



TR-4988: Oracle Database Backup, Recovery, and Clone on ANF with SnapCenter

NetApp database solutions

NetApp
July 31, 2025

Table of Contents

- TR-4988: Oracle Database Backup, Recovery, and Clone on ANF with SnapCenter 1
 - Purpose 1
 - Audience 1
 - Solution test and validation environment 1
 - Architecture 1
 - Hardware and software components 2
 - Oracle database configuration in the lab environment. 3
 - Key factors for deployment consideration 3
 - Solution deployment. 3
 - Prerequisites for deployment 3
 - SnapCenter installation and setup 5
 - Database backup 11
 - Database recovery 19
 - Database clone 27
 - Where to find additional information. 40

TR-4988: Oracle Database Backup, Recovery, and Clone on ANF with SnapCenter

Allen Cao, Niyaz Mohamed, NetApp

This solution provides overview and details for automated Oracle deployment in Microsoft Azure NetApp Files as primary database storage with NFS protocol and Oracle database is deployed as container database with dNFS enabled. Database deployed in Azure is protected using SnapCenter UI tool for simplified database management.

Purpose

NetApp SnapCenter software is an easy-to-use enterprise platform to securely coordinate and manage data protection across applications, databases, and file systems. It simplifies backup, restore, and clone lifecycle management by offloading these tasks to application owners without sacrificing the ability to oversee and regulate activity on the storage systems. By leveraging storage-based data management, it enables increased performance and availability, as well as reduced testing and development times.

In TR-4987, [Simplified, Automated Oracle Deployment on Azure NetApp Files with NFS](#), we demonstrate automated Oracle deployment on Azure NetApp Files (ANF) in Azure cloud. In this documentation, we showcase Oracle database protection and management on ANF in Azure cloud with a very user-friendly SnapCenter UI tool.

This solution addresses the following use cases:

- Backup and recovery of Oracle database deployed on ANF in Azure cloud with SnapCenter.
- Manage database snapshots and clone copies to accelerate application development and improve data lifecycle management.

Audience

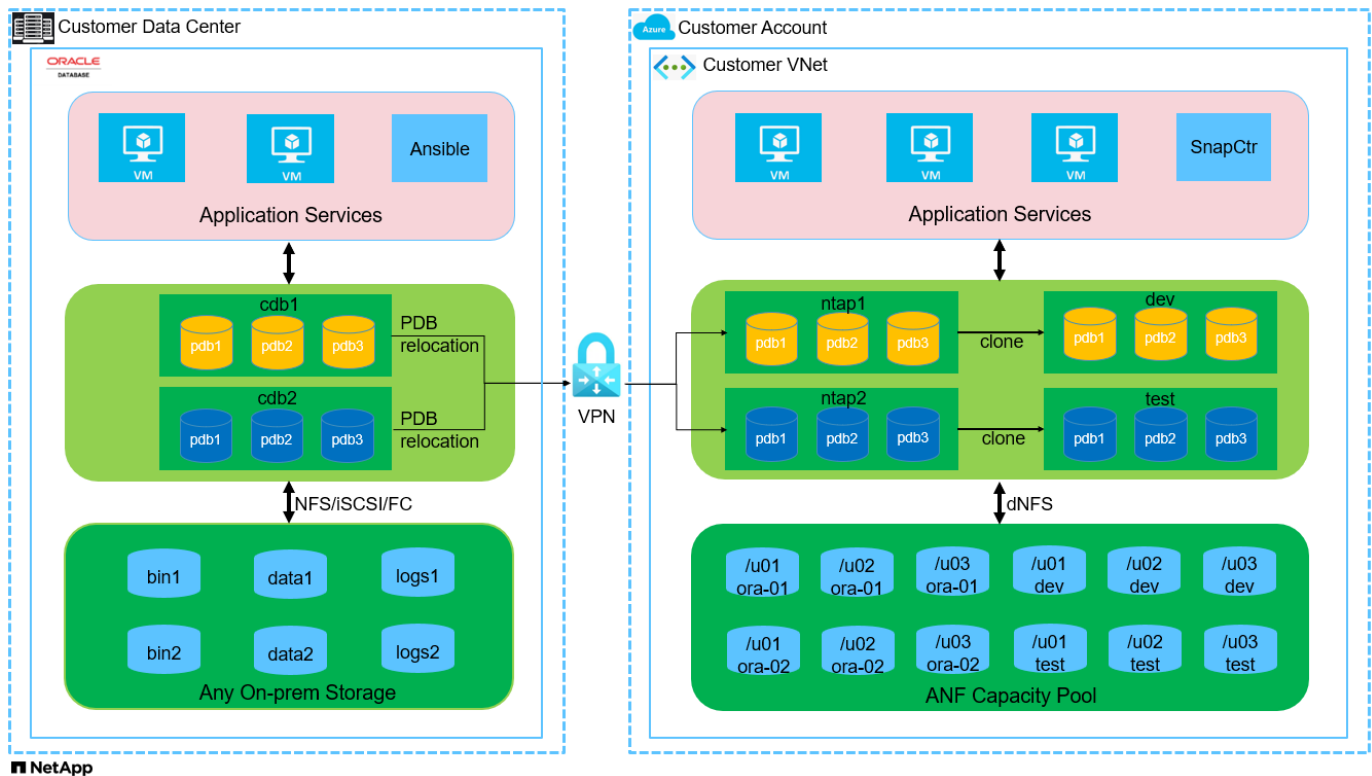
This solution is intended for the following people:

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

Solution test and validation environment

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

Architecture



Hardware and software components

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

Oracle database configuration in the lab environment

Server	Database	DB Storage
ora-01	NTAP1(NTAP1_PDB1,NTAP1_PDB2,NTAP1_PDB3)	/u01, /u02, /u03 NFS mounts on ANF capacity pool
ora-02	NTAP2(NTAP2_PDB1,NTAP2_PDB2,NTAP2_PDB3)	/u01, /u02, /u03 NFS mounts on ANF capacity pool

