab or browser.
GENERAL	
Imprint	

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.

\leftarrow C \bigcirc https://app.atr	roposs.com/#/awr-module A* 🔂 🗘 🗇 🍃 🌜
A	Charts and Stuff
atroposs	dioSize (1) sql Couldn't download - Network issue
Meet your expert!	Contraction Removed dbSize.sql Copen file
MODULES	See more
AWR Module	Overfilting Fitting Read Write Mixed AAS > CPUs AAS < CPUs
Database Assessment	Overfitting Threshold Read/Write Threshold Average Active Sessions (AAS)
App Analyse Module	30 • 60% • x = user calls per transaction x = read to throughput % Pricing Overview Reset Data
SQL Module	Application overfitting: x > 30 Read workload: x > 60% Azure PostgreSQL Charts and Table Write workload: x < 40%
Powerapp Module	Mixed workload: 40% < x < 60%
Netapp ANF Module	Table Summary Azure VM Comparison
DEV-MODULES	Table Summary
Development Module	Reset Table Format
GENERAL	ID Filename Type Status DB InstID InstancReleas(DB CPUs Used Memor/Used Used Vetwor/Overfit/DB IO AAS ConsoliConso name (gb CPUs (mb) Memor/Overfit/DB IO Bandwi Read Cluster/SKU (gb CPUs (mb) Throug(mbps) (%)
Imprint	A A A A A A A A A A A A A A A A A A A
Frontend and Webassembly by: PRODYNA (20)	1 awrrpt_1_27_28.html SI PASSELCDB1 1 cdb1 19.0.0. 8 7 31496 16384 54797. 643.48 13.19 5.35 97.12 77.07 1 cdb1 2.55 97.12 77.07 8 6 68

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.

e Home Insert Drav	v Page Layout Formulas	Data Review View	Automate Help											Commer	nts 🖻 S
et From Web	Existing Connections All	💌 🗟 Workbook Links	Organization Stock		graphy v Sort	tw.	Text to Columns 😹 Data V	/alidation ~	📳 Mana	olidate ionships age Data Model		Forecast Gr Sheet	oup Ungroup		
Get & Transform \checkmark : $\times \checkmark f_{x}$		Queries & Connections		Data Types		Sort & Filter		Data Too	bls		Fore	cast	Outlin	e G	5
A A		с	D	E	F	G	н	E	J	К	L N	I N	0	Р	Q
	-														
			DB data size	DB log size				_		_	_	_			
QL script output:			DD data size	DB tog size											
cure VM Sizing:	Instance type	VCPU	CPU usage (%)	RAM (MB)											
Ū		0.0	0	0											
icrosoft ANF Sizing:	IO throughput (MB/s)	IOPS 0	DB data volume (GiB) 0	DB logs volume (GiB) 0	IO profile (read/total) 0.00%										
orage for snapshots and fr pace	eeli		Allow 50% free space	Allow 50% free space		-									

4. Highlight column A and fields 1 and 2, click on Data, then Text to Columns to open the Text Wizzard. Choose Delimited, then Next to next screen.

Home Insert Dr	aw Page Layout Formulas	Data Review View	Automate Help										P	omments
t From Web		Gueries & Connections E Properties Workbook Links Queries & Connections	- 1	icks Currencies	Geography 😨	Ā	Filter Resppty Filter Advanced rt & Filter	Text to		En Consolidate Relationships	odel King King King King King King King King	Sheet ~	Ungroup Subto	+19 19
✓ : × ✓ f	d=filename=parent=db_typ	e=status=db_snap_begin_t	ime=db_snap_end_time	db_edition=db_rel	ease=db_cdb=dl	b_rac=db_inst_n		me*db_uname	=db_inst_name=	host_name=platfor	m=host_cpu_num=e	ib_cpu_count=c	lb_cpu_num≖d	b_cpu_usage_p
A	В	С	D	E		F	G	н	1		L M	N	0 P	Q
	e=status=db_snap_begin_time= =none=SI=PASSED=24-Apr-24 16													
L script output:			DB data size	DB log size										
				Convert Text to	Columns Wizard -	Step 1 of 3		7 X						
ire VM Sizing:	Instance type	VCPU	CPU usage (%)			it your data is Delim								
2		0 0	0	0 If this is correct, Original data ty		lose the data type t	nat best describes your da	ita.	_					
crosoft ANF Sizing:	IO throughput (MB/s)	IOPS	DB data volume (GiB) C Choose the file	type that best des									
te: allocate additional	0	0	0				abs separate each field. ith spaces between each f	Gald						
rage for snapshots and	ING		Allow 50% free spac	Preview of sele										
				1 id=filena 2 l=avrrpt_ 3 4 5	me-parent-db_t 1_1100_1101.ht	ype=status=db_ nl=none=SI=PAS	snap_begin_time=db_ SED=24-Apr-24 16:12	snap_en * :11=24-: 						
						Cancel	ack Next >	Einish						
				1										

