



# **SnapCenter for Databases**

## **NetApp Solutions**

NetApp  
March 19, 2024

This PDF was generated from [https://docs.netapp.com/us-en/netapp-solutions/databases/snapctr\\_ora\\_azure\\_anf.html](https://docs.netapp.com/us-en/netapp-solutions/databases/snapctr_ora_azure_anf.html) on March 19, 2024. Always check docs.netapp.com for the latest.

# Table of Contents

- SnapCenter for Databases ..... 1
  - TR-4988: Oracle Database Backup, Recovery, and Clone on ANF with SnapCenter ..... 1
  - TR-4977: Oracle Database backup, restore and clone with SnapCenter Services - Azure ..... 40
  - TR-4964: Oracle Database backup, restore and clone with SnapCenter Services - AWS ..... 74
  - Hybrid Cloud Database Solutions with SnapCenter ..... 107

# SnapCenter for Databases

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

Allen Cao, Niyaz Mohamed, NetApp

### 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

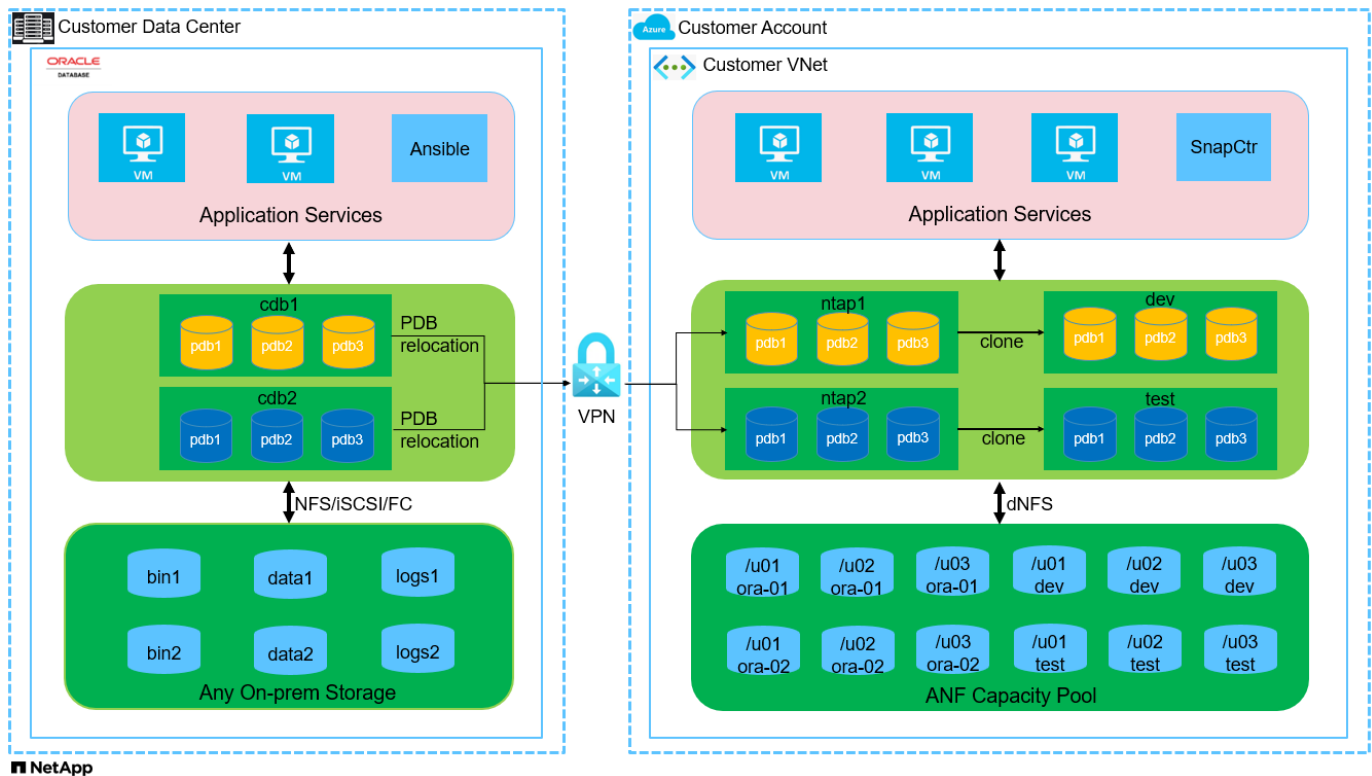
This solution is intended for the following people:

- 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:

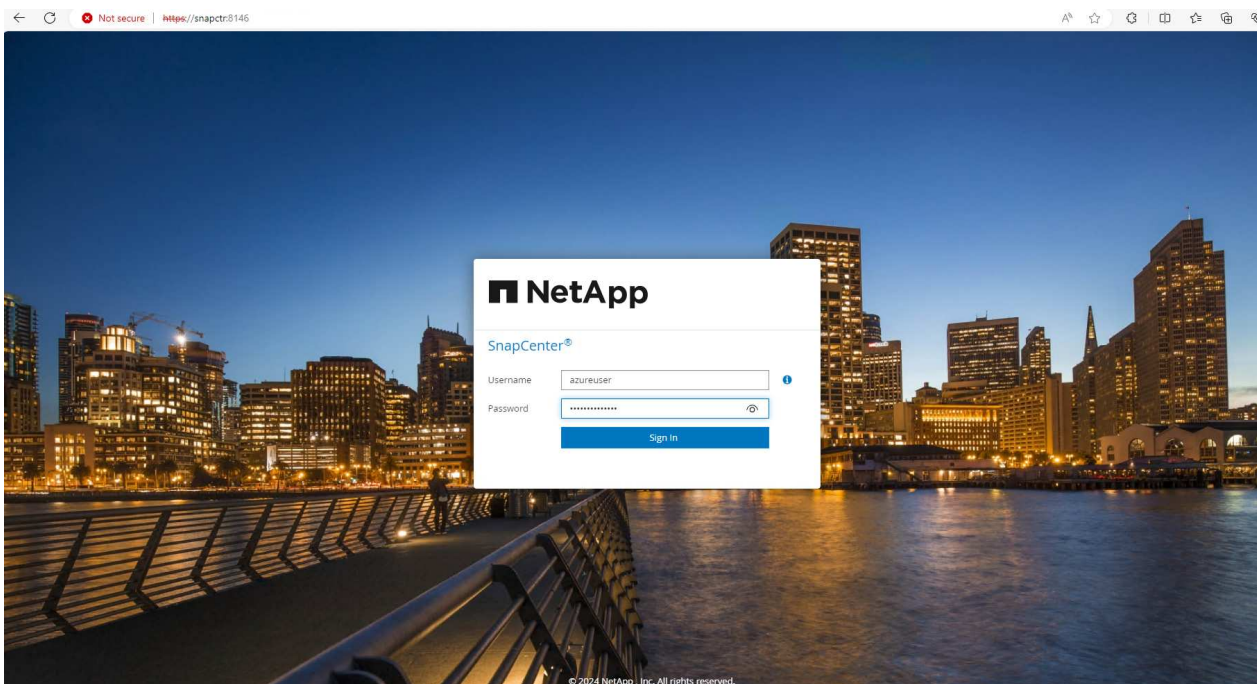
[Oracle Database Backup, Recovery, and Clone on ANF with SnapCenter](#)

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.



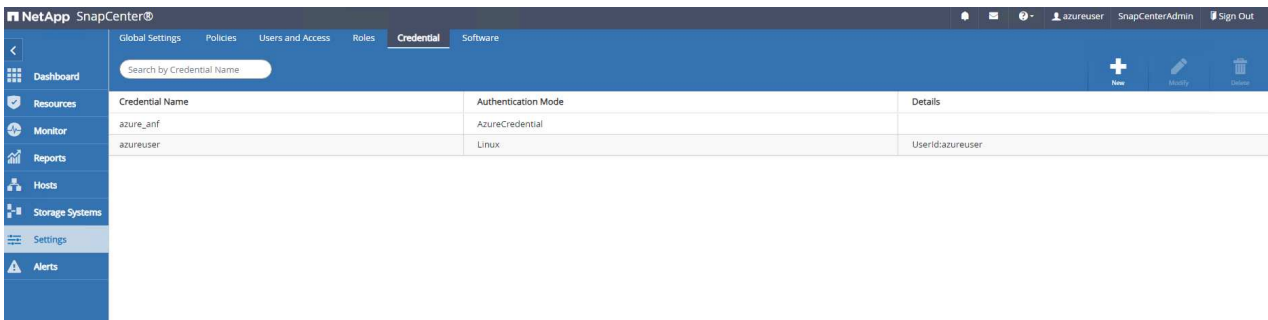
7. Add additional users to SnapCenter if needed.



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

XRlRk1QCaE0Hg==  
 -----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

ANFAVSAcct

ANFAVSRG

azure\_anf

Add Azure NetApp Account

Credential

azure\_anf

Subscription

Hybrid Cloud TME Onprem

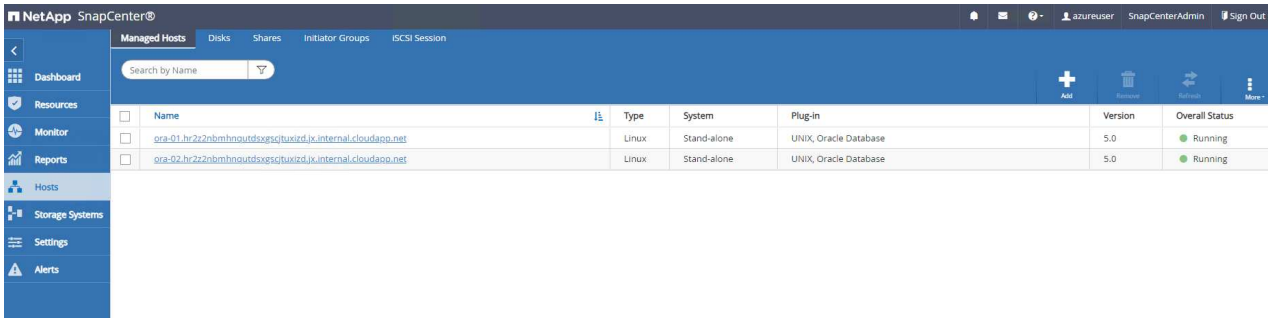
NetApp Account

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
ora-01.hr2z2nbmhnoutdssqstuwzdx.internal.cloudapp.net	Linux	Stand-alone	UNIX, Oracle Database	5.0	Running
ora-02.hr2z2nbmhnoutdssqstuwzdx.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
i

Installation Path
i

☒ Skip optional preinstall checks
i

☒ Add all hosts in the oracle RAC

Custom Plug-ins

Choose a File

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

NetApp SnapCenter®
azureuser SnapCenterAdmin Sign Out

Global Settings Policies Users and Access Roles Credential Software

Oracle Database
+
✎
📄
i
🗑

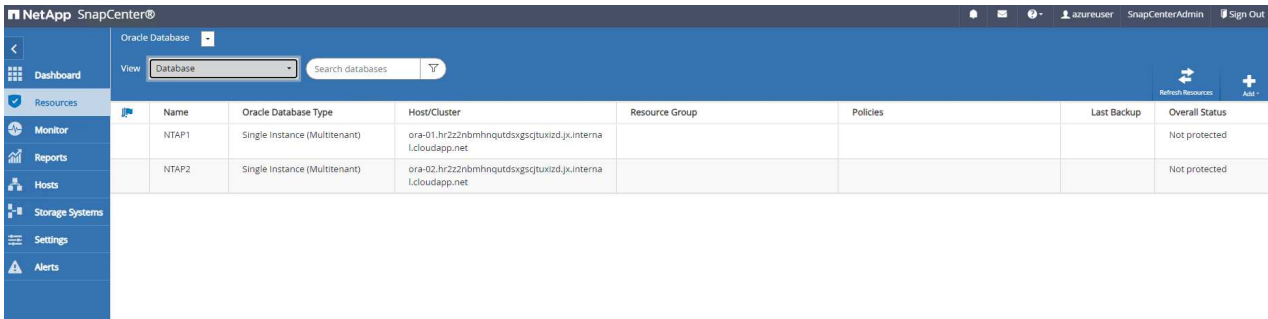
Name	Backup Type	Schedule Type	Replication	Verification
Oracle archivelogs backup	LOG, ONLINE	Hourly		
Oracle full online backup	FULL, ONLINE	Hourly		

Dashboard
Resources
Monitor
Reports
Hosts
Storage Systems
Settings
Alerts

## Database backup

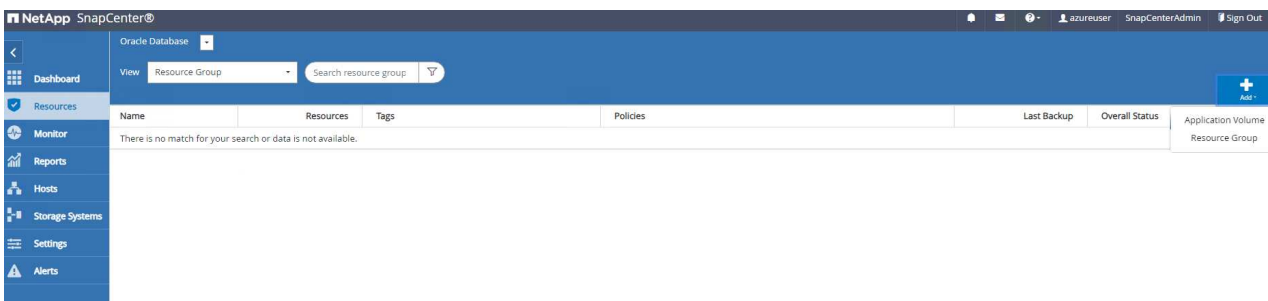
A NetApp snapshot backup creates a point-in-time image of the database volumes that you can use to restore in case of a system failure or data loss. Snapshot backups take very little time, usually less than a minute. The backup image consumes minimal storage space and incurs negligible performance overhead 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.hr22nbnmhnqutdsxsgtuxizd.jx.interna l.cloudapp.net				Not protected
NTAP2	Single Instance (Multitenant)	ora-02.hr22nbnmhnqutdsxsgtuxizd.jx.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.hr2z2nbmhnqutdsxgsqtuxizd.jk.internal.cloudapp.net)
NTAP2 (ora-02.hr2z2nbmhnqutdsxgsqtuxizd.jk.internal.cloudapp.net)

>
<

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

1

2

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.

New Resource Group

1
2
3
4
5
6

Name
Resources
Policies
Verification
Notification
Summary

Resource group name
Tags
Policy
Plug-in
Verification enabled for policy
Send email

full\_online\_bkup
oradata
Oracle full online backup: Hourly
SnapCenter Plug-in for Oracle Database
None
No

Previous
Finish

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

NetApp SnapCenter®

Dashboard
Resources
Monitor
Reports
Hosts
Storage Systems
Settings
Alerts

Oracle Database
View
Resource Group
Search resource group

Name
Resources
Tags
Policies
Last Backup
Overall Status

full\_online\_bkup
archivelog\_bkup

2
2

oradata
oralog

Oracle full online backup
Oracle archivelogs backup

02/06/2024 6:00:44 PM
02/06/2024 5:59:25 PM

Completed
Completed

10. Click on a resource group to reveal the resources it includes. Besides the scheduled backup job, an one-off backup can be triggered by clicking on Backup Now.

NetApp SnapCenter®

Oracle Database
Search resource groups

full\_online\_bkup Details
search

Modify Resource Group
Backup Now
Maintenance
Delete

Name
Resource Name
Type
Host

full\_online\_bkup
archivelog\_bkup

NTAP1
NTAP2

Oracle Database
Oracle Database

ora-01.hr222nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net
ora-02.hr222nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net

16

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

Manage Copies

3 Backups0 ClonesLocal copies

Summary Card

3 Backups1 Data Backup2 Log Backups0 Clones0 Snapshots Locked

Primary Backup(s)

search

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

New
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.

20

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

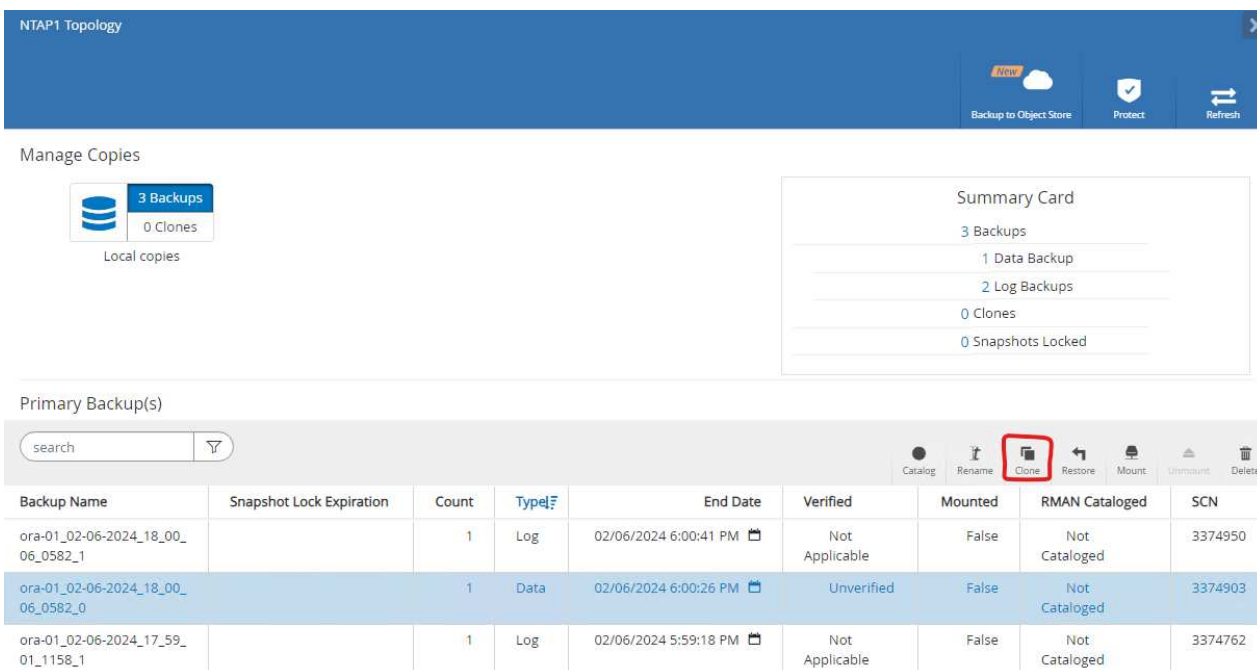
Cancel job

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

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

Clone from NTAP1

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Capacity Pool Max. Throughput (MiB/s)

Complete Database Clone

Clone SID

Exclude PDBs

PDB Clone

ntap1dev

Type to find PDBs

Previous

Next

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

Clone from NTAP1

×

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Select the host to create a clone

Clone host 

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	+	<div>+</div> <div>Reset</div>
▶ RedoGroup 2	×	200	MB	1	+	
▶ RedoGroup 3	×	200	MB	1	+	

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 user

None

+

i

Database port

1521

Oracle Home Settings

i

Oracle Home

/u01/app/oracle/product/19.0.0/NTAP2

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.