Key factors for deployment consideration

- **SnapCenter deployment.** SnapCenter can deploy in a Windows domain or Workgroup environment. For domain-based deployment, the domain user account should be a domain administrator account, or the domain user belongs to the local administrator's group on the SnapCenter hosting server.
- **Name resolution.** SnapCenter server needs to resolve the name to the IP address for each managed target database server host. Each target database server host must resolve the SnapCenter server name to the IP address. If a DNS server is unavailable, add naming to local host files for resolution.
- **Resource group configuration.** Resource group in SnapCenter is a logical grouping of similar resources that can be backed up together. Thus, it simplifies and reduces the number of backup jobs in a large database environment.
- **Separate full database and archive log backup.** Full database backup includes data volumes and log volumes consistent group snapshots. A frequent full database snapshot incurs higher storage consumption but improves RTO. An alternative is less frequent full database snapshots and more frequent archive logs backup, which consumes less storage and improves RPO but may extend RTO. Consider your RTO and RPO objectives when setting up the backup scheme. There is also a limit (1023) of the number of snapshot backups on a volume.
- **Privileges delegation.** Leverage role based access control that is built-in within SnapCenter UI to delegate privileges to application and database teams if desired.

Solution deployment

The following sections provide step-by-step procedures for SnapCenter deployment, configuration, and Oracle database backup, recovery, and clone on Azure NetApp Files in the Azure cloud.

Prerequisites for deployment

Deployment requires existing Oracle databases running on ANF in Azure. If not, follow the steps below to create two Oracle databases for solution validation. For details of Oracle database deployment on ANF in Azure cloud with automation, referred to TR-4987: [Simplified, Automated Oracle Deployment on Azure NetApp Files with NFS](#)

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



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

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

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

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



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

5. Clone a copy of the NetApp Oracle deployment automation toolkit for NFS. Follow instructions in [TR-4887](#) to execute the playbooks.

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

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

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

7. Watch the following video:

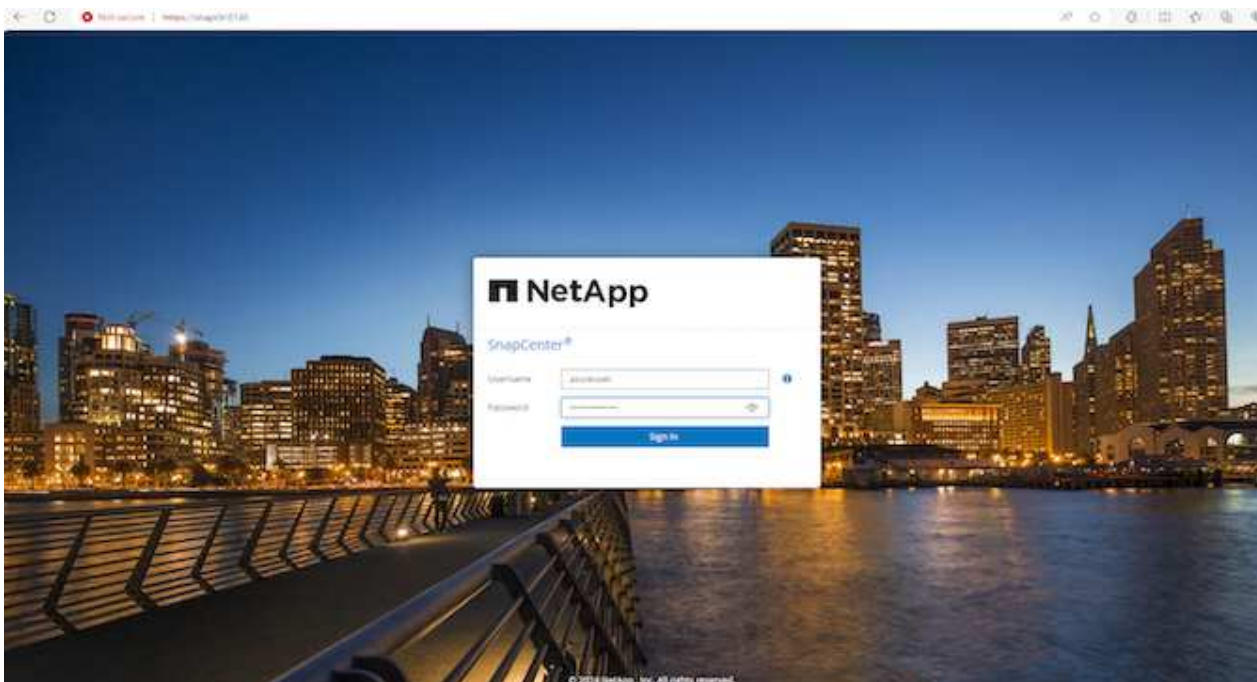
[Oracle Database Backup, Recovery, and Clone on ANF with SnapCenter](#)

8. Review the `Get Started` online menu.

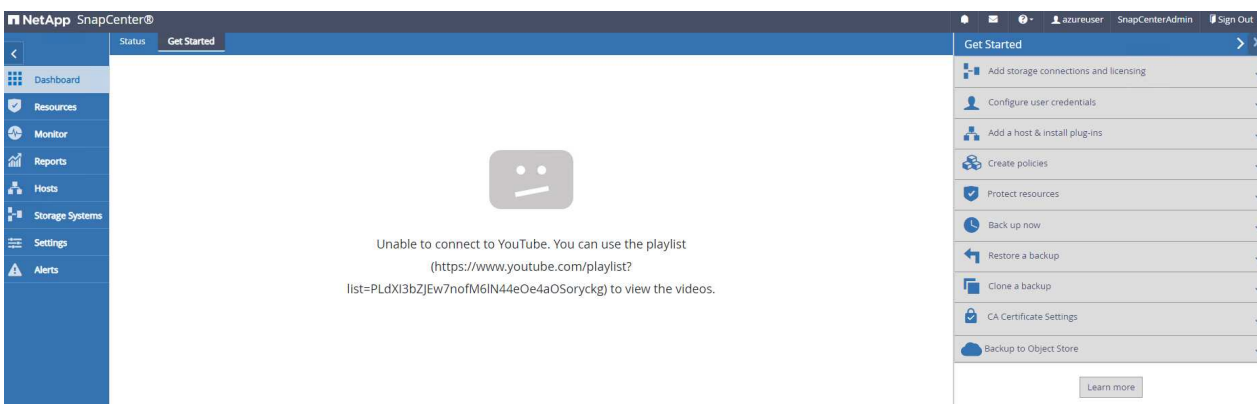
SnapCenter installation and setup

We recommend to go through online [SnapCenter Software documentation](#) before proceeding to SnapCenter installation and configuration: . Following provides a high level summary of steps for installation and setup of SnapCenter software for Oracle on Azure ANF.