5. Check Other, then enter '=' as Delimiters. Click on Next to next screen.

AutoSave (On 🌒 😭	ワ・ペッジ oracle_db_sizing	_template_and.xist + Saved >		🔎 Séarch Cao, Alten 📓 — 🗇
ile Home Insert Dra	aw Page Layout Formulas	Data Review View	Automate Help	(무 Comments) 영
Get From Text/CSV Get From Web Data ~ I From Table/Range Get & Transfo		Queries & Connections Properties Workbook Links Queries & Connections	Organization Stoc	
\vee 1 X \vee t	x id=filename=parent=db_tvpe	=status=db snap begin t	ime=db_snap_end_time=	db_edition=db_release=db_cdb=db_rac=db_inst_num=db_inst_id=db_name=db_uname=db_inst_name=host_name=platform=host_cpu_num=db_cpu_count=db_cpu_num=db_cpu_usage_pc
A id=filename=parent=db_typ	8 e=status=db_snap_begin_time=d	C lb_snap_end_time=db_ed	D ition=db_release=db_cdb	E F M N O P G 1 K M N O P G 1 control inst. JATr-MTAP1ong, Bit-Linux, 46 64-bit-4rina, -4-663 = 1558-2049-4069-258 -4406-768 44268 76-4442, 622-2145 64-124, 02-115, 07-15, 07-15, 01-121, 2119, 99-11 1 control inst. JATr-MTAP1ong, Bit-Linux, 46 64-bit-4rina, -4-663 = 1558-2049-4069-258 44066 76-6442, 682-4248
SQL script output:			DB data size	DB tog size
SQL SCRIPT OUTPUT:				Convert Text to Columns Wizard - Step 2 of 3 ? X
Azure VM Sizing:	Instance type	vCPU 0 0	CPU usage (%) 0	This scene lets you set the definitiency your data contains. You can see how your test is affected in the provise below. Definiters Definiters
Microsoft ANF Sizing:	IO throughput (MB/s) 0	IOPS 0	DB data volume (GIB) 0	[] Jab [] Semicolon] Treat consecutive delimiters as one
Note: allocate additional storage for snapshots and : space	free		Allow 50% free space	Some to gualifier · ·
				Data geniew
				5.4 filename parent fn_type status fn_tange bedin A kwrzpe_l_1180_1181.hemi nome BI PASED N-Aper-24 16:1
				Cancel < Sack Net > Enion

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.