Clone from NTAP1

×

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Specify scripts to run before clone operation ⓘ

Prescript full path

/var/opt/snapcenter/spl/scripts/

Enter Prescript path

Arguments

Script timeout

60

secs

Database Parameter settings

processes	320	×	▲
remote_login_passwordfile	EXCLUSIVE	×	+
sga_target	3G	×	
undo_tablespace	UNDOTBS1	×	▼

Reset

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

- 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
		ntap1dev	Single Instance (Multitenant)	ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net				Not protected
		NTAP2	Single Instance (Multitenant)	ora-02.hr2z2nbmhnqutdsxgscjtuxizd.jx.internal.cloudapp.net	archivelog_bkup full_online_bkup	Oracle archivelogs backup Oracle full online backup	02/06/2024 7:29:19 PM	Backup succeeded

11. Validate clone database on DB server host. For a cloned development database, database archive mode should be turned off.

36

```

[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

[https://docs.netapp.com/us-en/netapp-solutions/databases/automation\\_ora\\_anf\\_nfs.html](https://docs.netapp.com/us-en/netapp-solutions/databases/automation_ora_anf_nfs.html)

# TR-4977: Oracle Database backup, restore and clone with SnapCenter Services - Azure

Allen Cao, Niyaz Mohamed, NetApp

## Purpose

SnapCenter Services is the SaaS version of the classic SnapCenter database management UI tool that is available through the NetApp BlueXP cloud management console. It is an integral part of the NetApp cloud-backup, data-protection offering for databases such as Oracle and HANA running on 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

#### Software

RedHat Linux	Red Hat Enterprise Linux 8.7 (LVM) - x64 Gen2	Deployed RedHat subscription for testing
Oracle Database	Version 19.18	Applied RU patch p34765931_190000_Linux-x86-64.zip
Oracle OPatch	Version 12.2.0.1.36	Latest patch p6880880_190000_Linux-x86-64.zip

## Key factors for deployment consideration

- **Connector to be deployed in the same virtual network / subnet as databases and Azure NetApp Files.** When possible, the connector should be deployed in the same Azure virtual networks and resource groups, which enables connectivity to the Azure NetApp Files storage and the Azure compute instances.
- **An Azure user account or Active Directory service principle created at Azure portal for SnapCenter connector.** Deploying a BlueXP Connector requires specific permissions to create and configure a virtual machine and other compute resources, to configure networking, and to get access to the Azure subscription. It also requires permissions to later create roles and permissions for the Connector to operate. Create a custom role in Azure with permissions and assign to the user account or service principle. Review the following link for details: [Set up Azure permissions](#).
- **A ssh key pair created in the Azure resource group.** The ssh key pair is assigned to the Azure VM user for logging into the connector host and also the database VM host for deploying and executing a plug-in. BlueXP console UI uses the ssh key to deploy SnapCenter service plugin to database host for one-step plugin installation and application host database discovery.
- **A credential added to the BlueXP console setting.** To add Azure NetApp Files storage to the BlueXP working environment, a credential that grants permissions to access Azure NetApp Files from the BlueXP console needs to be set up in the BlueXP console setting.
- **java-11-openjdk installed on the Azure VM database instance host.** SnapCenter service installation requires java version 11. It needs to be installed on application host before plugin deployment attempt.

## Solution deployment

There is extensive NetApp documentation with a broader scope to help you protect your cloud-native application data. The goal of this documentation is to provide step-by-step procedures that cover SnapCenter Service deployment with the BlueXP console to protect your Oracle database deployed on an Azure NetApp Files storage and an Azure compute instance.

To get started, complete the following steps:

- Read the general instructions [Protect your cloud native applications data](#) and the sections related to Oracle and Azure NetApp Files.
- Watch the following video walkthrough

[Video of deployment of Oracle and ANF](#)

## Prerequisites for SnapCenter service deployment



Deployment requires the following prerequisites.


1. A primary Oracle database server on an Azure VM instance with an Oracle database fully deployed and running.
2. An Azure NetApp Files storage service capacity pool deployed in Azure that has capacity to meet the database storage needs listed in hardware component section.
3. A secondary database server on an Azure VM instance that can be used for testing the cloning of an Oracle database to an alternate host for the purpose of supporting a dev/test workload or any use cases that requires a full data set of production Oracle database.
4. For additional information for Oracle database deployment on Azure NetApp Files and Azure compute instance, see [Oracle Database Deployment and Protection on Azure NetApp Files](#).

## Onboarding to BlueXP preparation


1. Use the link [NetApp BlueXP](#) to sign up for BlueXP console access.
2. Create an Azure user account or an Active Directory service principle and grant permissions with role in Azure portal for Azure connector deployment.
3. To set up BlueXP to manage Azure resources, add a BlueXP credential with details of an Active Directory service principal that BlueXP can use to authenticate with Azure Active Directory (App client ID), a client secret for the service principal application (Client Secret), and the Active Directory ID for your organization (Tenant ID).
4. You also need the Azure virtual network, resources group, security group, an SSH key for VM access, etc. ready for connector provisioning and database plugin installation.

## Deploy a connector for SnapCenter services

1. Login to the BlueXP console.



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



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

## Provider

Choose the cloud provider where you want to run the BlueXP Connector:



[Deploy the Connector on your premises](#)

Continue

4. Skip the **Permission**, **Authentication**, and **Networking** steps if you already have them set up in your Azure account. If not, you must configure these before proceeding. From here, you could also retrieve the permissions for the Azure policy that is referenced in the previous section "[Onboarding to BlueXP preparation](#)."

## Deploying a BlueXP Connector

The BlueXP Connector is a crucial component for the day-to-day use of BlueXP.

It's used to connect BlueXP's services to your hybrid-cloud environments.

The BlueXP Connector can then manage the resources and processes within your public cloud environment.

Before you begin the deployment process, ensure that you have completed the required preparations. This guide will enable you to focus on the minimum requirements for BlueXP Connector installation.

### Permissions

Ensure that the Azure user or service principal you've provided has sufficient permissions

### Authentication

Choose between two methods: an

[Azure user account](#) or an

[Active Directory service principal](#)

### Networking

Ensure that you have details on the VNet and subnet in which the BlueXP Connector will reside

[Skip to Deployment](#)

[Previous](#)

[Continue](#)



- Click on **Skip to Deployment** to configure your connector **Virtual Machine Authentication**. Add the SSH key pair you have created in Azure resource group during onboarding to BlueXP preparation for connector OS authentication.

1

VM Authentication

2

Details

3

Network

4

Security Group

5

Review

## Virtual Machine Authentication

You are logged in with Azure user: [acao@netapp.com](#)

Tenant: Hybrid Cloud TME

## Subscription

Hybrid Cloud TME Onprem

## Location

South Central US

## Resource Group

☐ Create New ☒ Use Existing

## Resource Group

ANFAVSRG

## Authentication Method

☐ Password ☒ Public Key

## User Name

azureuser

## Enter SSH Public Key

-----BEGIN RSA PRIVATE KEY----- MIIGSAIBAAKCA...

[Previous](#)[Next](#)

6. Provide a name for the connector instance, select **Create** and accept default **Role Name** under **Details**, and choose the subscription for the Azure account.

Add BlueXP Connector - Azure
More Information

VM Authentication
2 Details
3 Network
4 Security Group
5 Review

### Details

Connector Instance Name
AzureConnector

Add Tags to Connector Instance

#### Connector Role

☒ Create
☐ Attach existing
☐ Manual

Role Name
BlueXP Operator-5519248

Subscriptions to apply with the role
Hybrid Cloud TME Onprem

Previous
Next

- Configure networking with the proper **VNet**, **Subnet**, and disable **Public IP** but ensure that the connector has the internet access in your Azure environment.

Add BlueXP Connector - Azure
More Information

VM Authentication
Details
3 Network
4 Security Group
5 Review

### Network

#### Connectivity

VNet
ANFAVSVal

Subnet
VM\_Sub

Public IP
Disable

**Notice:** Ensure that the subnet has internet connectivity through a NAT device or proxy server so that the Connector can communicate with Azure services.

#### Proxy Configuration (Optional)

HTTP Proxy
Example: http://172.16.254.1:8080

Define Credentials for this Proxy

Upload a root certificate

Previous
Next

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

**Add BlueXP Connector - Azure** More Information ×

✓ VM Authentication   ✓ Details   ✓ Network   **4** Security Group   5 Review

### Security Group

The security group must allow inbound HTTP, HTTPS and SSH access.

Assign a security group: ☒ Create a new security group   ☐ Select an existing security group

<b>HTTP</b> (Port 80)	<b>HTTPS</b> (Port 443)	<b>SSH</b> (Port 22)
Source Type <input type="text" value="Anywhere"/>	Source Type <input type="text" value="Anywhere"/>	Source Type <input type="text" value="Anywhere"/>
Source (CIDR) <input type="text" value="0.0.0.0/0"/>	Source (CIDR) <input type="text" value="0.0.0.0/0"/>	Source (CIDR) <input type="text" value="0.0.0.0/0"/>

Previous Next 📄

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 Connector - Azure

More Information

VM Authentication

Details

Network

Security Group

5 Review

Review

Code for Terraform Automation

BlueXP Connector Name	AzureConnector
Subscription	Hybrid Cloud TME Onprem
Location	South Central US
Resource Group	Existing - ANFAVSRG
Role	New - BlueXP Operator-5519248
Authentication Method	Password (user: azureuser)
VNet	ANFAVSVAl
Subnet	VM_Sub
Public IP	Enable
Proxy	None
Security Group	HTTP: 0.0.0.0/0, HTTPS: 0.0.0.0/0, SSH: 0.0.0.0/0

Previous

Add

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

NetApp BlueXP

Q BlueXP Search

Account Automation-to...

Workspace Azure-DB

Connector AzureConnector

2

?

Canvas

My working environments

My estate

+ Add Working Environment

Enable Services

Azure Blob Storage  
20 Storage Accounts

Amazon S3  
0 Buckets

Working Environments


Amazon S3  
0 Buckets

Azure Blob Storage  
20 Storage Accounts




**Define a credential in BlueXP for Azure resources access**

1. Click on setting icon on top right corner of BlueXP console to open **Account credentials** page, click **Add credentials** to start credential configuration workflow.



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



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

**NetApp BlueXP** Q BlueXP Search Account Automation-team Workspace Azure-DB Connector AzureConnector 2 ? !

**Add Credentials** 1 Credentials Type 2 Define Credentials 3 Marketplace Subscription 4 Review X

**Define Microsoft Azure Credentials**  
Learn more about Azure application credentials

Credentials Name ? Client Secret

Azure\_Hybrid\_TME .....

Application (client) ID Directory (tenant) ID

2fbc9be5-a259-4539-bb57-036b176f5cc7 9bb0aab6-5c98-419b-9cfd-7a38bd496...

☒ I have verified that the Azure role assigned to the Active Directory service principal matches BlueXP policy requirements.

Previous Next !

#### 4. Review and Add.

**NetApp BlueXP** Q BlueXP Search Account Automation-team Workspace Azure-DB Connector AzureConnector 2 ? !


**Add Credentials** 1 Credentials Type 2 Define Credentials 3 Review X

**Review**

Credentials Type	Azure
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

Previous Add !

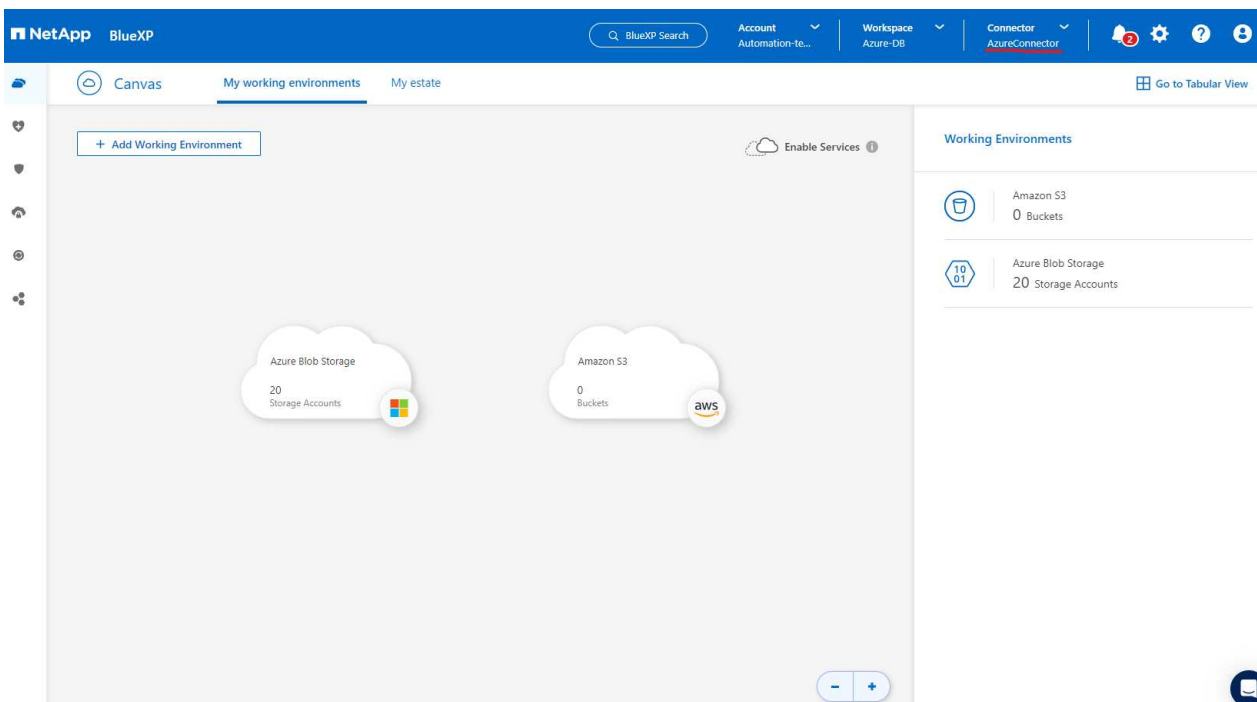
5. You may also need to associate a **Marketplace Subscription** with the credential.




## SnapCenter services setup

With the Azure credential configured, SnapCenter services can now be set up with the following procedures:

1. Back to Canvas page, from **My Working Environment** click **Add working Environment** to discover Azure NetApp Files deployed in Azure.



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



3. Name **Working Environment** and choose **Credential Name** created in previous section, and click **Continue**.

NetApp BlueXP

Q BlueXP Search Account Automation-team Workspace Azure-DB Connector AzureConnector

Add Azure NetApp Files Wizard Azure NetApp Files Details

Back

Azure NetApp Files Credentials

Working Environment Name  
AzureNFile

Select the credentials that provides BlueXP with the permissions that it needs to manage ANF.

Credentials Name  
Azure\_nfile

To add a new set of credential, go to the Credentials Page.

How to get the credentials

- Create an Azure AD client that meets the required permissions
- Enter credentials on credentials Page:
  - Application (client) ID** - The ID of an Active Directory service principal that BlueXP can use to authenticate with Azure Active Directory.
  - Client Secret** - The value of a client secret for the service principal application.
  - Directory (tenant) ID** - The Active Directory ID for your organization. This is sometimes referred to as a Tenant ID.

To learn more about Azure NetApp Files, [click here](#)

Continue

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

NetApp BlueXP

Q BlueXP Search Account Automation-te... Workspace Azure-DB Connector AzureConnector

Canvas My working environments My estate Go to Tabular View

+ Add Working Environment

Enable Services

AzureNFile Azure NetApp Files  
16 Volumes 7.08 TiB Capacity


Amazon S3  
0 Buckets

Azure Blob Storage  
20 Storage Accounts


Working Environments

- 1 Azure NetApp Files  
7.08 TiB Provisioned Capacity
- Amazon S3  
0 Buckets
- Azure Blob Storage  
20 Storage Accounts

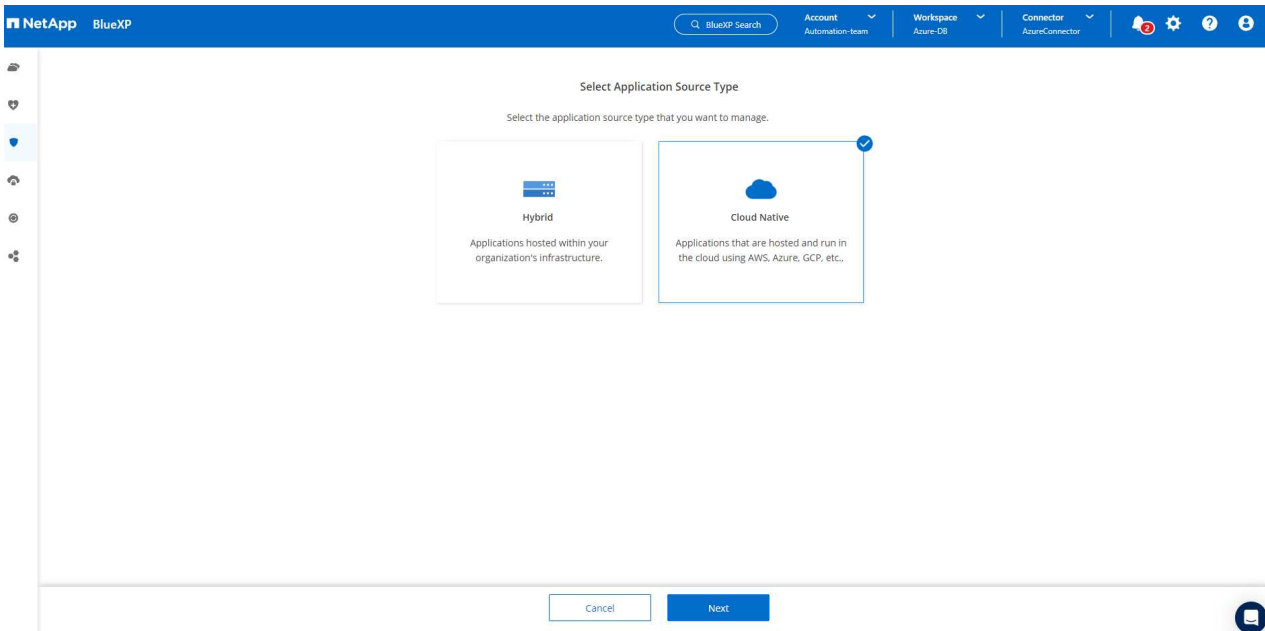
5. Click on **Azure NetApp Files** icon, then **Enter Working Environment** to view Oracle database volumes deployed in Azure NetApp Files storage.



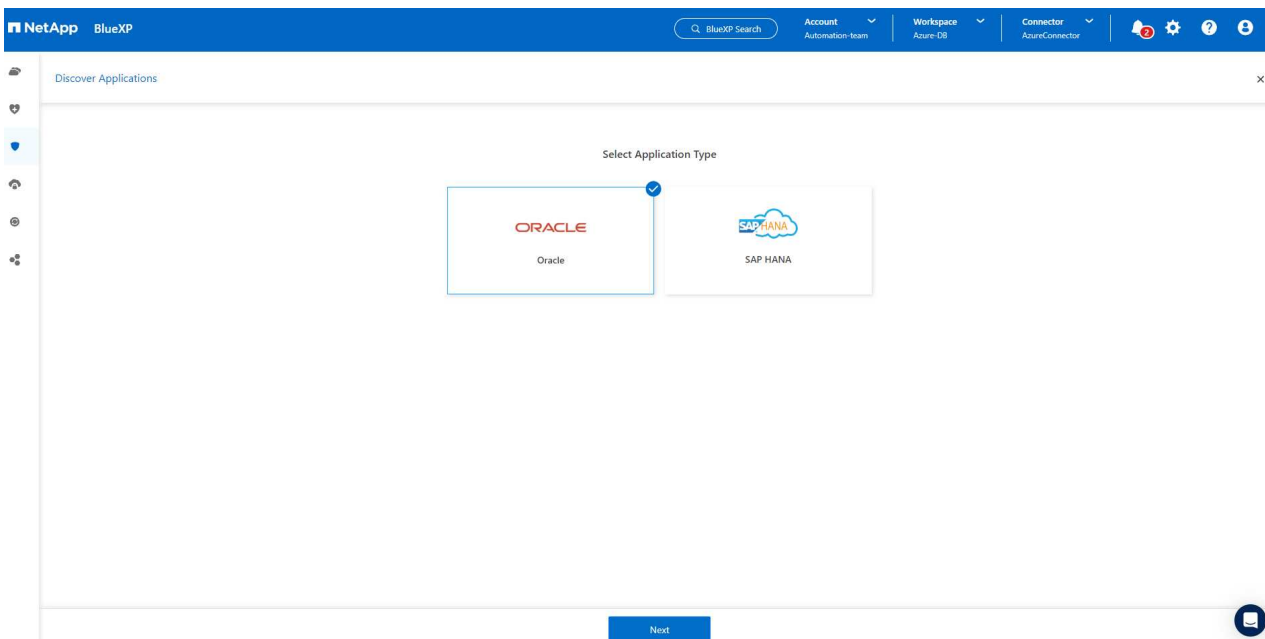
- From the left-hand sidebar of the console, hover your mouse over the protection icon, and then click **Protection > Applications** to open the Applications launch page. Click **Discover Applications**.



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



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



9. Select **Using SSH** and provide the Oracle Azure VM details such as **IP address**, **Connector**, Azure VM management **Username** such as azureuser. Click on **Add SSH Private Key** to paste in the SSH key pair that you used to deploy the Oracle Azure VM. You will also be prompted to confirm the fingerprint.



NetApp BlueXP

BlueXP Search

AccountAutomation team

WorkspaceAzure DB

ConnectorAzureConnector

Discover Applications

1 Host Details

2 Configuration

3 Review

Select host type

Provide the following details to add host and discover applications

Host Installation Type

☐ Manual

☒ Using SSH

Host FQDN or IP

172.30.137.142

Connector

AzureConnector

Username

azureuser

SSH Port

22

Plug-in Port

8145

Add SSH Private Key Optional

Previous

Next

Discover Applications

1 Host Details

2 Configuration

3 Review

Select host type

Provide the following details to add host and discover applications

Host Installation Type

☐ Manual

☒ Using SSH

Validate fingerprint

Algorithm

ssh-rsa

Fingerprint

AAAAE2VjZHNhLXNoYTItbmlzdHAyNTYAAAIbmlzdHAyNTYAAAB...

☒ By proceeding further, I confirm that the above fingerprint for host is valid.


Proceed

Cancel


Previous

Next

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



11. Review and click on **Discover Applications** to install a plugin on the Oracle Azure VM and discover Oracle database on the VM in one step.



12. Discovered Oracle databases on Azure VM are added to **Applications**, and the **Applications** page lists the number of hosts and Oracle databases within the environment. The database **Protection Status** initially shows as **Unprotected**.

NetApp
BlueXP

BlueXP Search

Account Automation-te...

Workspace Azure-DB

Connector AzureConnector

Backup and recovery
Volumes
Restore
Applications
Virtual Machines
Kubernetes
Job Monitoring
Reports

Cloud Native

Oracle

3

Hosts

3

ORACLE

0

Clone

Application Protection

0

Protected

3

Unprotected

3 Databases

Filter By

Manage Databases

Settings

Name	Host Name	Policy Name	Protection Status
NTAP	172.30.137.142		Unprotected
db1	172.30.15.99		Unprotected
db1test	172.30.15.124		Unprotected

1 - 3 of 3

This completes the initial setup of SnapCenter services for Oracle. The next three sections of this document describe Oracle database backup, restore, and clone operations.

### Oracle database backup

61

1. Our test Oracle database in Azure VM is configured with three volumes with an aggregate total storage about 1.6 TiB. This gives context about the timing for the snapshot backup, restore, and clone of a database of this size.