1. From SnapCenter Windows server, download and install latest java JDK from [Get Java for desktop applications](#).
2. From SnapCenter Windows server, download and install latest version (currently 5.0) of SnapCenter installation executable from NetApp support site: [NetApp | Support](#).
3. After SnapCenter server installation, launch browser to login to SnapCenter with Windows local admin user or domain user credential via port 8146.



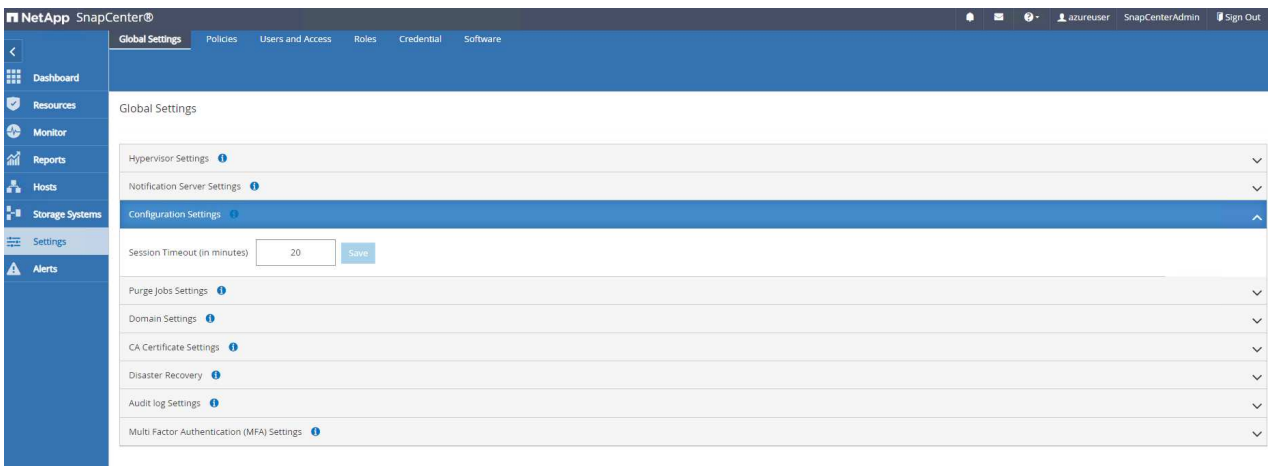
4. Review Get Started online menu.




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




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



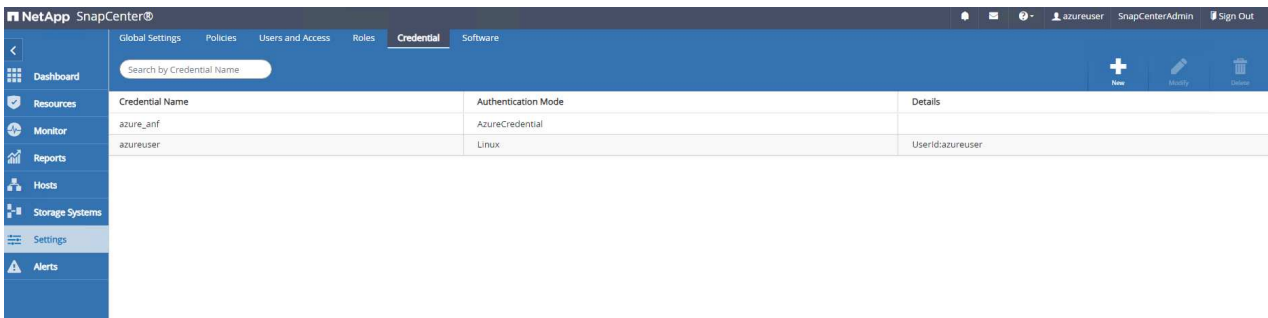
7. Add additional users to SnapCenter if needed.



8. The Roles tab list the built-in roles that can be assigned to different SnapCenter users. Custom roles also can be created by admin user with desired privileges.



- From Settings-Credential, create credentials for SnapCenter management targets. In this demo use case, they are linux user for login to Azure VM and ANF credential for capacity pool access.



Credential

Credential Name

azureuser

Authentication Mode

Linux

Authentication Type

☐ Password Based
 ☒ SSH Key Based

Username

azureuser

SSH Private Key

XRlrK1QCaEOHg==
 -----END RSA PRIVATE KEY-----

☒ Use sudo privileges

Cancel

OK

Credential

Credential Name

azure_anf

Authentication Mode

Azure Credential

Azure Details

Tenant ID

Enter Tenant Id

Client ID

Enter Client Id

Client Secret Key

Enter client secret key

Cancel

OK

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

NetApp SnapCenter®

ONTAP Storage

Azure NetApp Files

Dashboard

Resources

Monitor

Reports

Hosts

Storage Systems

Settings

Alerts

Search by NetApp Account

NetApp Account

Resource Group

Credential

<input type="checkbox"/>	ANFAVSAcct	ANFAVSRG	azure_anf
--------------------------	------------	----------	-----------