e Home Insert Dr	aw Page Lavout Form	nulas Data Review View	Automate Help									Comments	ts 🖻 St
From Text/CSV From Web ta From Table/Range Get & Transfe	From Picture >	Refrech All ~ Dy Workbook Links Queries & Connections		ccks Currencies Geograph Data Types	ny ⊽ Ž↓ ŽŽ Sort	Filter		move Duplicates	S= Consolidate Relationships Manage Data Mo	del Analysis - She Forecast	ast Group Ungn tt -		
~) i [× ~ J	c id=filename=parent=c	db_type=status=db_snap_begin_ti	ime=db_snap_end_tim	e=db_edition=db_release=db_	cdb*db_rac*db_inst	_num=db_inst_id=db_na	ame*db_uname*c	db_inst_name*	host_name=platforn	n=host_cpu_num=db_	pu_count≖db_cpi	num*db_cpu_us	isage_pct*
A	В	c	D	E	F	G	н	1.1	J K	L M	N O	P C	Q
		_time=db_snap_end_time=db_edi r-24 16:12:11=24-Apr-24 16:27:54=1											
				Convert Text to Columns V	fizard - Step 3 of 3		? ×						
zure VM Sizing:	Instance type	VCPU	CPU usage (%)	F This screen lets you select ea	ach column and set the [ata Format.							
care viri sizing.		0 0	0	0 Column data format									
icrosoft ANF Sizing:	IO throughput (MB/s) 0) IOPS 0	DB data volume (Gil			rts numeric values to numb Il remaining values to text.	ers, date values						-
ote: allocate additional		-		O Do not import column		Advanced							
orage for snapshots and bace	free		Allow 50% free spac	e A Destination: SAS1			1						
Jace			Allow John nee space	C P Strugton 1931									
				Data preview									
				PenerGeneral id filename 1 awrrpt_1_1180_		eneral Ceneral Ceneral b_type status db_snap I PASSED 24-Apr	p_begin≜						
					Cancel	< Back Next >	Einish						

e Home Insert Dra		Review View											C.	Comments	
Get From Text/CSV	Recent Sources	Queries & Connections Properties Workbook Links es & Connections	Organization Stock	s Currencies Geo	graphy ਦ X↓ XXX Graphy E X↓ Sort Fi	Resppty	Text to Columns 😸 Data	n Fill Ber C ove Duplicates Her R Validation ~ M Data Tools	onsolidate elationships lanage Data Mo		at-If Forecas ysis ~ Sheet Forecast	t Group	Ungroup Sut	btotal	
> : × ✓ 𝔅															
	В	С	D	E	F	G	н	1 1	к	L	м	N	0	P Q	
d	filename 1 awrrpt_1_1180_1181.html	parent none	db_type SI	status PASSED		db_snap_end_time 4/24/2024 16:2		db_releas db_cdb 19.0.0.0 YES	db_rac NO					inst_n:host_n P1 ora_01	
iQL script output:			DB data size	DB log size											
uzure VM Sizing:	Instance type SI	vCPU 4	CPU usage (%) 68.9	RAM (MB) 15628											
licrosoft ANF Sizing:	IO throughput (MB/s) 239.09	10PS 22749.64	DB data volume (GiB)		IO profile (read/total) 62.89%		-				-				
lote: allocate additional torage for snapshots and f		12740104													
			Allow 50% free space												
torage for snapshots and f															
torage for snapshots and f															
torage for snapshots and f															
torage for snapshots and f															
torage for snapshots and f															

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
     6:
            where first time >= sysdate - 14
new
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.

Get Data ~	From Web	Recent Sources	Refresh All × Workbook Links	organization Stock	cs Currencies Geo	ography 🔤 🗛		Text to Columns 😽 Dat	a Validation 👻 🔞	Consolidate Relationships Manage Data Model	What-If Analysis ~	Sheet	oup Ungroup	
	Get & Transf		Queries & Connections		Data Types	S	ort & Filter		Data Tools		Forec	ist	Outline	R.
C11														
1 id	A	6 filename 1 awrrpt_1_1180_1181	parent .html none	D db_type SI	E status PASSED	F db_snap_begin_time 4/24/2024 16	G db_snap_end_time :12 4/24/2024 16:2		db_releas db_cd 19.0.0.0 YES	K db_rac db NO	L M _inst_nidb_ins 1			P C b_inst_n-host_ TAP1 ora_0
3 4 5 SQL	script output:			DB data size 159.05	DB log size 93.83									
6 7 Azur	e VM Sizing:	Instance type	VCPU	CPU usage (%)	RAM (MB)									
9	•	SI IO throughput (MB/s	4 s) IOPS	68.9 DB data volume (GiB)	15828 DB logs volume (GiB)	IO profile (read/total)								
11 Note	osoft ANF Sizing: : allocate additional ige for snapshots and ie	239.09	22749.64	318.1 Allow 50% free space	187.66 Allow 50% free space	62.89%								
13														
14														
16														
17														
18 19														
20														
21														
22 23														
23														
25														
26														
27 28														
29														
30														
31 32														
33														
34														
35														
<	> ora_azure_	anf_sizing +						: •						

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 requirment 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.