```
[oracle@acao-ora01 ~]$ df -h
Filesystem                Size      Used Avail Use% Mounted on
devtmpfs                  7.9G         0  7.9G   0% /dev
tmpfs                     7.9G         0  7.9G   0% /dev/shm
tmpfs                     7.9G      17M  7.9G   1% /run
tmpfs                     7.9G         0  7.9G   0% /sys/fs/cgroup
/dev/mapper/rootvg-rootlv 40G       23G    15G  62% /
/dev/mapper/rootvg-usrlv  9.8G     1.6G    7.7G  18% /usr
/dev/sda2                 496M     115M   381M  24% /boot
/dev/mapper/rootvg-varlv  7.9G     787M    6.7G  11% /var
/dev/mapper/rootvg-homelv 976M     323M    586M  36% /home
/dev/mapper/rootvg-optlv  2.0G     9.6M    1.8G   1% /opt
/dev/mapper/rootvg-tmplv  2.0G      22M    1.8G   2% /tmp
/dev/sda1                 500M     6.8M   493M   2% /boot/efi
172.30.136.68:/ora01-u01 100G      23G     78G  23% /u01
172.30.136.68:/ora01-u03 500G     117G    384G  24% /u03
172.30.136.68:/ora01-u02 1000G     804G    197G  81% /u02
tmpfs                     1.6G         0  1.6G   0% /run/user/1000
[oracle@acao-ora01 ~]$
```

1. To protect database, click the three dots next to the database **Protection Status**, and then click **Assign Policy** to view the default preloaded or user defined database protection policies that can be applied to your Oracle databases. Under **Settings - Policies**, you have option to create your own policy with a customized backup frequency and backup data-retention window.

The screenshot shows the NetApp BlueXP interface. The top navigation bar includes 'Backup and recovery', 'Volumes', 'Restore', 'Applications' (selected), 'Virtual Machines', 'Kubernetes', 'Job Monitoring', and 'Reports'. Below the navigation bar, there are filters for 'Cloud Native' and 'Oracle'. A summary section shows '4 Hosts', '3 ORACLE', and '0 Clone'. An 'Application Protection' summary shows '0 Protected' and '3 Unprotected'. Below this, a table lists databases with their protection status. A dropdown menu for the 'NTAP' database is open, showing 'View Details' and 'Assign Policy' (highlighted with a red box).

Name	Host Name	Policy Name	Protection Status
NTAP	172.30.137.142		Unprotected
db1	172.30.15.99		Unprotected
db1tst	172.30.15.124		Unprotected

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

The screenshot shows the 'Assign Policy' dialog in the NetApp BlueXP interface. The title is 'Assign Policy' with a subtitle 'Assign a policy to start taking backups of the database "NTAP"'. Below the title, there are 4 policies listed in a table. The 'my\_full\_bkup' policy is selected with a blue checkmark. At the bottom, there are 'Cancel' and 'Assign' buttons.

Policy Name	Backup Type	Schedules
<input type="radio"/> 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
<input type="radio"/> 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, C
<input type="radio"/> 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, Jul, Aug, Sep, C
<input checked="" type="radio"/> my_full_bkup	FullBackup	Hourly: Repeats Every 6 Hrs, Keeps 3 Days

- 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.

The screenshot shows the NetApp BlueXP interface with the 'Applications' tab selected. At the top, there are filters for 'Cloud Native' and 'Oracle'. Below these, there are three cards: '3 Hosts', '3 ORACLE', and '0 Clone'. To the right, an 'Application Protection' summary shows '1 Protected' and '2 Unprotected' databases. The main section is titled '3 Databases' and includes a 'Filter By' button and a search bar. A table lists the databases:

Name	Host Name	Policy Name	Protection Status
NTAP	172.30.137.142	my_full_bkup	Protected
db1	172.30.15.99		Unprotected
db1tst	172.30.15.124		Unprotected

A three-dot menu is open for the 'NTAP' database, showing options: 'View Details', 'On-Demand Backup' (highlighted with a red underline), 'Assign Policy', 'Un-assign Policy', and 'Restore'.

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.

The screenshot shows the NetApp BlueXP interface with the 'Job Monitoring' tab selected. The breadcrumb trail is 'Job Monitoring > Job Name: Backup of NTAP oracle database on host 172.30.137.142 with policy my\_full\_bkup and schedule Hourly'. The job name is 'Job Name: Backup of NTAP oracle database on host 172.30.137.142 with policy my\_full\_bkup and schedule H...' and the job ID is '61a12139-330e-4390-bca8-e7d15680869c'. Below this, there are four cards: 'Other Job Type', 'Jul 11 2023, 2:17:53 pm Start Time', 'Jul 11 2023, 2:21:38 pm End Time', and 'Success Job Status'. The 'Sub-Jobs(17)' section shows a table of sub-jobs:

Job Name	Job ID	Start Time	End Time	Duration
Backup of NTAP oracle database on host 172.30...	61a12139-330e-4390-bc...	Jul 11 2023, 2:17:53 pm	Jul 11 2023, 2:21:38 pm	4 Minutes
Applying Retention	27ff9d5f-68f0-4880-a48...	Jul 11 2023, 2:21:38 pm	Jul 11 2023, 2:21:38 pm	0 Second
Performing cleanup after backup	074c0689-097e-41aa-ac...	Jul 11 2023, 2:21:36 pm	Jul 11 2023, 2:21:38 pm	2 Seconds
Finalizing Oracle database log backup	348189d3-90b5-4cce-97...	Jul 11 2023, 2:21:36 pm	Jul 11 2023, 2:21:36 pm	0 Second

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

The screenshot shows the NetApp BlueXP interface. The top navigation bar includes 'Backup and recovery', 'Volumes', 'Restore', 'Applications' (selected), 'Virtual Machines', 'Kubernetes', 'Job Monitoring', and 'Reports'. The 'Applications' section is active, showing filters for 'Cloud Native' and 'Oracle'. Summary cards indicate 4 Hosts, 3 ORACLE, and 0 Clone. An 'Application Protection' card shows 2 Protected and 1 Unprotected databases. Below, a table lists databases with their protection status. A context menu is open for the 'db1tst' database, showing options like 'View Details', 'On-Demand Backup', 'Assign Policy', 'Un-assign Policy', and 'Restore'.

Name	Host Name	Policy Name	Protection Status
NTAP	172.30.137.142	my_full_bkup	Protected
db1	172.30.15.99	my_full_bkup	Protected
db1tst	172.30.15.124		Unprotected

- 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.

The screenshot shows the 'Database Details' page for the 'NTAP' database. It displays various attributes in a grid: Database Name (NTAP), Protection (Protected), Policy Names (my\_full\_bkup), Database Type (zEHlu7vkdyabnucxllbkKELkVXTyNcllients), Host Name (172.30.137.142), Host Storage (ANF), Database Version (Unreachable), Connector Id (zEHlu7vkdyabnucxllbkKELkVXTyNcllients), Clones (-), Parent Database (-), RMAN Catalog (Disabled), and RMAN catalog repository (-). Below this, a 'Backups' section shows a list of 14 backups. A table displays the first four backups with columns for Backup Name, Backup Type, SCN, RMAN Catalog, Backup Time, and a Delete link.

Backup Name	Backup Type	SCN	RMAN Catalog	Backup Time	
my_full_bkup_Hourly_NTAP_2023_07_13_12_04_28_8376...	Log	29192187	Not Cataloged	Jul 13, 2023, 8:06:22 am	Delete
my_full_bkup_Hourly_NTAP_2023_07_13_12_03_07_4363...	Data	29192136	Not Cataloged	Jul 13, 2023, 8:03:40 am	Delete
my_full_bkup_Hourly_NTAP_2023_07_13_06_04_28_5618...	Log	29178022	Not Cataloged	Jul 13, 2023, 2:05:50 am	Delete
my_full_bkup_Hourly_NTAP_2023_07_13_06_03_03_6371...	Data	29177972	Not Cataloged	Jul 13, 2023, 2:03:43 am	Delete

## 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.

The screenshot shows the NetApp BlueXP interface. The top navigation bar includes 'NetApp BlueXP', a search bar, and links for 'Account', 'Workspace', and 'Connector'. The left sidebar lists various categories like Storage, Health, Protection, Governance, Mobility, and Extensions. The main content area is titled 'Applications' and shows a summary of resources: 4 Cloud Native Hosts, 3 ORACLE databases, and 0 Clones. Below this, there's a section for '3 Databases' with a table listing them. A context menu is open for the 'db1tst' database, showing options like 'View Details', 'On-Demand Backup', 'Assign Policy', 'Un-assign Policy', and 'Restore'.

Name	Host Name	Policy Name	Protection Status
NTAP	172.30.137.142	my_full_bkup	Protected
db1	172.30.15.99	my_full_bkup	Protected
db1tst	172.30.15.124		Unprotected

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

The screenshot shows the 'Restore "NTAP"' dialog box in the NetApp BlueXP interface. The dialog has three steps: '1 Restore Point and Location', '2 Configuration', and '3 Review'. The current step is 'Restore Point and Location', which prompts the user to 'Specify the restore point to which the database should be restored.' A dropdown menu for 'Restore Point' is open, showing a list of timestamps: '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', and 'Jul 12, 2023, 2:03:31 am'. Below the dropdown, there are two 'location' fields. At the bottom of the dialog, there are 'Previous' and 'Next' buttons.

3. Choose your **Restore Location** to **original location** for an Oracle database in place restore and recovery.



NetApp BlueXP

Restore "NTAP"

1 Restore Point and Location 2 Configuration 3 Review

Restore Point and Location

Specify the restore point to which the database should be restored.

Restore Point  
Jul 13, 2023, 8:03:40 am

Restore to original location

Restore to alternate location

Previous Next

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

NetApp BlueXP

Restore "NTAP"

1 Restore Point and Location 2 Configuration 3 Review

Restore Scope

☒ All Data Files  
Data Files Restore

☒ Control Files  
Control Files Restore

Database state will be changed if needed for restore and recovery.

Recovery Scope


☒ All Logs ☐ Until System Change Number ☐ Date and Time ☐ No 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.



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



## Oracle database clone

Database clone procedures are similar to restore but to an alternate Azure VM with identical Oracle software stack pre-installed and configured.



Ensure that your Azure NetApp File storage has sufficient capacity for a cloned database the same size as the primary database to be cloned. The alternate Azure VM has been added to **Applications**.

1. Click the three-dot drop down menu for the particular database to be cloned in **Applications**, then click **Restore** to initiate clone workflow.

The screenshot shows the NetApp BlueXP interface. The top navigation bar includes 'NetApp BlueXP', a search bar, and links for 'Account', 'Workspace', and 'Connector'. The left sidebar shows various categories like 'Storage', 'Health', 'Protection', 'Governance', 'Mobility', and 'Extensions'. The main content area is titled 'Applications' and displays a summary of resources: 4 Cloud Native Hosts, 3 ORACLE, and 0 Clones. Below this, there's a section for '3 Databases' with a table listing them. A context menu is open for the 'db1tst' database, showing options like 'View Details', 'On-Demand Backup', 'Assign Policy', 'Un-assign Policy', and 'Restore' (which is highlighted).

Name	Host Name	Policy Name	Protection Status
NTAP	172.30.137.142	my_full_bkup	Protected
db1	172.30.15.99	my_full_bkup	Protected
db1tst	172.30.15.124		Unprotected

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

The screenshot shows the 'Restore "NTAP"' configuration page in NetApp BlueXP. The page has three steps: '1 Restore Point and Location', '2 Configuration', and '3 Review'. The current step is 'Restore Point and Location', which prompts the user to 'Specify the restore point to which the database should be restored.' A dropdown menu for 'Restore Point' is set to 'Jul 13, 2023, 8:03:40 am'. Below this, there are two options: 'Restore to original location' (with a database icon) and 'Restore to alternate location' (with a database icon and a checkmark). At the bottom, there are 'Previous' and 'Next' buttons.

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

The screenshot shows the 'Configuration' step in the 'Restore "NTAP"' workflow. The page has a blue header with 'NetApp BlueXP' and navigation links. The main content area is titled 'Configuration' and includes a sub-header 'Specify the alternate host details on which the database will be restored and throughput.' Below this is a form with the following fields:

- Host:** A dropdown menu showing '172.30.137.147'.
- SID:** A text input field containing 'NTAP1'.
- Oracle Home:** A text input field containing '/u01/app/oracle/product/19.0.0/clone'.
- Database Credentials:** A section labeled 'Optional' with an 'Add Credential' button.
- Maximum storage throughput (MiB/s):** A section labeled 'Optional' with a text input field containing 'Enter throughput (1-4500)'.

At the bottom of the form are 'Previous' and 'Next' buttons.

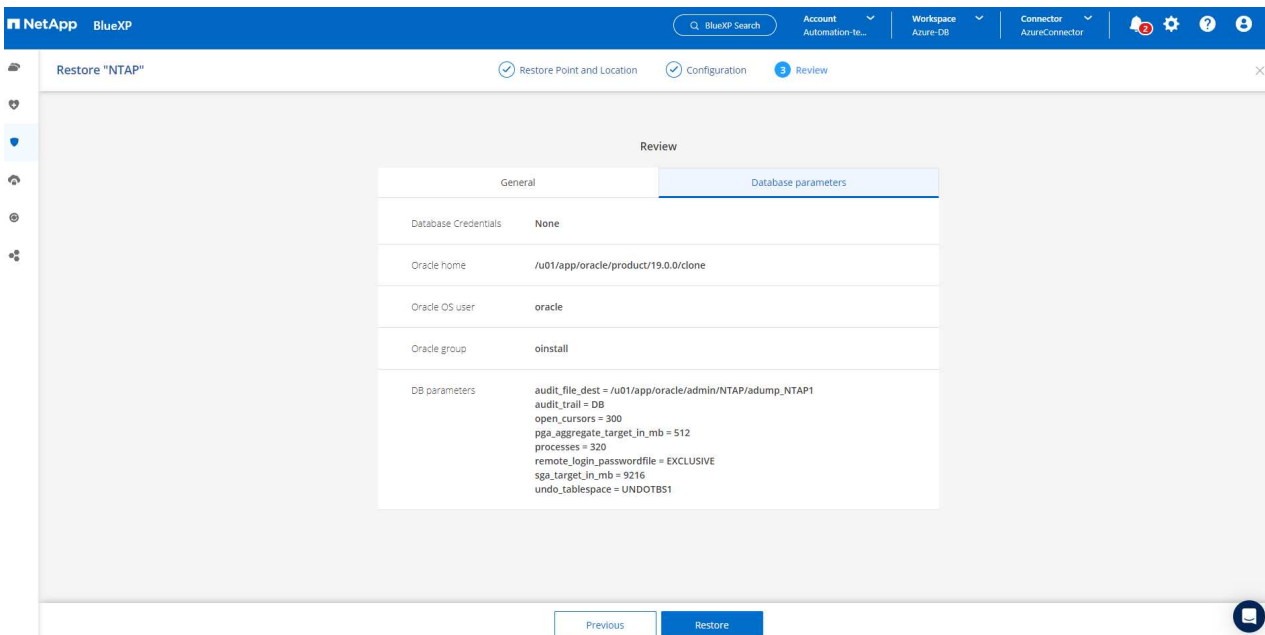
4. Review **General** page shows the details of cloned database such as SID, alternate host, data file locations, recovery scope etc.

The screenshot shows the 'Review' step in the 'Restore "NTAP"' workflow. The page has a blue header with 'NetApp BlueXP' and navigation links. The main content area is titled 'Review' and contains a table with two tabs: 'General' (selected) and 'Database parameters'.


General	
Backup Name	my_full_bkup_Hourly_NTAP_2023_07_13_12_03_07_43633_0
SID	NTAP1
Host	172.30.137.147
Datafile locations	/u02_NTAP1
Control files	/u02_NTAP1/NTAP1/control/control01.ctl
Redo logs	RedoGroup = 1 TotalSize = 1024 Path = /u02_NTAP1/NTAP1/redolog/redo01_01.log RedoGroup = 2 TotalSize = 1024 Path = /u02_NTAP1/NTAP1/redolog/redo02_01.log RedoGroup = 3 TotalSize = 1024 Path = /u02_NTAP1/NTAP1/redolog/redo03_01.log
Recovery scope	Until cancel using selected backup's archive logs
Recovery Point	Jul 13, 2023, 8:03:40 am
Location	Alternate Location

At the bottom of the table are 'Previous' and 'Restore' buttons.

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



6. Monitor the cloning job status from the **Job Monitoring** tab, we observed that it took 8 minutes to clone a 1.6 TiB Oracle database.



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

NetApp BlueXP

BlueXP Search

Account Automation-te...

Workspace Azure-DB

Connector AzureConnector

3

⚙

?

👤

Backup and recovery
Volumes
Restore
Applications
Virtual Machines
Kubernetes
Job Monitoring
Reports

Cloud Native

Oracle

Cloud

4

Hosts

ORACLE

4

ORACLE

Clone

0

Clone

Application Protection

2 Protected

2 Unprotected

4 Databases
Filter By +

Manage Databases

Settings

Name	Host Name	Policy Name	Protection Status	
NTAP	172.30.137.142	my_full_bkup	Protected	...
NTAP1	172.30.137.147		Unprotected	...
db1	172.30.15.99	my_full_bkup	Protected	...
db1tst	172.30.15.124		Unprotected	...

1 - 4 of 4

<<
<
1
>
>>

8. Validate the cloned database on the Oracle Azure VM that showed the cloned database was running as expected.