Add Azure NetApp Account

Credential

azure_anf

Subscription

Hybrid Cloud TME Onprem

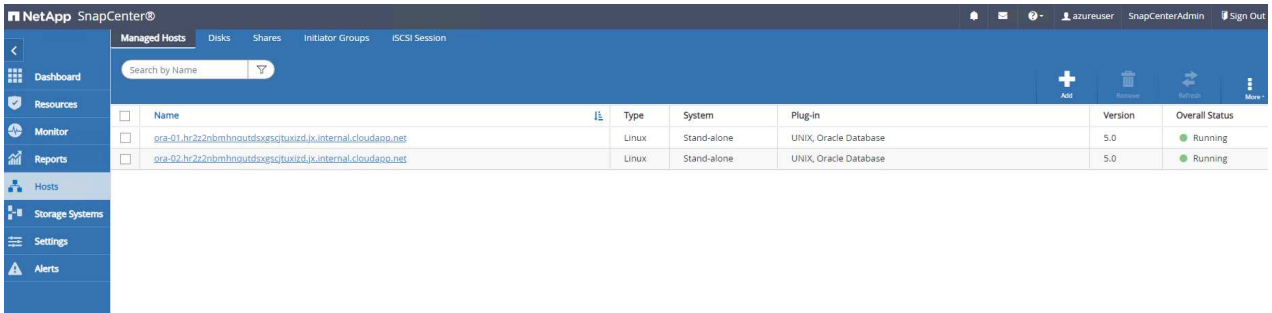
NetApp Account

ANFAVSAcct (ResourceGroup: ANFAVSRG)

Submit

Cancel

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



	Name	Type	System	Plug-in	Version	Overall Status
<input type="checkbox"/>	ora-01.hr222nbnh0u0d0xsgtjucrdjx.internal.cloudapp.net	Linux	Stand-alone	UNIX, Oracle Database	5.0	Running
<input type="checkbox"/>	ora-02.hr222nbnh0u0d0xsgtjucrdjx.internal.cloudapp.net	Linux	Stand-alone	UNIX, Oracle Database	5.0	Running

Add Host

Host Type

Host Name

Credentials  

Select Plug-ins to Install SnapCenter Plug-ins Package 5.0 for Linux

- ☒ Oracle Database
- ☐ SAP HANA
- ☐ Unix File Systems

 [More Options](#): Port, Install Path, Custom Plug-Ins...

More Options

Port

8145

Installation Path

/opt/NetApp/snapcenter

☒

Skip optional preinstall checks

☒

Add all hosts in the oracle RAC

Custom Plug-ins

Choose a File

Browse

Upload

No plug-ins found.

Save

Cancel

- Once host plugin is installed on DB server VM, databases on the host are auto discovered and visible in Resources tab. Back to Settings-Policies, create backup policies for full Oracle database online backup and archive logs only backup. Refer to this document [Create backup policies for Oracle databases](#) for detailed step by step procedures.

NetApp SnapCenter®

Global Settings

Policies

Users and Access

Roles

Credential

Software

Dashboard

Resources

Monitor

Reports

Hosts

Storage Systems

Settings

Alerts

Oracle Database

Search by Name

+

Modify

Copy

Details

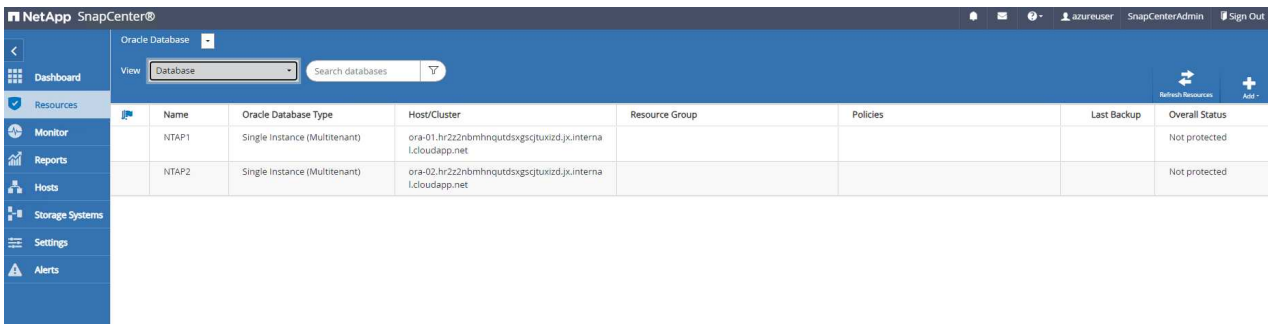
Delete

Name	Backup Type	Schedule Type	Replication	Verification
Oracle archivelogs backup	LOG, ONLINE	Hourly		
Oracle full online backup	FULL, ONLINE	Hourly		

Database backup

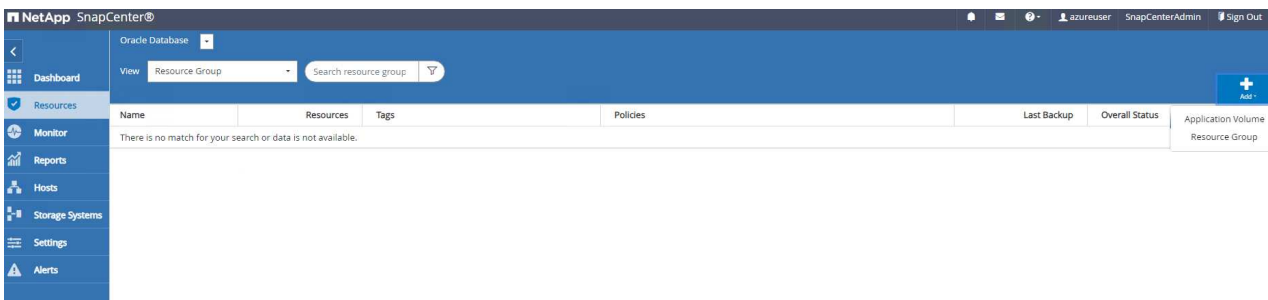
A NetApp snapshot backup creates a point-in-time image of the database volumes that you can use to restore in case of a system failure or data loss. Snapshot backups take very little time, usually less than a minute. The backup image consumes minimal storage space and incurs negligible performance overhead because it records only changes to files since the last snapshot copy was made. Following section demonstrates the implementation of snapshots for Oracle database backup in SnapCenter.

1. Navigating to Resources tab, which lists the databases discovered once SnapCenter plugin installed on database VM. Initially, the Overall Status of database shows as Not protected.



Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Overall Status
NTAP1	Single Instance (Multitenant)	ora-01.hr2z2nbnhngutdsxsgtuxizd.jc.interna l.cloudapp.net				Not protected
NTAP2	Single Instance (Multitenant)	ora-02.hr2z2nbnhngutdsxsgtuxizd.jc.interna l.cloudapp.net				Not protected

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



Name	Resources	Tags	Policies	Last Backup	Overall Status	Application Volume
There is no match for your search or data is not available.						Resource Group

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

New Resource Group

1
2
3
4
5
6
NameResourcesPoliciesVerificationNotificationSummary

Provide a name and tags for the resource group

Name
full_online_bkup

Tags
oradata

☒ Use custom name format for Snapshot copy

\$HostName

Backup settings

Exclude archive log destinations from backup

PreviousNext

4. Add resources to your Resource Group. Grouping of similar resources can simplify database management in a large environment.

New Resource Group

1
2
3
4
5
6
NameResourcesPoliciesVerificationNotificationSummary

Add resources to Resource Group

Host
All

Available Resources

search available resources

Selected Resources

NTAP1 (ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.
NTAP2 (ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.

>
<

PreviousNext

5. Select the backup policy and set a schedule by click on '+' sign under Configure Schedules.



Select one or more policies and configure schedules

Oracle full online backup + ⓘ

Configure schedules for selected policies

Policy	Applied Schedules	Configure Schedules
Oracle full online backup	None	+

Total 1

Previous Next

Add schedules for policy Oracle full online backup

Hourly

Start date

02/06/2024 05:55 pm



☐ Expires on

03/06/2024 05:51 pm



Repeat every

2



hours

0

mins



The schedules are triggered in the SnapCenter Server time zone.



Cancel

OK

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

New Resource Group

1

2

3

4

5

6

Name

Resources

Policies

Verification

Notification

Summary

Configure verification schedules

Policy

123

Schedule Type

Applied Schedules

Configure Schedules

There is no match for your search or data is not available.

Total 0

Previous

Next

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

New Resource Group

1

2

3

4

5

6

Name

Resources

Policies

Verification

Notification

Summary

Provide email settings ⓘ

Select the service accounts or people to notify regarding protection issues.

Email preference

Never

From

From email

To

Email to

Subject

Notification

☐ Attach job report

Previous

Next

8. Summary of new resource group.

Backup

×

Create a backup for the selected resource group

Resource Group

full_online_bkup

Policy

Oracle full online backup ▾

i

☐ Verify after backup

Cancel

Backup

11. Click on the running job to open a monitoring window, which allows the operator to track the job progress in real-time.

Job Details



Backup of Resource Group 'full_online_bkup' with policy 'Oracle full online backup'

✓ ▾ Backup of Resource Group 'full_online_bkup' with policy 'Oracle full online backup'

✓ ▶ ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net

✓ ▶ ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net

📌 Task Name: Backup of Resource Group 'full_online_bkup' with policy 'Oracle full online backup' Start Time: 02/06/2024 6:00:05 PM End Time: 02/06/2024 6:00:44 PM

View Logs

Cancel Job

Close

12. A snapshot backup set appears under database topology once a successful backup job finishes. A full database backup set includes a snapshot of the database data volumes and a snapshot of the database log volumes. A log-only backup contains only a snapshot of the database log volumes.

NetApp SnapCenter

azureuser SnapCenterAdmin Sign Out

Oracle Database

Search resource groups

full_online_bkup Details

search

NTAP1 Topology

Close

Search resource groups

Search

Resource Name

NTAP1

NTAP2

Name

full_online_bkup

archivelog_bkup

Manage Copies

3 Backups

0 Clones

Local copies

Summary Card

3 Backups

1 Data Backup

2 Log Backups

0 Clones

0 Snapshots Locked

Primary Backup(s)

search

Cancel Remove Clone Restore Move Compact Delete

Backup Name	Snapshot Lock Expiration	Count	Type	End Date	Verified	Mounted	RMAN Cataloged	SCN
ora-01_02-06-2024_18_00_06_0582_1		1	Log	02/06/2024 6:00:41 PM	Not Applicable	False	Not Cataloged	3374950
ora-01_02-06-2024_18_00_06_0582_0		1	Data	02/06/2024 6:00:26 PM	Unverified	False	Not Cataloged	3374903
ora-01_02-06-2024_17_59_01_1158_1		1	Log	02/06/2024 5:59:18 PM	Not Applicable	False	Not Cataloged	3374762