```

[oracle@acao-ora02 admin]$ cat /etc/oratab
#

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

# A colon, ':', is used as the field terminator.  A new line terminates
# the entry.  Lines beginning with a pound sign, '#', are comments.
#
# Entries are of the form:
#   $ORACLE_SID:$ORACLE_HOME:<N|Y>:
#
# The first and second fields are the system identifier and home
# directory of the database respectively.  The third field indicates
# to the dbstart utility that the database should , "Y", or should not,
# "N", be brought up at system boot time.
#
# Multiple entries with the same $ORACLE_SID are not allowed.
#
#
# SnapCenter Plug-in for Oracle Database generated entry (DO NOT REMOVE THIS LINE)
NTAPI:/u01/app/oracle/product/19.0.0/clone:N
[oracle@acao-ora02 admin]$ export ORACLE_SID=NTAPI
[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
-----
NTAPI         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

Allen Cao, Niyaz Mohamed, NetApp

### Purpose

SnapCenter Services is the SaaS version of the classic SnapCenter database management UI tool that is available through the NetApp BlueXP cloud management console. It is an integral part of the NetApp cloud-backup, data-protection offering for databases such as Oracle and HANA running on NetApp cloud storage. This SaaS-based service simplifies traditional SnapCenter standalone server deployment that generally requires a Windows server operating in a Windows domain environment.

In this documentation, we demonstrate how you can set up SnapCenter Services to backup, restore, and clone Oracle databases deployed to Amazon FSx for ONTAP storage and EC2 compute instances. Although it is much easier to set up and use, SnapCenter Services deliver key functionalities that are available in the legacy SnapCenter UI tool.

This solution addresses the following use cases:

- Database backup with snapshots for Oracle databases hosted in Amazon FSx for ONTAP
- Oracle database recovery in the case of a failure
- Fast and storage-efficient cloning of primary databases for a dev/test environment or other use cases

### Audience

This solution is intended for the following audiences:

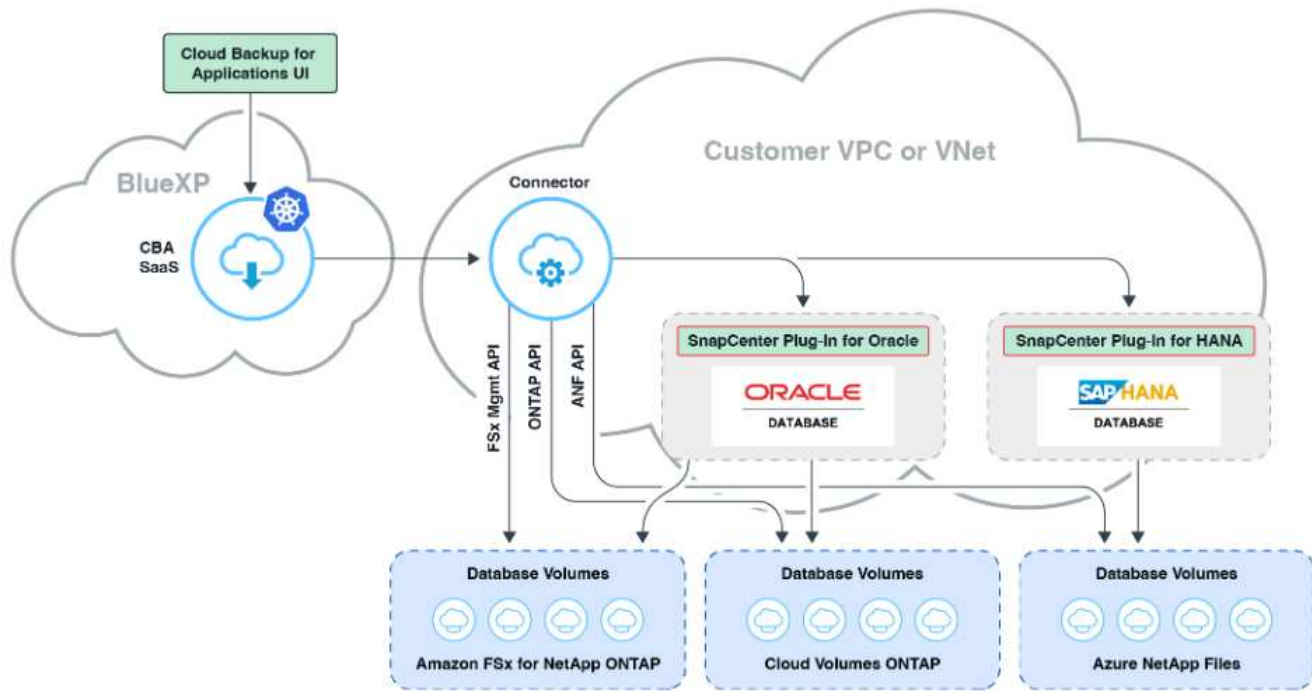
- The DBA who manages Oracle databases running on Amazon FSx for ONTAP storage
- The solution architect who is interested in testing Oracle database backup, restore, and clone in the public AWS cloud
- The storage administrator who supports and manages the Amazon FSx for ONTAP storage
- The application owner who owns applications that are deployed to Amazon FSx for ONTAP storage

### Solution test and validation environment

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

### Architecture





This image provides a detailed picture of BlueXP backup and recovery for applications within the BlueXP console, including the UI, the connector, and the resources it manages.

## Hardware and software components

### Hardware

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

### Software

RedHat Linux	RHEL-8.6.0_HVM-20220503-x86_64-2-Hourly2-GP2	Deployed RedHat subscription for testing
Oracle Grid Infrastructure	Version 19.18	Applied RU patch p34762026_190000_Linux-x86-64.zip
Oracle Database	Version 19.18	Applied RU patch p34765931_190000_Linux-x86-64.zip
Oracle OPatch	Version 12.2.0.1.36	Latest patch p6880880_190000_Linux-x86-64.zip
SnapCenter Service	Version	v2.3.1.2324

## Key factors for deployment consideration

- **Connector to be deployed in the same VPC as database and FSx.** When possible, the connector should be deployed in the same AWS VPC, which enables connectivity to the FSx storage and the EC2 compute instance.
- **An AWS IAM policy created for SnapCenter connector.** The policy in JSON format is available in the detailed SnapCenter service documentation. When you launch connector deployment with the BlueXP console, you are also prompted to set up the prerequisites with details of required permission in JSON format. The policy should be assigned to the AWS user account that owns the connector.
- **The AWS account access key and the SSH key pair created in the AWS account.** The SSH key pair is assigned to the ec2-user for logging into the connector host and then deploying a database plug-in to the EC2 DB server host. The access key grants permission for provisioning the required connector with IAM policy above.
- **A credential added to the BlueXP console setting.** To add Amazon FSx for ONTAP to the BlueXP working environment, a credential that grants BlueXP permissions to access Amazon FSx for ONTAP is set up in the BlueXP console setting.
- **java-11-openjdk installed on the EC2 database instance host.** SnapCenter service installation requires java version 11. It needs to be installed on application host before plugin deployment attempt.

## Solution deployment

There is extensive NetApp documentation with a broader scope to help you protect your cloud-native application data. The goal of this documentation is to provide step-by-step procedures that cover SnapCenter Service deployment with the BlueXP console to protect your Oracle database deployed to Amazon FSx for ONTAP and an EC2 compute instance. This document fills in certain details that might be missing from more general instructions.

To get started, complete the following steps:

- Read the general instructions [Protect your cloud native applications data](#) and the sections related to Oracle and Amazon FSx for ONTAP.
- Watch the following video walkthrough.

## Solution Deployment

### Prerequisites for SnapCenter service deployment

Deployment requires the following prerequisites.

1. A primary Oracle database server on an EC2 instance with an Oracle database fully deployed and running.
2. An Amazon FSx for ONTAP cluster deployed in AWS that is hosting the database volumes above.
3. An optional database server on an EC2 instance that can be used for testing the cloning of an Oracle database to an alternate host for the purpose of supporting a dev/test workload or any use cases that requires a full data set of a production Oracle database.
4. If you need help to meet the above prerequisites for Oracle database deployment on Amazon FSx for ONTAP and EC2 compute instance, see [Oracle Database Deployment and Protection in AWS FSx/EC2 with iSCSI/ASM](#) or white paper [Oracle Database Deployment on EC2 and FSx Best Practices](#)

## Onboarding to BlueXP preparation

1. Use the link [NetApp BlueXP](#) to sign up for BlueXP console access.
2. Login to your AWS account to create an IAM policy with proper permissions and assign the policy to the AWS account that will be used for BlueXP connector deployment.


The screenshot displays the AWS IAM console interface. The left-hand navigation pane is titled 'Identity and Access Management (IAM)' and includes sections for 'Access management' (User groups, Users, Roles), 'Policies' (Identity providers, Account settings), 'Access reports' (Access analyzer, Archive rules, Analyzers, Settings, Credential report, Organization activity, Service control policies (SCPs)), and a search bar for IAM. The main content area is titled 'Summary' for a policy named 'snapcenter'. It shows the 'Policy ARN' as 'arn:aws:iam::541696183547:policy/snapcenter' and the 'Description' as 'Policy to grant snapcenter service permission to create connector in AWS.' Below this, there are tabs for 'Permissions', 'Policy usage', 'Tags', 'Policy versions', and 'Access Advisor'. The 'Permissions' tab is active, displaying a 'Policy summary' and a 'JSON' button. The JSON string for the policy is shown in a text area, starting with '{' and containing a 'Version' of '2012-10-17' and a 'Statement' array with a single action allowing various IAM and EC2 operations. The bottom of the sidebar shows the 'AWS account ID: 541696183547'.

The policy should be configured with a JSON string that is available in NetApp documentation. The JSON string can also be retrieved from the page when connector provisioning is launched and you are prompted for the prerequisites permissions assignment.


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

## Deploy a connector for SnapCenter services


1. Login to the BlueXP console. For a shared account, it is a best practice to create an individual workspace by clicking **Account > Manage Account > Workspace** to add a new workspace.



2. Click **Add a Connector** to launch the connector provisioning workflow.



1. Choose your cloud provider (in this case, **Amazon Web Services**).



1. Skip the **Permission**, **Authentication**, and **Networking** steps if you already have them set up in your AWS account. If not, you must configure these before proceeding. From here, you could also retrieve the permissions for the AWS policy that is referenced in the previous section "[Onboarding to BlueXP preparation](#)."

## Deploying a Connector

The Connector is a crucial component for the day-to-day use of Cloud Manager.  
It's used to connect Cloud Manager's services to your hybrid-cloud environments.  
The Connector can then manage the resources and processes within your public cloud environment.

Before you begin the deployment process, ensure that you have completed the required preparations. This guide will enable you to focus on the minimum requirements for Connector installation.

### Permissions

Set up an IAM role with the required permissions

### Authentication

Choose between two AWS authentication methods: AWS keys or assuming an IAM role

### Networking

Obtain details about the VPC and subnet in which the Connector will reside

[Skip to Deployment](#)

[Previous](#)

[Continue](#)



1. Enter your AWS account authentication with **Access Key** and **Secret Key**.

- 1 AWS Credentials
- 2 Details
- 3 Network
- 4 Security Group
- 5 Review

## AWS Authentication

Region

us-east-1 | US East (N. Virginia)

Select the Authentication Method: ☐ Assume Role ☒ AWS Keys

AWS Access Key

AKIA6JRXA6ZVGVFSHMO3

AWS Secret Key

.....

Want to launch an instance without AWS Credentials? [▼](#)

[Previous](#)

[Next](#)



2. Name the connector instance and select **Create Role** under **Details**.

**Add Connector - AWS**More Information ×

✓ AWS Credentials 2 Details 3 Network 4 Security Group 5 Review

### Details

Connector Instance Name

SnapCenterSvs

+

 Add Tags to Connector Instance

Connector Role

☒ Create Role ☐ Select an existing Role

Role Name

Cloud-Manager-Operator-VZzSSP9-SnapCenter

☐ AWS Managed Encryption

Master Key: aws/ebs (default) [Change Key](#)

Previous

Next

1. Configure networking with the proper **VPC**, **Subnet**, and SSH **Key Pair** for connector access.

Add BlueXP Connector - AWS

More Information

✓ AWS Credentials

✓ Details

3 Network

4 Security Group

5 Review

Network

Connectivity

VPC

vpc-0b522d5e982a50ceb - 172.30.15.0/25

Subnet

172.30.15.0/25 | priv-subnet-01

Key Pair

sufi\_new

Public IP

Use subnet settings (Disable)

Notice:

Ensure that the subnet has internet connectivity through a NAT device or proxy server so that the Connector can communicate with AWS services.

Proxy Configuration (Optional)

HTTP Proxy

Example: http://172.16.254.1:8080

Define Credentials for this Proxy

Upload a root certificate

Previous

Next

2. Set the **Security Group** for the connector.

Add BlueXP Connector - AWS

More Information

✓ AWS Credentials

✓ Details

✓ Network

4 Security Group

5 Review

Security Group

The security group must allow inbound HTTP, HTTPS and SSH access.

Assign a security group: ☐ Create a new security group ☒ Select an existing security group

1 Security Group

Security Group Name	Description
<input checked="" type="radio"/> default	default VPC security group

Previous

Next



3. Review the summary page and click **Add** to start connector creation. It generally takes about 10 mins to complete deployment. Once completed, the connector instance appears in the AWS EC2 dashboard.

Add BlueXP Connector - AWS

More Information

✓ AWS Credentials

✓ Details

✓ Network

✓ Security Group

5 Review

Review

[Code for Terraform Automation](#)

BlueXP Connector Name	aws-snapctr-us-east
AWS Access Key	AKIA4H43ZT56IWWR3TI
Region	us-east-1
VPC	vpc-0b522d5e982a50ceb - 172.30.15.0/25
Subnet	172.30.15.0/25   priv-subnet-01
Key Pair	sufi_new
Public IP	Use subnet settings (Disable)
Proxy	None
Security Group	default

Previous

Add

Define a credential in BlueXP for AWS resources access

1. First, from AWS EC2 console, create a role in **Identity and Access Management (IAM)** menu **Roles**, **Create role** to start role creation workflow.

The screenshot shows the AWS IAM console 'Roles' page. The left sidebar contains the 'Identity and Access Management (IAM)' menu with 'Roles' selected. The main content area displays a table of roles with columns for 'Role name', 'Trusted entities', and 'Last activity'. The 'Create role' button is located in the top right corner.

2. In **Select trusted entity** page, choose **AWS account**, **Another AWS account**, and paste in the BlueXP account ID, which can be retrieved from BlueXP console.