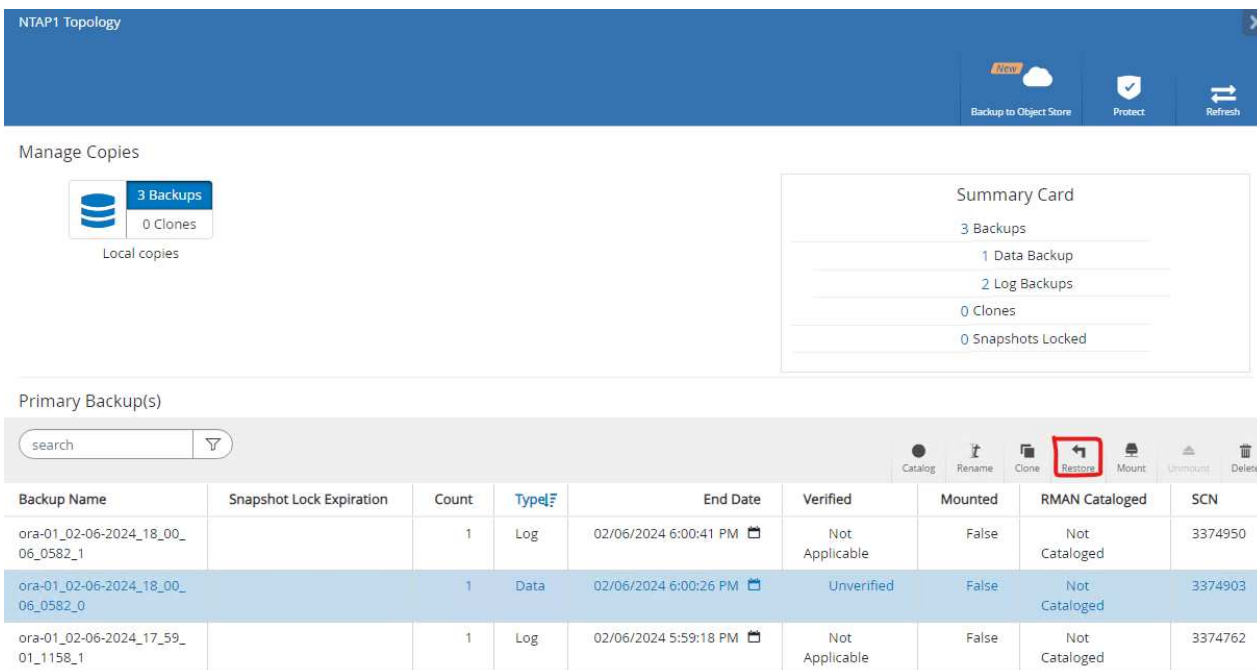
Total 2

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 Object Store Protect Refresh

Manage Copies

3 Backups
0 Clones
Local copies

Summary Card

3 Backups

1 Data Backup

2 Log Backups

0 Clones

0 Snapshots Locked

Primary Backup(s)

search

Catalog Rename Clone **Restore** Mount Unmount Delete

Backup Name	Snapshot Lock Expiration	Count	Type	End Date	Verified	Mounted	RMAN Cataloged	SCN
ora-01_02-06-2024_18_00_06_0582_1		1	Log	02/06/2024 6:00:41 PM	Not Applicable	False	Not Cataloged	3374950
ora-01_02-06-2024_18_00_06_0582_0		1	Data	02/06/2024 6:00:26 PM	Unverified	False	Not Cataloged	3374903
ora-01_02-06-2024_17_59_01_1158_1		1	Log	02/06/2024 5:59:18 PM	Not Applicable	False	Not Cataloged	3374762

2. Select **Restore Scope**. For a container database, SnapCenter is flexible to perform a full container database (All Datafiles), pluggable databases, or tablespaces level restore.

Restore NTAP1

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Restore Scope ⓘ

☒ All Datafiles

☐ Pluggable databases (PDBs)

☐ Pluggable database (PDB) tablespaces

☐ Control files

Database State

☒ Change database state if needed for restore and recovery

Restore Mode ⓘ

☐ Force in place restore

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

Previous

Next

3. Select Recovery Scope. All logs means to apply all available archive logs in the backup set. Point-in-time recovery by SCN or timestamp are also available.

Restore NTAP1

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Choose Recovery Scope

☒ All Logs

☐ Until SCN (System Change Number)

☐ Date and Time

☐ No recovery

Specify external archive log files locations

Previous

Next

4. The `PreOps` allows execution of scripts against database before restore/recovery operation.

Restore NTAP1

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Specify optional scripts to run before performing a restore job

Prescript full path

/var/opt/snapcenter/spl/scripts/

Enter Prescript path

Arguments

Script timeout

60

secs

Previous

Next

5. The PostOps allows execution of scripts against database after restore/recovery operation.

Restore NTAP1

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Specify optional scripts to run after performing a restore job

Postscript full path

/var/opt/snapcenter/spl/scripts/

Enter Postscript path

Arguments

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

Previous

Next

6. Notification via email if desired.

Restore NTAP1

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Provide email settings ⓘ

Email preference

Never

From

From email

To

Email to

Subject

Notification

☐ Attach job report

⚠ If you want to send notifications for Restore jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server.

Previous

Next

7. Restore job summary

Restore NTAP1

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Summary

Backup name	ora-01_02-06-2024_18_00_06_0582_0
Backup date	02/06/2024 6:00:26 PM
Restore scope	All DataFiles
Recovery scope	All Logs
Options	Change database state if necessary , Open the database or container database in READ-WRITE mode after recovery
Prescript full path	None
Prescript arguments	
Postscript full path	None
Postscript arguments	
Send email	No

Previous

Finish

8. Click on running job to open Job Details window. The job status can also be opened and viewed from the Monitor tab.

Job Details



Restore 'ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net\NTAP1'

✓ ▾ Restore 'ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net\NTAP1'

✓ ▾ ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net

- ✓ ▶ Prescripts
- ✓ ▶ Mount log backups
- ✓ ▶ Pre Restore
- ✓ ▶ Restore
- ✓ ▶ Post Restore
- ✓ ▶ Unmount log backups
- ✓ ▶ Postscripts
- ✓ ▶ Post Restore Cleanup
- ✓ ▶ Data Collection

📌 Task Name: ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net Start Time: 02/06/2024 4:04:55 PM End Time: 02/06/2024 4:08:42 PM

View Logs

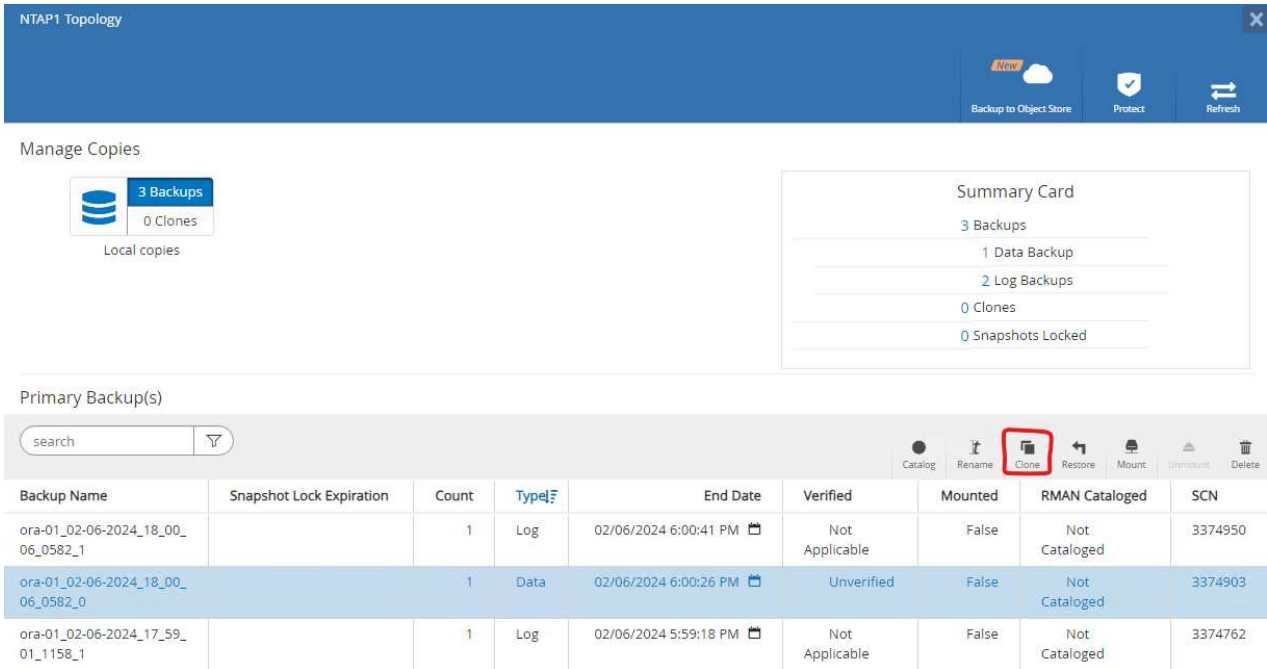
Cancel Job

Close

Database clone

Database clone via SnapCenter is accomplished by creating a new volume from a snapshot of a volume. The system uses the snapshot information to clone a new volume using the data on the volume when the snapshot was taken. More importantly, it is quick (a few minutes) and efficient compared with other methods to make a cloned copy of the production database to support development or testing. Thus, dramatically improve your database application lifecycle management. The following section demonstrates the workflow of database clone with SnapCenter UI.

1. From Resources tab, open the database Primary Backup(s) page. Choose the snapshot of database data volume, then click on clone button to launch database clone workflow.



NTAP1 Topology

Manage Copies

3 Backups
0 Clones
Local copies

Summary Card

- 3 Backups
- 1 Data Backup
- 2 Log Backups
- 0 Clones
- 0 Snapshots Locked

Primary Backup(s)

search

Clone

Backup Name	Snapshot Lock Expiration	Count	Type	End Date	Verified	Mounted	RMAN Cataloged	SCN
ora-01_02-06-2024_18_00_06_0582_1		1	Log	02/06/2024 6:00:41 PM	Not Applicable	False	Not Cataloged	3374950
ora-01_02-06-2024_18_00_06_0582_0		1	Data	02/06/2024 6:00:26 PM	Unverified	False	Not Cataloged	3374903
ora-01_02-06-2024_17_59_01_1158_1		1	Log	02/06/2024 5:59:18 PM	Not Applicable	False	Not Cataloged	3374762

2. Name the clone database SID. Optionally, for a container database, clone can be done at PDB level as well.

Clone from NTAP1

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Capacity Pool Max. Throughput (MiB/s)

Complete Database Clone

Clone SID

Exclude PDBs

ntap1dev

Type to find PDBs

PDB Clone

Previous

Next

3. Select the DB server where you want to place your cloned database copy. Keep the default file locations unless you want to name them differently.

29

Clone from NTAP1

×

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Select the host to create a clone

Clone host

ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.inter

Datafile locations ⓘ

/u02_ntap1dev

Reset

Control files ⓘ

/u02_ntap1dev/ntap1dev/control/control01.ctl

/u02_ntap1dev/ntap1dev/control/control02.ctl

Reset

Redo logs ⓘ

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

Reset

Previous

Next

- 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.

30

Clone from NTAP1

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Database Credentials for the clone

Credential name for sys userNone+ ⓘ

Database port1521

Oracle Home Settings ⓘ

Oracle Home/u01/app/oracle/product/19.0.0/NTAP2

Oracle OS Useroracle

Oracle OS Groupoinstall

Previous

Next

5. The `PreOps` allows execution of scripts before clone operation. Database parameters can be adjusted to meet a clone DB needs as versus a production database, such as reduced SGA target.

Clone from NTAP1

1

Name

2

Locations

3

Credentials

4

5

6

7

Specify scripts to run before clone operation ⓘ

Prescript full path

/var/opt/snapcenter/spl/scripts/
Enter Prescript path

Arguments

Script timeout

60
secs

Database Parameter settings

Previous
Next

- The PostOps allows execution of scripts against database after clone operation. Clone database recovery can be SCN, timestamp based, or Until cancel (rolling forward database to last archived log in the backup set).

Clone from NTAP1



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

☒ Recover Database

☒ Until Cancel






☐ Date and Time



Date-time format: MM/DD/YYYY hh:mm:ss

☐ Until SCN (System Change Number)



Specify external archive log locations   

☒ Create new DBID 

☒ Create tempfile for temporary tablespace 

 Enter SQL queries to apply when clone is created

 Enter scripts to run after clone operation 

Previous

Next

7. Notification via email if desired.

1 Name

Provide email settings ⓘ

2 Locations

Email preference

Never ▾

3 Credentials

From

From email

4 PreOps

To

Email to

5 PostOps

Subject

Notification

6 Notification

☐ Attach job report

7 Summary



If you want to send notifications for Clone jobs, an SMTP server must be configured. Continue to the Summary page to save your information, and then go to Settings>Global Settings>Notification Server Settings to configure the SMTP server.