The screenshot shows the AWS IAM console 'Select trusted entity' page. The left sidebar displays the path 'IAM > Roles > Create role'. The main content area shows the 'Select trusted entity' page with options for 'AWS service', 'AWS account', 'SAML 2.0 federation', and 'Custom trust policy'. The 'AWS account' option is selected, and the 'Another AWS account' sub-option is chosen. The 'Account ID' field is filled with '992013314444'.

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

IAM > Roles > Create role

Step 1  
Select trusted entity

Step 2  
**Add permissions**

Step 3  
Name, review, and create

### Add permissions Info

**Permissions policies** (Selected 1/889) Info  
Choose one or more policies to attach to your new role.

Q Filter policies by property or policy name and press enter. 4 matches

'fsx' X Clear filters

	Policy name	Type	Description
<input type="checkbox"/>	AmazonFSxReadOnlyAccess	AWS ma...	Provides read only access to Amazon FSx.
<input checked="" type="checkbox"/>	AmazonFSxFullAccess	AWS ma...	Provides full access to Amazon FSx and access to related AWS services.
<input type="checkbox"/>	AmazonFSxConsoleReadOnlyAccess	AWS ma...	Provides read only access to Amazon FSx and access to related AWS services via the AWS Management Console.
<input type="checkbox"/>	AmazonFSxConsoleFullAccess	AWS ma...	Provides full access to Amazon FSx and access to related AWS services via the AWS Management Console.

► **Set permissions boundary - optional** Info  
Set a permissions boundary to control the maximum permissions this role can have. This is not a common setting, but you can use it to delegate permission management to others.

Cancel Previous **Next**

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

IAM > Roles > Create role

Step 1  
Select trusted entity

Step 2  
Add permissions

Step 3  
**Name, review, and create**

### Name, review, and create

**Role details**

**Role name**  
Enter a meaningful name to identify this role.  
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.  
Maximum 1000 characters. Use alphanumeric and +, -, @, \_ characters.

**Step 1: Select trusted entities** Edit

```

1 {
2   "Version": "2012-10-17",
3   "Statement": [
4     {
5       "Effect": "Allow",
6       "Action": "sts:AssumeRole",
7       "Principal": {
8         "AWS": "952013314444"
9       },
10      "Condition": {}
11    }
12  ]
13 }
```

5. Back to BlueXP console, click on setting icon on top right corner of the console to open **Account credentials** page, click **Add credentials** to start credential configuration workflow.

**NetApp BlueXP** Q BlueXP Search Account Automation-te... Workspace Database-2 Connector acio-aws-conn... 🔔 ⚙️ ? ⓘ

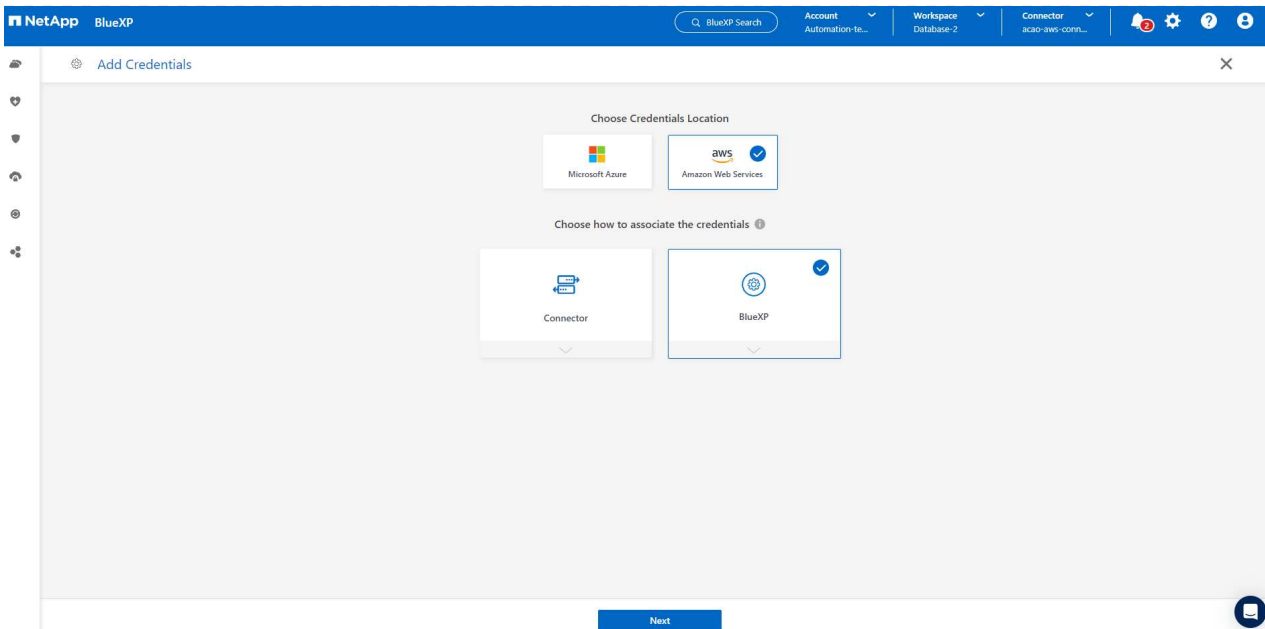
**Credentials** Account credentials User credentials

BlueXP and the Connector use account-level credentials to deploy and manage resources in your cloud environment.


5 Credentials Add credentials

Account	Type	Role
shantanucreds	Assume Role	BlueXP
210811600188	AWS Account ID	nkarthik_kafka_nfs_role_FSxN
		Assume Role


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



- 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.



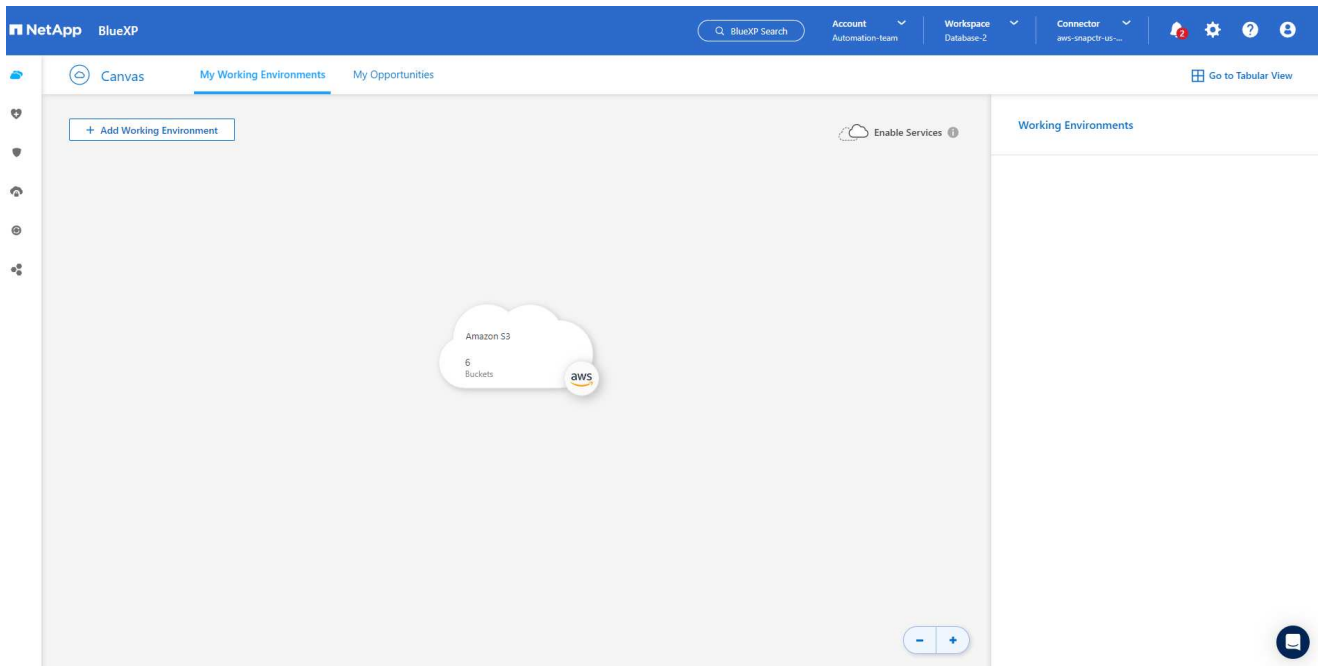
- Review and **Add**.




**SnapCenter services setup**

With the connector deployed and the credential added, SnapCenter services can now be set up with the following procedure:

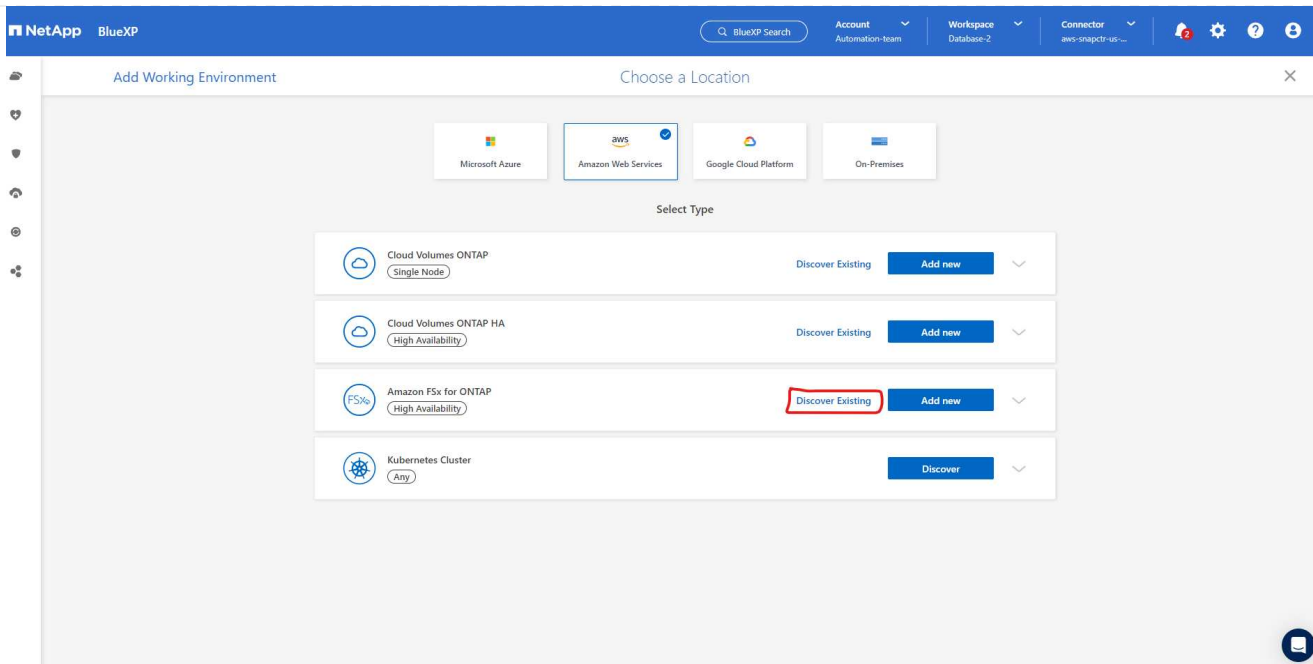
1. From **My Working Environment** click **Add working Environment** to discover FSx deployed in AWS.




1. Choose **Amazon Web Services** as the location.




1. Click **Discover Existing** next to **Amazon FSx for ONTAP**.




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



2. Choose the AWS region where Amazon FSx for ONTAP is deployed, select the FSx cluster that is hosting the Oracle database and click Add.




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




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







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




1. From the left-hand sidebar of the console, hover your mouse over the protection icon, and then click **Protection > Applications** to open the Applications launch page. Click **Discover Applications**.



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



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




1. Fill in the AWS EC2 Oracle application host details. Choose **Using SSH** as **Host Installation Type** for one step plugin installation and database discovery. Then, click on **Add SSH Private Key**.

The screenshot shows the 'Add Host' form in NetApp BlueXP. The form is titled 'Add Host' and contains the following fields and options:


- Host Installation Type:** Radio buttons for 'Manual' and 'Using SSH' (selected).
- Host FQDN or IP:** Text input field containing '172.30.15.58'.
- Connector:** Dropdown menu showing 'aws-snapctr-us-east'.
- Username(Sudo):** Text input field containing 'ec2-user'.
- SSH Port:** Text input field containing '22'.
- Plug-in Port:** Text input field containing '8145'.
- Add SSH Private Key Optional:** A button with a plus icon and a link icon.

At the bottom of the form, there are 'Cancel' and 'Next' buttons.

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



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



4. Click on **Next** to install an Oracle database plugin and discover the Oracle databases on the EC2 host. Discovered databases are added to **Applications**. The database **Protection Status** shows as **Unprotected** when initially discovered.

NetApp
BlueXP

BlueXP Search

AccountAutomation team

WorkspaceDatabase-2

Connectoraws-snapctr-us...

Backup and recovery

Volumes

Restore

Applications

Virtual Machines

Kubernetes

Job Monitoring

Cloud Native

Oracle

1

Hosts

1

ORACLE

0

Clone

Application Protection

0

Protected

1

Unprotected

1 Databases

Filter By

Manage Databases

Settings

Name	Host Name	Policy Name	Protection Status
db1	172.30.15.58		Unprotected

1 - 1 of 1

This completes the initial setup of SnapCenter services for Oracle. The next three sections of this document describe Oracle database backup, restore, and clone operations.

## Oracle database backup

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

The screenshot shows the NetApp BlueXP interface. The top navigation bar includes 'Backup and recovery', 'Volumes', 'Restore', 'Applications' (selected), 'Virtual Machines', 'Kubernetes', and 'Job Monitoring'. Below the navigation bar, there are filters for 'Cloud Native' and 'Oracle'. A summary section shows '1 Hosts', '1 ORACLE', and '0 Clone'. An 'Application Protection' section shows '0 Protected' and '1 Unprotected'. A table lists databases with columns 'Name', 'Host Name', 'Policy Name', and 'Protection Status'. The table shows one database 'db1' with host '172.30.15.58' and status 'Unprotected'. A 'Settings' dropdown menu is open, showing 'Policies', 'About', and 'Hosts'.

Name	Host Name	Policy Name	Protection Status
db1	172.30.15.58		Unprotected

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

The screenshot shows the NetApp BlueXP interface with the 'Applications > Policies' view. The top navigation bar is the same as the previous screenshot. Below the navigation bar, there are filters for 'Cloud Native' and 'Oracle'. A 'Create Policy' button is visible. A table lists policies with columns 'Policy Name', 'Backup Type', and 'Schedules and Retention'. The table shows four policies: 'Oracle Full Backup for Bronze', 'Oracle Full Backup for Gold', 'Oracle Full Backup for Silver', and 'my\_full\_bkup'. Each policy has a 'FullBackup' type and specific schedules and retention settings. A 'Create Policy' button is visible.

Policy Name	Backup Type	Schedules and Retention
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 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 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, Jul, Aug, Sep, Oct, Nov
my_full_bkup	FullBackup	Hourly: Repeats Every 1 Hr, Keeps 3 Days

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

The screenshot shows the NetApp BlueXP interface. The top navigation bar includes 'Backup and recovery', 'Volumes', 'Restore', 'Applications' (selected), 'Virtual Machines', 'Kubernetes', and 'Job Monitoring'. The 'Applications' section displays a summary for 'Cloud Native' and 'Oracle' environments. Below this, a table lists databases. The database 'db1' is shown with a protection status of 'Unprotected'. A context menu is open over the 'Unprotected' status, with the 'Assign Policy' option highlighted.


Name	Host Name	Policy Name	Protection Status
db1	172.30.15.58		Unprotected

1. Choose the policy to assign to the database.

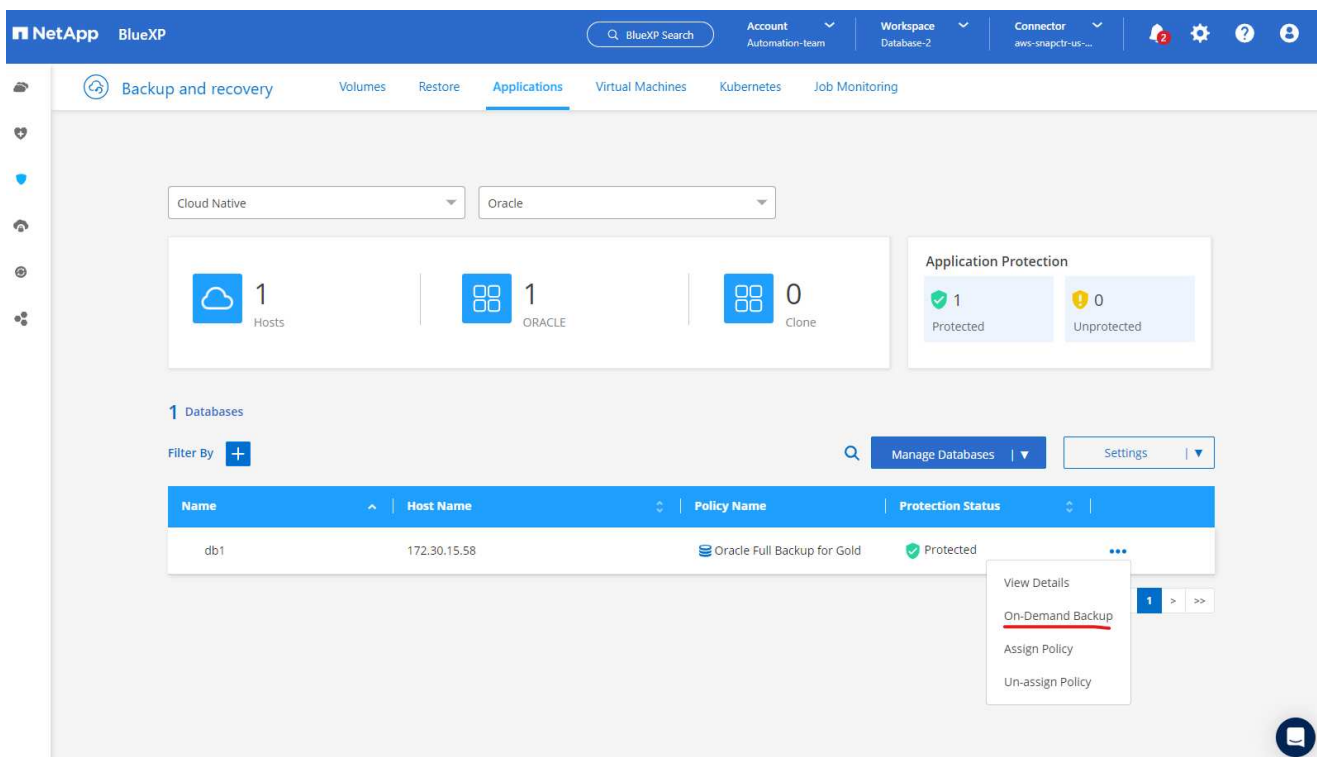
The screenshot shows the 'Assign Policy' dialog in the NetApp BlueXP interface. The dialog title is 'Assign Policy' with the subtitle 'Assign a policy to start taking backups of the database "db1"'. It displays a list of 4 policies. The policy 'my\_full\_bkup' is selected, indicated by a blue checkmark in the selection column.