Previous

Next

8. Clone job summary.

Clone from NTAP1

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Summary

Clone from backup	ora-01_02-06-2024_18_00_06_0582_0
Clone SID	ntap1dev
Capacity Pool Max. Throughput (MiB/s)	none
Clone server	ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net
Exclude PDBs	none
Oracle home	/u01/app/oracle/product/19.0.0/NTAP2
Oracle OS user	oracle
Oracle OS group	oinstall
Datafile mountpaths	/u02_ntap1dev
Control files	/u02_ntap1dev/ntap1dev/control/control01.ctl /u02_ntap1dev/ntap1dev/control/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

Finish

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.cloudapp.net Start Time: 02/06/2024 6:21:59 PM End Time: 02/06/2024 6:28:10 PM

View Logs

Cancel Job

Close

10. Cloned database registers with SnapCenter immediately.

NetApp SnapCenter®								
Oracle Database		View Database Search databases						
	Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Overall Status	
	NTAP1	Single Instance (Multitenant)	ora-01.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net	archivelog_bkup full_online_bkup	Oracle archivelogs backup Oracle full online backup	02/06/2024 7:29:18 PM	Backup succeeded	
	<u>ntap1dev</u>	Single Instance (Multitenant)	ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net				Not protected	
	NTAP2	Single Instance (Multitenant)	ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net	archivelog_bkup full_online_bkup	Oracle archivelogs backup Oracle full online backup	02/06/2024 7:29:19 PM	Backup succeeded	

11. Validate clone database on DB server host. For a cloned development database, database archive

36

mode should be turned off.

```
[azureuser@ora-02 ~]$ sudo su
[root@ora-02 azureuser]# su - oracle
Last login: Tue Feb  6 16:26:28 UTC 2024 on pts/0

[oracle@ora-02 ~]$ uname -a
Linux ora-02 4.18.0-372.9.1.el8.x86_64 #1 SMP Fri Apr 15 22:12:19
EDT 2022 x86_64 x86_64 x86_64 GNU/Linux
[oracle@ora-02 ~]$ df -h
```

Filesystem	Size	Used	Avail
Use% Mounted on			
devtmpfs	7.7G	0	7.7G
0% /dev			
tmpfs	7.8G	0	7.8G
0% /dev/shm			
tmpfs	7.8G	49M	7.7G
1% /run			
tmpfs	7.8G	0	7.8G
0% /sys/fs/cgroup			
/dev/mapper/rootvg-rootlv	22G	17G	5.6G
75% /			
/dev/mapper/rootvg-usrlv	10G	2.0G	8.1G
20% /usr			
/dev/mapper/rootvg-homelv	1014M	40M	975M
4% /home			
/dev/sda1	496M	106M	390M
22% /boot			
/dev/mapper/rootvg-varlv	8.0G	958M	7.1G
12% /var			
/dev/sda15	495M	5.9M	489M
2% /boot/efi			
/dev/mapper/rootvg-tmplv	12G	8.4G	3.7G
70% /tmp			
tmpfs	1.6G	0	1.6G
0% /run/user/54321			
172.30.136.68:/ora-02-u03	250G	2.1G	248G
1% /u03			
172.30.136.68:/ora-02-u01	100G	10G	91G
10% /u01			
172.30.136.68:/ora-02-u02	250G	7.5G	243G
3% /u02			
tmpfs	1.6G	0	1.6G
0% /run/user/1000			
tmpfs	1.6G	0	1.6G
0% /run/user/0			

```
172.30.136.68:/ora-01-u02-Clone-020624161543077 250G 8.2G 242G
4% /u02_ntapldev
```

```
[oracle@ora-02 ~]$ cat /etc/oratab
#
# This file is used by ORACLE utilities.  It is created by root.sh
# and updated by either Database Configuration Assistant while
# creating
# a database or ASM Configuration Assistant while creating ASM
# instance.
#
# A colon, ':', is used as the field terminator.  A new line
# terminates
# the entry.  Lines beginning with a pound sign, '#', are comments.
#
# Entries are of the form:
#   $ORACLE_SID:$ORACLE_HOME:<N|Y>:
#
# The first and second fields are the system identifier and home
# directory of the database respectively.  The third field indicates
# to the dbstart utility that the database should , "Y", or should
# not,
# "N", be brought up at system boot time.
#
# Multiple entries with the same $ORACLE_SID are not allowed.
#
#
NTAP2:/u01/app/oracle/product/19.0.0/NTAP2:Y
# SnapCenter Plug-in for Oracle Database generated entry (DO NOT
# REMOVE THIS LINE)
ntapldev:/u01/app/oracle/product/19.0.0/NTAP2:N
```

```
[oracle@ora-02 ~]$ export ORACLE_SID=ntapldev
[oracle@ora-02 ~]$ sqlplus / as sysdba
```

```
SQL*Plus: Release 19.0.0.0.0 - Production on Tue Feb 6 16:29:02 2024
Version 19.18.0.0.0
```

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

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



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

NAME	OPEN_MODE	LOG_MODE
NTAP1DEV	READ WRITE	ARCHIVELOG

```
SQL> shutdown immediate;
```

Database closed.

Database dismounted.

ORACLE instance shut down.

```
SQL> startup mount;
```

ORACLE instance started.

Total System Global Area 3221223168 bytes

Fixed Size 9168640 bytes

Variable Size 654311424 bytes

Database Buffers 2550136832 bytes

Redo Buffers 7606272 bytes

Database mounted.

```
SQL> alter database noarchivelog;
```

Database altered.

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

Database altered.

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

NAME	OPEN_MODE	LOG_MODE
NTAP1DEV	READ WRITE	NOARCHIVELOG

```
SQL> show pdbs
```

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

```
SQL> alter pluggable database all open;
```

Where to find additional information

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

- Azure NetApp Files

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

- SnapCenter Software documentation

<https://docs.netapp.com/us-en/snapcenter/index.html>

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

[Deployment Procedure](#)

Copyright information

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

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

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

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

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

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

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

Trademark information

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