Policy Name	Backup Type	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, C
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, Jul, Aug, Sep, C
my_full_bkup	FullBackup	Hourly: Repeats Every 1 Hr, Keeps 3 Days

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



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



1. The database backups details can be viewed by clicking **View Details** from the menu list. This includes the backup name, backup type, SCN, and backup date. A backup set covers a snapshot for both data volume and log volume. A log volume snapshot takes place right after a database volume snapshot. You can apply a filter if you are looking for a particular backup in a long list.



NetAppBlueXP

Q BlueXP SearchAccountAutomation-teamWorkspaceDatabase-2Connectoraws-snapctr-us-...2⚙️?👤

Backup and recoveryVolumesRestoreApplicationsVirtual MachinesKubernetesJob Monitoring

Applications > Database Details

Database Details

db1Database Name

172.30.15.58Host Name

-Clones

ProtectedProtection

FSxHost Storage

-Parent Database

Oracle Full Backup for GoldPolicy Names

UnreachableDatabase Version

Database Type

bKed8yv2T19Bj0V5QyqvA...Agent Id

8 Backups

Filter By + 🔍 Select Timeframe ▼

Backup Name	Backup Type	SCN	Backup Date	
Oracle_Full_Backup_for_Gold_Weekly_db1_2023_03_24_19_12_18_60900_1	Log	2589354	Mar 24, 2023, 3:12:34 pm	Delete
Oracle_Full_Backup_for_Gold_Weekly_db1_2023_03_24_19_11_51_51476_0	Data	2589306	Mar 24, 2023, 3:12:18 pm	...
Oracle_Full_Backup_for_Gold_Hourly_db1_2023_03_24_18_10_31_71953_1	Log	2586621	Mar 24, 2023, 2:10:45 pm	Delete
Oracle_Full_Backup_for_Gold_Hourly_db1_2023_03_24_18_10_03_70535_0	Data	2586557	Mar 24, 2023, 2:10:31 pm	...

Oracle database restore and recovery

1. For a database restore, choose the right backup, either by the SCN or backup time. Click the three dots from the database data backup, and then click **Restore** to initiate database restore and recovery.

The screenshot shows the 'Database Details' page in the NetApp BlueXP interface. The page displays database information for 'db1' and a list of 6 backups. The 'Restore' button is highlighted in the backup list.

Backup Name	Backup Type	SCN	Backup Date	
Oracle_Full_Backup_for_Gold_Hourly_db1_2023_03_24_18_10_31_71953_1	Log	2586621	Mar 24, 2023, 2:10:45 pm	Delete
Oracle_Full_Backup_for_Gold_Hourly_db1_2023_03_24_18_10_03_70535_0	Data	2586557	Mar 24, 2023, 2:10:31 pm	...
Oracle_Full_Backup_for_Gold_Hourly_db1_2023_03_24_15_37_04_98851_1	Log	2580577	Mar 24, 2023, 11:37:11	Restore
Oracle_Full_Backup_for_Gold_Hourly_db1_2023_03_24_15_36_33_27205_0	Data	2580524	Mar 24, 2023, 11:37:00	Delete Clone

1. Choose your restore setting. If you are sure that nothing has changed in the physical database structure after the backup (such as the addition of a data file or a disk group), you can use the **Force in place restore** option, which is generally faster. Otherwise, do not check this box.

The screenshot shows the 'Restore Settings' page in the NetApp BlueXP interface. The page displays the 'Restore Scope' and 'Recovery Scope' settings. The 'Force in place restore' checkbox is checked.

**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.

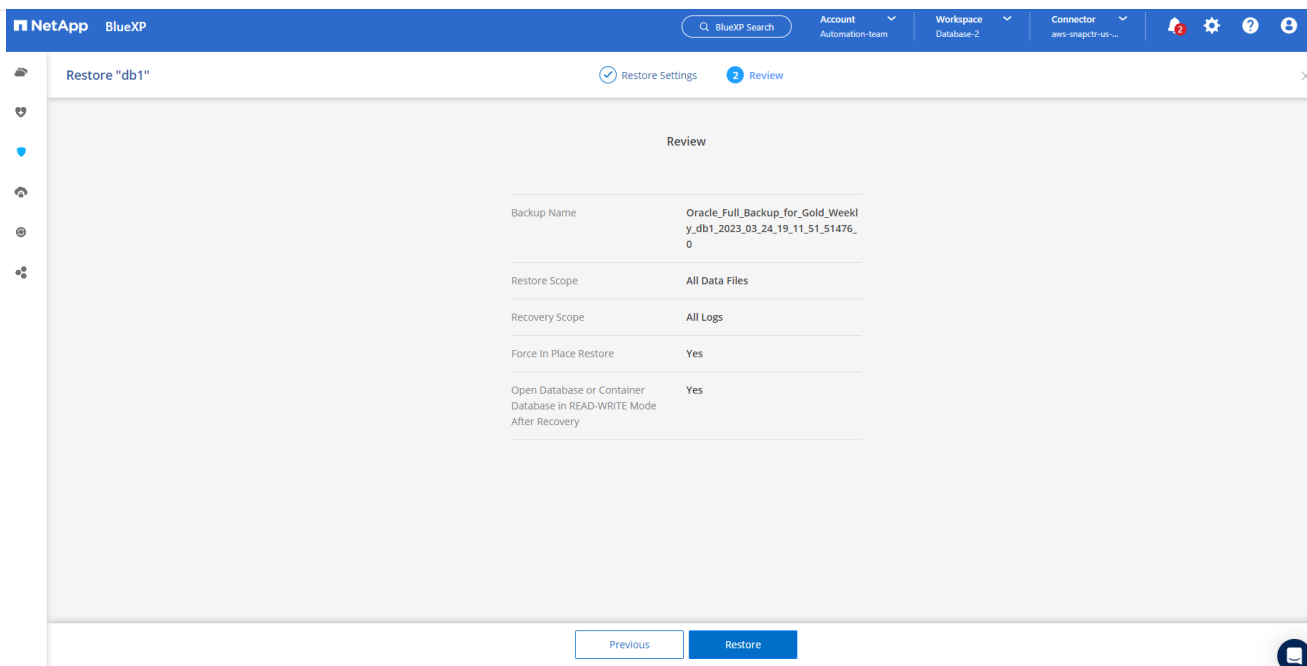
**Recovery Scope**

- ☒ All Logs
- ☐ Until System Change Number
- ☐ Date and Time
- ☐ No Recovery


Archive Log Files Locations: /mnt/log\_location001

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

1. Review and start database restore and recovery.



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



NetAppBlueXP

BlueXP Search

AccountAutomation team

WorkspaceDatabase-2

Connectoraws-snapctr-us-...

2

?

3

Backup and recovery

Volumes

Restore

Applications

Virtual Machines

Kubernetes

Job Monitoring

Job Monitoring > Job Id: 1fdca0bd-a9c8-45aa-9d7a-05a07cb291f4

Job Details

Job Id: 1fdca0bd-a9c8-45aa-9d7a-05a07cb291f4

Expand All


Sub-Jobs(6)

Job Name	Job ID	Start Time	End Time	Duration
Restore for Oracle Database db1 using backup ...	1fdca0bd-a9c8-45aa-9d...	Mar 24 2023, 3:16:28 pm	Mar 24 2023, 3:23:33 pm	7 Minutes
Post Restore Cleanup	2096a8e4-889d-4b2a-9...	Mar 24 2023, 3:23:18 pm	Mar 24 2023, 3:23:32 pm	14 Seconds
Post Restore	fb7b1171-9f6f-4228-9e...	Mar 24 2023, 3:20:06 pm	Mar 24 2023, 3:23:19 pm	3 Minutes
Restore	0f4580d0-6598-458b-a7...	Mar 24 2023, 3:17:49 pm	Mar 24 2023, 3:20:07 pm	2 Minutes

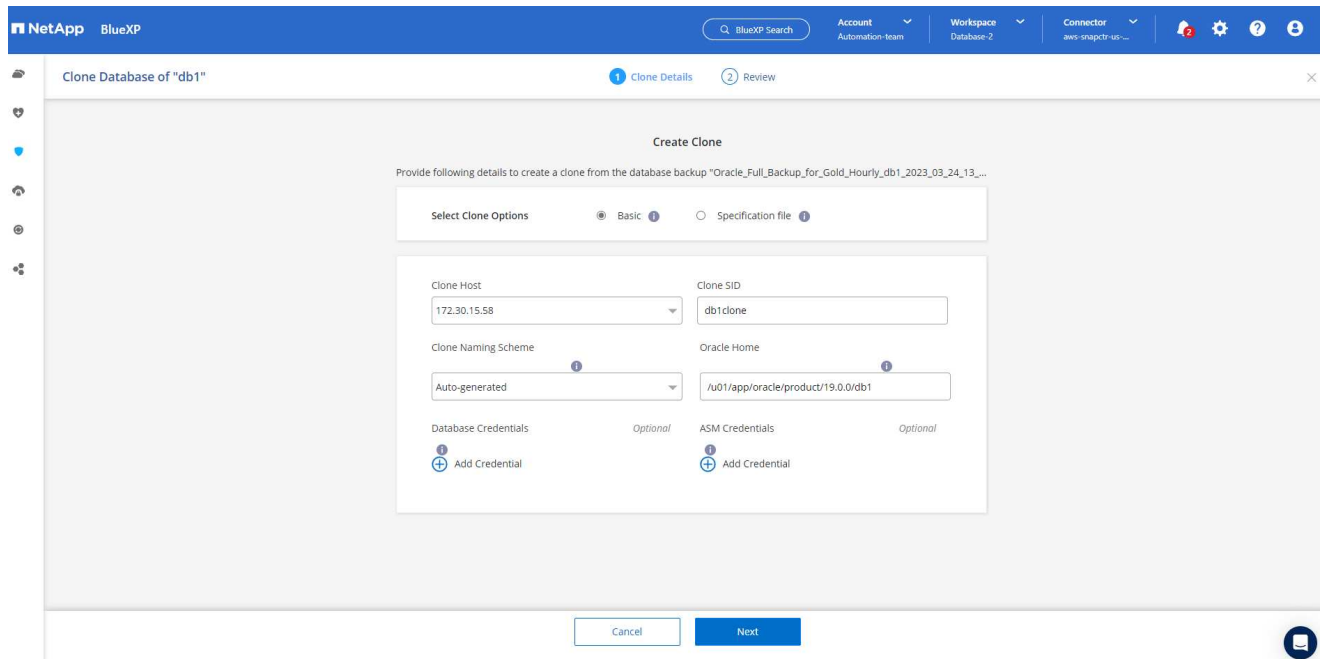
Oracle database clone

To clone a database, launch the clone workflow from the same database backup details page.

1. Select the right database backup copy, click the three dots to view the menu, and choose the **Clone** option.



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



1. Alternatively, select **Specification file**, which gives you the option of downloading the current init file, making changes, and then uploading it back to the job.

NetApp BlueXP

BlueXP Search

Account Automation team

Workspace Database-2

Connector aws-snapctr-us...

Clone Database of "db1"

1 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

☐ Basic
 ☒ Specification file

Generate specification file to modify input parameters and use for clone.

Download File

Specification File

db1\_3\_24\_2023\_10\_14\_spec.json

Browse

Clone Host

172.30.15.58

Clone SID

db1clone

Database Credentials

Optional

Add Credential

ASM Credentials

Optional

Add Credential

Cancel

Next

1. Review and launch the job.

NetApp BlueXP

BlueXP Search

Account Automation team

Workspace Database-2

Connector aws-snapctr-us...

Clone Database of "db1"

Clone Details

2 Review

Review

General	Database parameters
Backup Name	Oracle_Full_Backup_for_Gold_Hourly_db1_2023_03_24_13_34_07_26748_0
Clone SID	db1clone
Clone Host	172.30.15.58
Datafile locations	DATA_db1clone
Control files	+DATA_db1clone/db1clone/control/control01.ctl
Redo logs	RedoGroup = 1 TotalSize = 1024 Path = +DATA_db1clone/db1clone/redolog/redo01_01.log RedoGroup = 2 TotalSize = 1024 Path = +DATA_db1clone/db1clone/redolog/redo02_01.log RedoGroup = 3 TotalSize = 1024 Path = +DATA_db1clone/db1clone/redolog/redo03_01.log
Recovery scope	Until cancel using selected backup's archive logs

Previous

Clone

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

104

NetAppBlueXP

BlueXP Search

AccountAutomation-team

WorkspaceDatabase-2

Connectoraws-snapc1r108-...

Backup and recovery

Volumes

Restore

Applications

Virtual Machines

Kubernetes

Job Monitoring

Job Monitoring > Job Id: cd30abaf-fbe2-4052-a6db-4bf965a8d29b

Job Details

Job Id: cd30abaf-fbe2-4052-a6db-4bf965a8d29b

Expand All

Sub-Jobs(2)

Job Name	Job ID	Start Time	End Time	Duration
Cloning Oracle Database db1 as db1clone on h...	cd30abaf-fbe2-4052-a6...	Mar 24 2023, 1:30:36 pm		--
Running pre scripts	51f152c1-853a-4ec6-a4f...	Mar 24 2023, 1:30:41 pm	Mar 24 2023, 1:30:41 pm	0 Second
Validating clone request	f93a6c44-2eb2-4c5e-9f...	Mar 24 2023, 1:30:35 pm	Mar 24 2023, 1:30:42 pm	7 Seconds

1. Validate the cloned database on the EC2 instance host.

```
#
# Multiple entries with the same $ORACLE_SID are not allowed.
#
#
+ASM:/u01/app/oracle/product/19.0.0/grid:N
db1:/u01/app/oracle/product/19.0.0/db1:N
# SnapCenter Plug-in for Oracle Database generated entry (DO NOT REMOVE THIS LINE)
db1clone:/u01/app/oracle/product/19.0.0/db1:N
[oracle@ip-172-30-15-58 ~]$ crsctl stat res -t
-----
Name                Target  State        Server                    State details
-----
Local Resources
-----
ora.DATA.dg
      ONLINE  ONLINE      ip-172-30-15-58          STABLE
ora.DATA_DB1CLONE.dg
      ONLINE  ONLINE      ip-172-30-15-58          STABLE
ora.LISTENER.lsnr
      ONLINE  ONLINE      ip-172-30-15-58          STABLE
ora.LOGS.dg
      ONLINE  ONLINE      ip-172-30-15-58          STABLE
ora.LOGS_SCO_2748138658.dg
      ONLINE  ONLINE      ip-172-30-15-58          STABLE
ora.asm
      ONLINE  ONLINE      ip-172-30-15-58          Started,STABLE
ora.ons
      OFFLINE OFFLINE      ip-172-30-15-58          STABLE
-----
Cluster Resources
-----
ora.cssd
      1        ONLINE  ONLINE      ip-172-30-15-58          STABLE
ora.db1.db
      1        ONLINE  ONLINE      ip-172-30-15-58          Open,HOME=/u01/app/o
                             racle/product/19.0.0
                             /db1,STABLE
ora.db1clone.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.driver.afd
      1        ONLINE  ONLINE      ip-172-30-15-58          STABLE
ora.evmd
      1        ONLINE  ONLINE      ip-172-30-15-58          STABLE
-----
[oracle@ip-172-30-15-58 ~]$
```

```
[oracle@ip-172-30-15-58 ~]$ export ORACLE_HOME=/u01/app/oracle/product/19.0.0/db1
[oracle@ip-172-30-15-58 ~]$ export ORACLE_SID=db1clone
[oracle@ip-172-30-15-58 ~]$ export PATH=$ORACLE_HOME/bin:$PATH
[oracle@ip-172-30-15-58 ~]$ sqlplus / as sysdba
```

```
SQL*Plus: Release 19.0.0.0.0 - Production on Fri Mar 24 18:32:21 2023
Version 19.18.0.0.0
```

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

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

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

```
NAME          OPEN_MODE
-----
DB1CLONE      READ WRITE
```

```
SQL>
```



## Additional information

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

- Set up and administer BlueXP

<https://docs.netapp.com/us-en/cloud-manager-setup-admin/index.html>

- BlueXP backup and recovery documentation

<https://docs.netapp.com/us-en/cloud-manager-backup-restore/index.html>

- Amazon FSx for NetApp ONTAP

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

- Amazon EC2

[https://aws.amazon.com/pm/ec2/?trk=36c6da98-7b20-48fa-8225-4784bcd9843&sc\\_channel=ps&s\\_kwid=AL!4422!3!467723097970!e!!g!!aws%20ec2&ef\\_id=Cj0KCQiA54KfBhCKARIsAJzSrdqWQrghn6l71jiWzSeaT9Uh1-vY-VfhJixF-xnv5rWwn2S7RqZOTQ0aAh7eEALw\\_wcB:G:s&s\\_kwid=AL!4422!3!467723097970!e!!g!!aws%20ec2](https://aws.amazon.com/pm/ec2/?trk=36c6da98-7b20-48fa-8225-4784bcd9843&sc_channel=ps&s_kwid=AL!4422!3!467723097970!e!!g!!aws%20ec2&ef_id=Cj0KCQiA54KfBhCKARIsAJzSrdqWQrghn6l71jiWzSeaT9Uh1-vY-VfhJixF-xnv5rWwn2S7RqZOTQ0aAh7eEALw_wcB:G:s&s_kwid=AL!4422!3!467723097970!e!!g!!aws%20ec2)

## Hybrid Cloud Database Solutions with SnapCenter

### TR-4908: Hybrid Cloud Database Solutions with SnapCenter Overview

Alan Cao, Felix Melligan, NetApp

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


- Database dev/test operations in the hybrid cloud
- Database disaster recovery in the hybrid cloud

Today, many enterprise databases still reside in private corporate data centers for performance, security, and/or other reasons. This hybrid cloud database solution enables enterprises to operate their primary databases on site while using a public cloud for dev/test database operations as well as for disaster recovery to reduce licensing and operational costs.

Many enterprise databases, such as Oracle, SQL Server, SAP HANA, and so on, carry high licensing and operational costs. Many customers pay a one-time license fee as well as annual support costs based on the number of compute cores in their database environment, whether the cores are used for development, testing, production, or disaster recovery. Many of those environments might not be fully utilized throughout the application lifecycle.

The solutions provide an option for customers to potentially reduce their licensable cores count by moving their database environments devoted to development, testing, or disaster recovery to the cloud. By using public-cloud scale, redundancy, high availability, and a consumption-based billing model, the cost saving for licensing and operation can be substantial, while not sacrificing any application usability or availability.

Beyond potential database license-cost savings, the NetApp capacity-based CVO license model allows customers to save storage costs on a per-GB basis while empowering them with high level of database manageability that is not available from competing storage services. The following chart shows a storage cost comparison of popular storage services available in the public cloud.



This solution demonstrates that, by using the SnapCenter GUI-based software tool and NetApp SnapMirror technology, hybrid cloud database operations can be easily setup, implemented, and operated.

The following videos demonstrate SnapCenter in action:


- [Backup of an Oracle database across a Hybrid Cloud using SnapCenter](#)
- [SnapCenter- Clone DEV/TEST to AWS Cloud for an Oracle database](#)

Notably, although the illustrations throughout this document show CVO as a target storage instance in the public cloud, the solution is also fully validated for the new release of the FSx ONTAP storage engine for AWS.

To test drive the solution and use cases for yourself, a NetApp Lab-on-Demand SL10680 can be requested at following xref:./databases/ [TL\\_AWS\\_004 HCoD: AWS - NW,SnapCenter\(OnPrem\)](#).

### Solution Architecture

The following architecture diagram illustrates a typical implementation of enterprise database operation in a hybrid cloud for dev/test and disaster recovery operations.



In normal business operations, synchronized database volumes in the cloud can be cloned and mounted to dev/test database instances for applications development or testing. In the event of a failure, the synchronized database volumes in the cloud can then be activated for disaster recovery.

## SnapCenter Requirements

This solution is designed in a hybrid cloud setting to support on-premises production databases that can burst to all of the popular public clouds for dev/test and disaster recovery operations.

This solution supports all databases that are currently supported by SnapCenter, although only Oracle and SQL Server databases are demonstrated here. This solution is validated with virtualized database workloads, although bare-metal workloads are also supported.

We assume that production database servers are hosted on-premises with DB volumes presented to DB hosts from a ONTAP storage cluster. SnapCenter software is installed on-premises for database backup and data replication to the cloud. An Ansible controller is recommended but not required for database deployment automation or OS kernel and DB configuration syncing with a standby DR instance or dev/test instances in the public cloud.

## Requirements

Environment	Requirements
On-premises	Any databases and versions supported by SnapCenter
	SnapCenter v4.4 or higher
	Ansible v2.09 or higher
	ONTAP cluster 9.x
	Intercluster LIFs configured
	Connectivity from on-premises to a cloud VPC (VPN, interconnect, and so on)
	Networking ports open - ssh 22 - tcp 8145, 8146, 10000, 11104, 11105
Cloud - AWS	<a href="#">Cloud Manager Connector</a>
	<a href="#">Cloud Volumes ONTAP</a>
	Matching DB OS EC2 instances to On-prem
Cloud - Azure	<a href="#">Cloud Manager Connector</a>
	<a href="#">Cloud Volumes ONTAP</a>
	Matching DB OS Azure Virtual Machines to On-prem
Cloud - GCP	<a href="#">Cloud Manager Connector</a>
	<a href="#">Cloud Volumes ONTAP</a>
	Matching DB OS Google Compute Engine instances to on-premises

## Prerequisites configuration

Certain prerequisites must be configured both on-premises and in the cloud before the execution of hybrid cloud database workloads. The following section provides a high-level summary of this process, and the following links provide further information about necessary system configuration.

### On premises

- SnapCenter installation and configuration
- On-premises database server storage configuration
- Licensing requirements
- Networking and security
- Automation

### Public cloud

- A NetApp Cloud Central login
- Network access from a web browser to several endpoints
- A network location for a connector

- Cloud provider permissions
- Networking for individual services

Important considerations:

1. Where to deploy the Cloud Manager Connector?
2. Cloud Volume ONTAP sizing and architecture
3. Single node or high availability?

The following links provide further details:

[On Premises](#)

[Public Cloud](#)

### Prerequisites on-premises

The following tasks must be completed on-premises to prepare the SnapCenter hybrid-cloud database workload environment.

#### SnapCenter installation and configuration

The NetApp SnapCenter tool is a Windows-based application that typically runs in a Windows domain environment, although workgroup deployment is also possible. It is based on a multitiered architecture that includes a centralized management server (the SnapCenter server) and a SnapCenter plug-in on the database server hosts for database workloads. Here are a few key considerations for hybrid-cloud deployment.


- **Single instance or HA deployment.** HA deployment provides redundancy in the case of a single SnapCenter instance server failure.
- **Name resolution.** DNS must be configured on the SnapCenter server to resolve all database hosts as well as on the storage SVM for forward and reverse lookup. DNS must also be configured on database servers to resolve the SnapCenter server and the storage SVM for both forward and reverse lookup.
- **Role-based access control (RBAC) configuration.** For mixed database workloads, you might want to use RBAC to segregate management responsibility for different DB platform such as an admin for Oracle database or an admin for SQL Server. Necessary permissions must be granted for the DB admin user.
- **Enable policy-based backup strategy.** To enforce backup consistency and reliability.
- **Open necessary network ports on the firewall.** For the on-premises SnapCenter server to communicate with agents installed in the cloud DB host.
- **Ports must be open to allow SnapMirror traffic between on-prem and public cloud.** The SnapCenter server relies on ONTAP SnapMirror to replicate onsite Snapshot backups to cloud CVO storage SVMs.

After careful pre-installation planning and consideration, click this [SnapCenter installation workflow](#) for details of SnapCenter installation and configuration.


#### On-premises database server storage configuration


Storage performance plays an important role in the overall performance of databases and applications. A well-designed storage layout can not only improve DB performance but also make it easy to manage database backup and recovery. Several factors should be considered when defining your storage layout, including the size of the database, the rate of expected data change for the database, and the frequency with which you perform backups.

Directly attaching storage LUNs to the guest VM by either NFS or iSCSI for virtualized database workloads generally provides better performance than storage allocated via VMDK. NetApp recommends the storage layout for a large SQL Server database on LUNs depicted in the following figure.



The following figure shows the NetApp recommended storage layout for small or medium SQL Server database on LUNs.



 The Log directory is dedicated to SnapCenter to perform transaction log rollup for database recovery. For an extra large database, multiple LUNs can be allocated to a volume for better performance.

For Oracle database workloads, SnapCenter supports database environments backed by ONTAP storage that are mounted to the host as either physical or virtual devices. You can host the entire database on a single or multiple storage devices based on the criticality of the environment. Typically, customers isolate data files on dedicated storage from all other files such as control files, redo files, and archive log files. This helps administrators to quickly restore (ONTAP single-file SnapRestore) or clone a large critical database (petabyte scale) using Snapshot technology within few seconds to minutes.





For mission critical workloads that are sensitive to latency, a dedicated storage volume should be deployed to different types of Oracle files to achieve the best latency possible. For a large database, multiple LUNs (NetApp recommends up to eight) per volume should be allocated to data files.



For smaller Oracle databases, SnapCenter supports shared storage layouts in which you can host multiple databases or part of a database on the same storage volume or LUN. As an example of this layout, you can host data files for all the databases on a +DATA ASM disk group or a volume group. The remainder of the files (redo, archive log, and control files) can be hosted on another dedicated disk group or volume group (LVM). Such a deployment scenario is illustrated below.



To facilitate the relocation of Oracle databases, the Oracle binary should be installed on a separate LUN that is included in the regular backup policy. This ensures that in the case of database relocation to a new server host, the Oracle stack can be started for recovery without any potential issues due to an out-of-sync Oracle binary.

## Licensing requirements

SnapCenter is licensed software from NetApp. It is generally included in an on-premises ONTAP license. However, for hybrid cloud deployment, a cloud license for SnapCenter is also required to add CVO to SnapCenter as a target data replication destination. Please review following links for SnapCenter standard capacity-based license for details:

[SnapCenter standard capacity-based licenses](#)

## Networking and security

In a hybrid database operation that requires an on-premises production database that is burstable to cloud for dev/test and disaster recovery, networking and security is important factor to consider when setting up the environment and connecting to the public cloud from an on-premises data center.

Public clouds typically use a virtual private cloud (VPC) to isolate different users within a public-cloud platform. Within an individual VPC, security is controlled using measures such as security groups that are configurable based on user needs for the lockdown of a VPC.

The connectivity from the on-premises data center to the VPC can be secured through a VPN tunnel. On the VPN gateway, security can be hardened using NAT and firewall rules that block attempts to establish network

connections from hosts on the internet to hosts inside the corporate data center.

For networking and security considerations, review the relevant inbound and outbound CVO rules for your public cloud of choice:

- [Security group rules for CVO - AWS](#)
- [Security group rules for CVO - Azure](#)
- [Firewall rules for CVO - GCP](#)

#### **Using Ansible automation to sync DB instances between on-premises and the cloud - optional**

To simplify management of a hybrid-cloud database environment, NetApp highly recommends but does not require that you deploy an Ansible controller to automate some management tasks, such as keeping compute instances on-premises and in the cloud in sync. This is particularly important because an out-of-sync compute instance in the cloud might render the recovered database in the cloud error prone because of missing kernel packages and other issues.

The automation capability of an Ansible controller can also be used to augment SnapCenter for certain tasks, such as breaking up the SnapMirror instance to activate the DR data copy for production.

Follow these instructions to set up your Ansible control node for RedHat or CentOS machines: [RedHat/CentOS Ansible Controller Setup](#).

Follow these instructions 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.



## 2. SnapCenter plugin installation prerequisites

SnapCenter performs backup, restore, clone, and other functions by using a plugin agent running on the DB hosts. It connects to the database host and database via credentials configured under the Setting and Credentials tab for plugin installation and other management functions. There are specific privilege requirements based on the target host type, such as Linux or Windows, as well as the type of database.

DB hosts credentials must be configured before SnapCenter plugin installation. Generally, you want to use an administrator user accounts on the DB host as your host connection credentials for plugin installation. You can also grant the same user ID for database access using OS-based authentication. On the other hand, you can also employ database authentication with different database user IDs for DB management access. If you decide to use OS-based authentication, the OS admin user ID must be granted DB access. For Windows domain-based SQL Server installation, a domain admin account can be used to manage all SQL Servers within the domain.

Windows host for SQL server:

1. If you are using Windows credentials for authentication, you must set up your credential before installing plugins.
2. If you are using a SQL Server instance for authentication, you must add the credentials after installing plugins.
3. If you have enabled SQL authentication while setting up the credentials, the discovered instance or database is shown with a red lock icon. If the lock icon appears, you must specify the instance or database credentials to successfully add the instance or database to a resource group.
4. You must assign the credential to a RBAC user without sysadmin access when the following conditions are met:
  - The credential is assigned to a SQL instance.
  - The SQL instance or host is assigned to an RBAC user.
  - The RBAC DB admin user must have both the resource group and backup privileges.

Unix host for Oracle:

1. You must have enabled the password-based SSH connection for the root or non-root user by editing sshd.conf and restarting the sshd service. Password-based SSH authentication on AWS instance is turned off by default.
2. Configure the sudo privileges for the non-root user to install and start the plugin process. After installing the plugin, the processes run as an effective root user.
3. Create credentials with the Linux authentication mode for the install user.

4. You must install Java 1.8.x (64-bit) on your Linux host.
5. Installation of the Oracle database plugin also installs the SnapCenter plugin for Unix.

### 3. SnapCenter host plugin installation

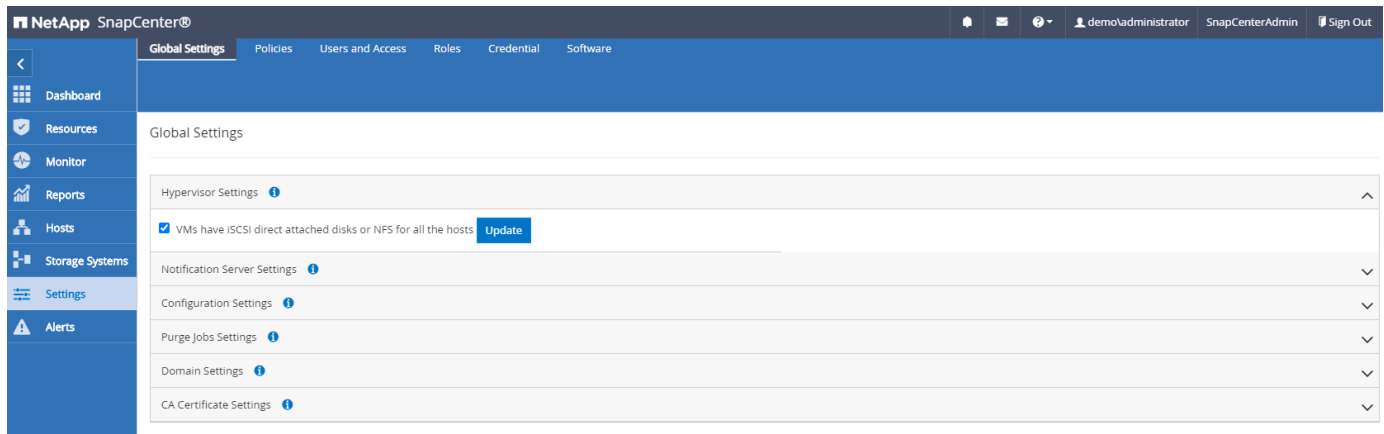


Before attempting to install SnapCenter plugins on cloud DB server instances, make sure that all configuration steps have been completed as listed in the relevant cloud section for compute instance deployment.

The following steps illustrate how a database host is added to SnapCenter while a SnapCenter plugin is installed on the host. The procedure applies to adding both on-premises hosts and cloud hosts. The following demonstration adds a Windows or a Linux host residing in AWS.


#### Configure SnapCenter VMware global settings

Navigate to Settings > Global Settings. Select "VMs have iSCSI direct attached disks or NFS for all the hosts" under Hypervisor Settings and click Update.




#### Add Windows host and installation of plugin on the host

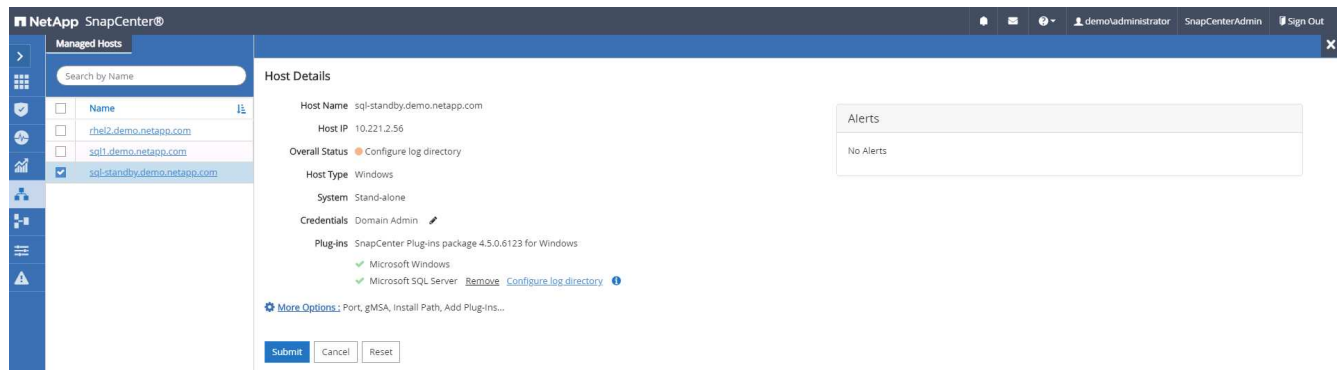
1. Log into SnapCenter with a user ID with SnapCenterAdmin privileges.
2. Click the Hosts tab from the left-hand menu, and then click Add to open the Add Host workflow.
3. Choose Windows for Host Type; the Host Name can be either a host name or an IP address. The host name must be resolved to the correct host IP address from the SnapCenter host. Choose the host credentials created in step 2. Choose Microsoft Windows and Microsoft SQL Server as the plugin packages to be installed.



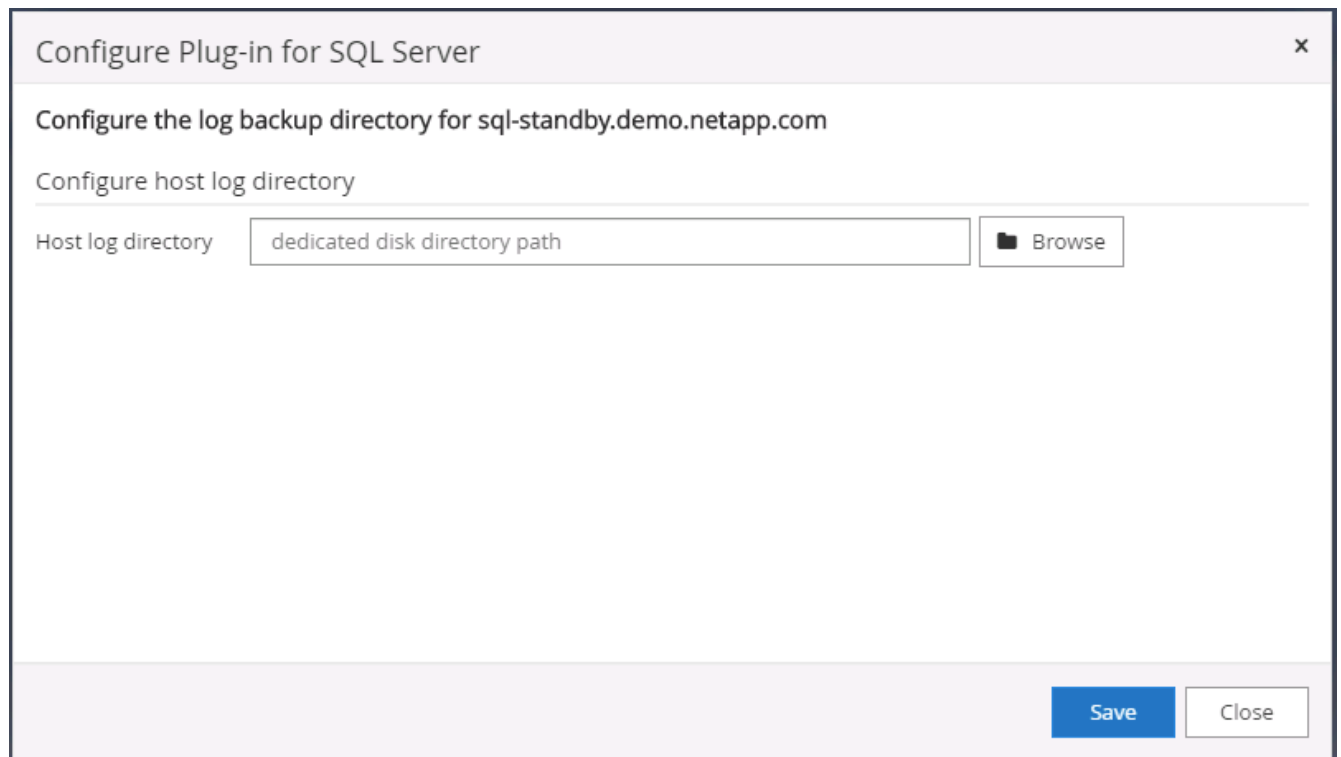
4. After the plugin is installed on a Windows host, its Overall Status is shown as "Configure log directory."



5. Click the Host Name to open the SQL Server log directory configuration.



6. Click "Configure log directory" to open "Configure Plug-in for SQL Server."



7. Click Browse to discover NetApp storage so that a log directory can be set; SnapCenter uses this log directory to roll up the SQL server transaction log files. Then click Save.

## Configure Plug-in for SQL Server

Configure the log backup directory for sql-standby.demo.netapp.com

Configure host log directory

Host log directory  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.

- After the log directory is configured, the Windows host plugin Overall Status is changed to Running.

NetApp SnapCenter®							
Managed Hosts							
Search by Name							
	Name	Type	System	Plug-in	Version	Overall Status	
<input type="checkbox"/>	rhel2.demo.netapp.com	Linux	Stand-alone	UNIX, Oracle Database	4.5	Running	
<input type="checkbox"/>	sql1.demo.netapp.com	Windows	Stand-alone	Microsoft Windows Server, Microsoft SQL Server	4.5	Running	
<input type="checkbox"/>	sql-standby.demo.netapp.com	Windows	Stand-alone	Microsoft Windows Server, Microsoft SQL Server	4.5	Running	

- 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 sqlldb that the host needs to be assigned to), and click Save to complete host resource assignment.

NetApp SnapCenter®					
Users and Access					
Search by Name					
	Name	Type	Roles	Domain	
<input type="checkbox"/>	administrator	User	SnapCenterAdmin	demo	
<input type="checkbox"/>	oracdba	User	App Backup and Clone Admin	demo	
<input type="checkbox"/>	sqlldb	User	App Backup and Clone Admin	demo	

Assign Assets

Asset Type
Host
search

<input type="checkbox"/>	Asset Name
<input type="checkbox"/>	rhel2.demo.netapp.com
<input type="checkbox"/>	sql1.demo.netapp.com
<input checked="" type="checkbox"/>	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.

demoadministrator
SnapCenterAdmin
Sign Out

Add Host

Host Type
Linux
Host Name
ora-standby
Credentials
admin

Select Plug-ins to Install
SnapCenter Plug-ins Package 4.5 for Linux
☒ Oracle Database
☐ SAP HANA

More Options
Port, 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

8145

Installation Path

/opt/NetApp/snapcenter

☒

Skip preinstall checks

☒

Add all hosts in the oracle RAC

Custom Plug-ins

Choose a File

Browse

Upload

No plug-ins found.

Save

Cancel

5. Click Submit to start the plugin installation. You are prompted to Confirm Fingerprint as shown below.

Confirm Fingerprint

Authenticity of the host cannot be determined

Host name	Fingerprint	Valid
ora-standby.demo.netapp.com	ssh-rsa 3072 5C:02:EF:6B:63:54:59:10:84:DF:4D:6B:AB:FB:61:67	

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.

NetApp SnapCenter®

Managed Hosts

Disks

Shares

Initiator Groups

ISCSI Session

Search by Name

+

+



+

+

	Name	Type	System	Plug-in	Version	Overall Status
<input type="checkbox"/>	ora-standby.demo.netapp.com	Linux	Stand-alone	UNIX, Oracle Database	4.5	Running
<input type="checkbox"/>	rhel2.demo.netapp.com	Linux	Stand-alone	UNIX, Oracle Database	4.5	Running
<input type="checkbox"/>	sql1.demo.netapp.com	Windows	Stand-alone	Microsoft Windows Server, Microsoft SQL Server	4.5	Running
<input type="checkbox"/>	sql-standby.demo.netapp.com	Windows	Stand-alone	Microsoft Windows Server, Microsoft SQL Server	4.5	Running

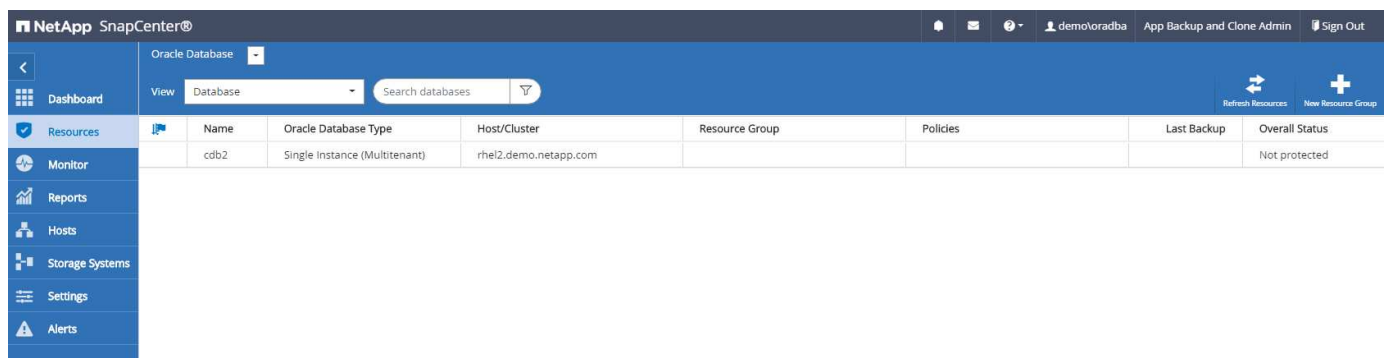
7. Assign the newly added host to the proper database management user ID (in our case, oradba).





#### 4. Database resource discovery

With successful plugin installation, the database resources on the host can be immediately discovered. Click the Resources tab in the left-hand menu. Depending on the type of database platform, a number of views are available, such as the database, resources group, and so on. You might need to click the Refresh Resources tab if the resources on the host are not discovered and displayed.



When the database is initially discovered, the Overall Status is shown as "Not protected." The previous screenshot shows an Oracle database not protected yet by a backup policy.

When a backup configuration or policy is set up and a backup has been executed, the Overall Status for the database shows the backup status as "Backup succeeded" and the timestamp of the last backup. The following screenshot shows the backup status of a SQL Server user database.



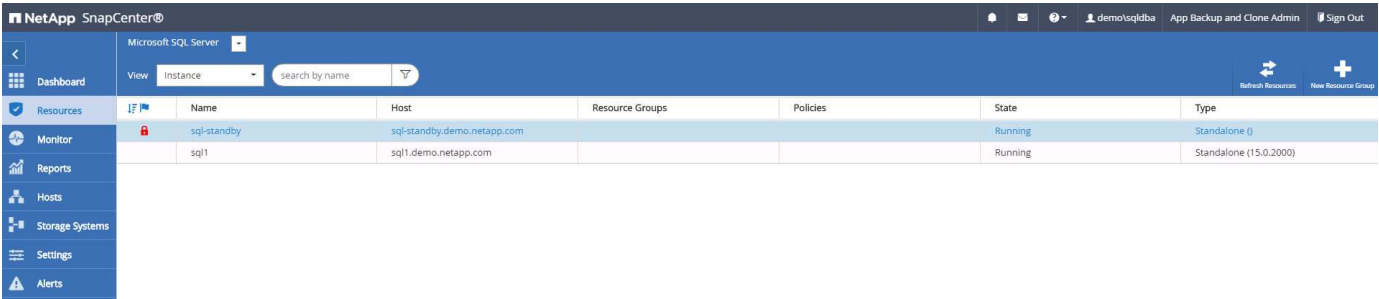
NetApp SnapCenter®

Microsoft SQL Server

View Database search by name

	Name	Instance	Host	Last Backup	Overall Status	Type
	master	sql1	sql1.demo.netapp.com		Not available for backup	System database
	model	sql1	sql1.demo.netapp.com		Not available for backup	System database
	msdb	sql1	sql1.demo.netapp.com		Not available for backup	System database
	tempdb	sql1	sql1.demo.netapp.com		Not available for backup	System database
	tpcc	sql1	sql1.demo.netapp.com	09/14/2021 2:35:07 PM	Backup succeeded	User database

If database access credentials are not properly set up, a red lock button indicates that the database is not accessible. For example, if Windows credentials do not have sysadmin access to a database instance, then database credentials must be reconfigured to unlock the red lock.

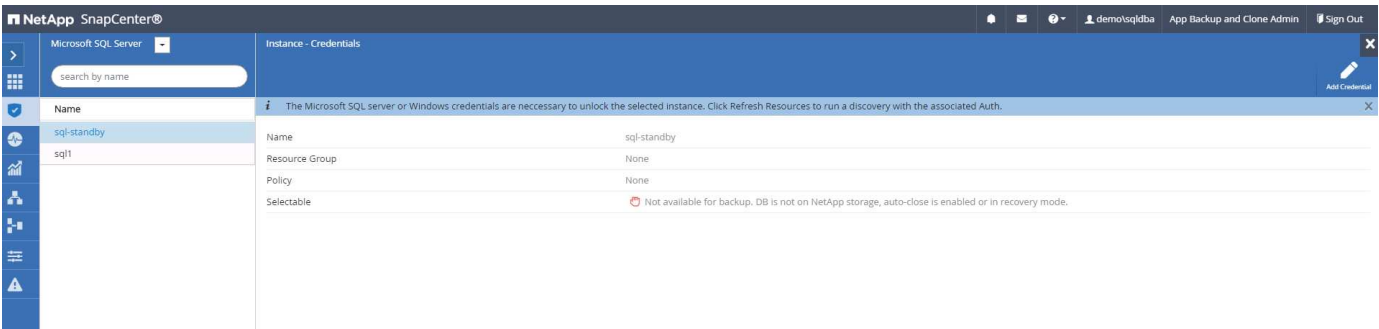


NetApp SnapCenter®

Microsoft SQL Server

View Instance search by name

	Name	Host	Resource Groups	Policies	State	Type
	sql-standby	sql-standby.demo.netapp.com			Running	Standalone ()
	sql1	sql1.demo.netapp.com			Running	Standalone (15.0.2000)



NetApp SnapCenter®

Microsoft SQL Server

Instance - Credentials

search by name

The Microsoft SQL server or Windows credentials are necessary to unlock the selected instance. Click Refresh Resources to run a discovery with the associated Auth.

Name	sql-standby
Resource Group	None
Policy	None
Selectable	Not available for backup. DB is not on NetApp storage, auto-close is enabled or in recovery mode.

After the appropriate credentials are configured either at the Windows level or the database level, the red lock disappears and SQL Server Type information is gathered and reviewed.



NetApp SnapCenter®

Microsoft SQL Server

View Instance search by name

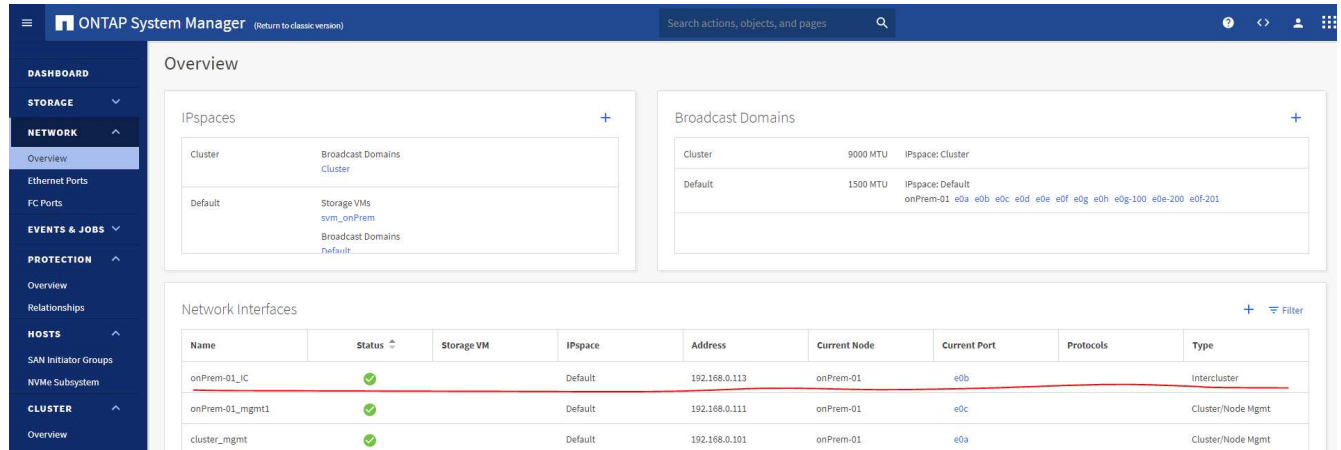
	Name	Host	Resource Groups	Policies	State	Type
	sql1	sql1.demo.netapp.com			Running	Standalone (15.0.2000)
	sql-standby	sql-standby.demo.netapp.com			Running	Standalone (15.0.2000)

## 5. Setup storage cluster peering and DB volumes replication

To protect your on-premises database data using a public cloud as the target destination, on-premises ONTAP cluster database volumes are replicated to the cloud CVO using NetApp SnapMirror technology. The replicated target volumes can then be cloned for DEV/OPS or disaster recovery. The following high-level steps enable you to set up cluster peering and DB volumes replication.

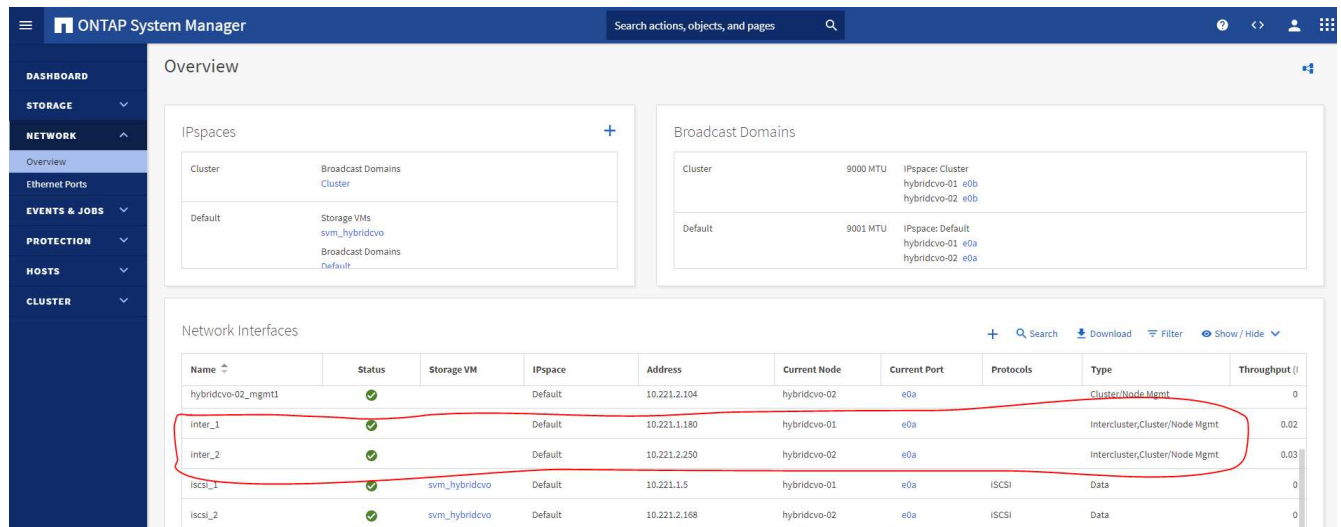
1. Configure intercluster LIFs for cluster peering on both the on-premises cluster and the CVO cluster instance. This step can be performed with ONTAP System Manager. A default CVO deployment has intercluster LIFs configured automatically.

On-premises cluster:



Name	Status	Storage VM	IPspace	Address	Current Node	Current Port	Protocols	Type
onPrem-01_IC	✓		Default	192.168.0.113	onPrem-01	e0b		Intercluster
onPrem-01_mgmt1	✓		Default	192.168.0.111	onPrem-01	e0c		Cluster/Node Mgmt
cluster_mgmt	✓		Default	192.168.0.101	onPrem-01	e0a		Cluster/Node Mgmt

Target CVO cluster:




Name	Status	Storage VM	IPspace	Address	Current Node	Current Port	Protocols	Type	Throughput (I)
hybridcvo-02_mgmt1	✓		Default	10.221.2.104	hybridcvo-02	e0a		Cluster/Node Mgmt	0
inter_1	✓		Default	10.221.1.180	hybridcvo-01	e0a		Intercluster,Cluster/Node Mgmt	0.02
inter_2	✓		Default	10.221.2.250	hybridcvo-02	e0a		Intercluster,Cluster/Node Mgmt	0.03
iscsi_1	✓	svm_hybridcvo	Default	10.221.1.5	hybridcvo-01	e0a	ISCSI	Data	0
iscsi_2	✓	svm_hybridcvo	Default	10.221.2.168	hybridcvo-02	e0a	ISCSI	Data	0


2. With the intercluster LIFs configured, cluster peering and volume replication can be set up by using drag-and-drop in NetApp Cloud Manager. See ["Getting Started - AWS Public Cloud"](#) for details.

Alternatively, cluster peering and DB volume replication can be performed by using ONTAP System Manager as follows:

3. Log into ONTAP System Manager. Navigate to Cluster > Settings and click Peer Cluster to set up cluster peering with the CVO instance in the cloud.



4. Go to the Volumes tab. Select the database volume to be replicated and click Protect.



5. Set the protection policy to Asynchronous. Select the destination cluster and storage SVM.

ONTAP System Manager

(Return to classic version)

Search actions, objects, and pages

DASHBOARD

STORAGE

Overview

Applications

Volumes

LUNS

NVMe Namespaces

Shares

Qtrees

Quotas

Storage VMs

Tiers

NETWORK

Overview

Ethernet Ports

FC Ports

EVENTS & JOBS

PROTECTION

HOSTS

CLUSTER

Protect Volumes

PROTECTION POLICY

Asynchronous

Source

CLUSTER

onPrem

STORAGE VM

svm\_onPrem

SELECTED VOLUMES

rhel2\_u03

Destination

CLUSTER

hybridcvo

STORAGE VM

svm\_hybridcvo

Destination Settings

2 matching labels

VOLUME NAME

PREFIX

vol\_

SUFFIX

<SourceVolumeName>\_dest

☐ Override default storage service name

Configuration Details

☒ Initialize relationship
 ☐ Enable FabricPool

Save

Cancel

- Validate that the volume is synced between the source and target and that the replication relationship is healthy.

Volumes

+ Add

Delete

Protect

More

rhel2\_u03

All Volumes

Edit

More

onPrem\_data

rhel2\_u01

rhel2\_u02

rhel2\_u03

rhel2\_u0309232119421203118

Overview

Snapshot Copies


Clone Hierarchy

SnapMirror (Local or Remote)

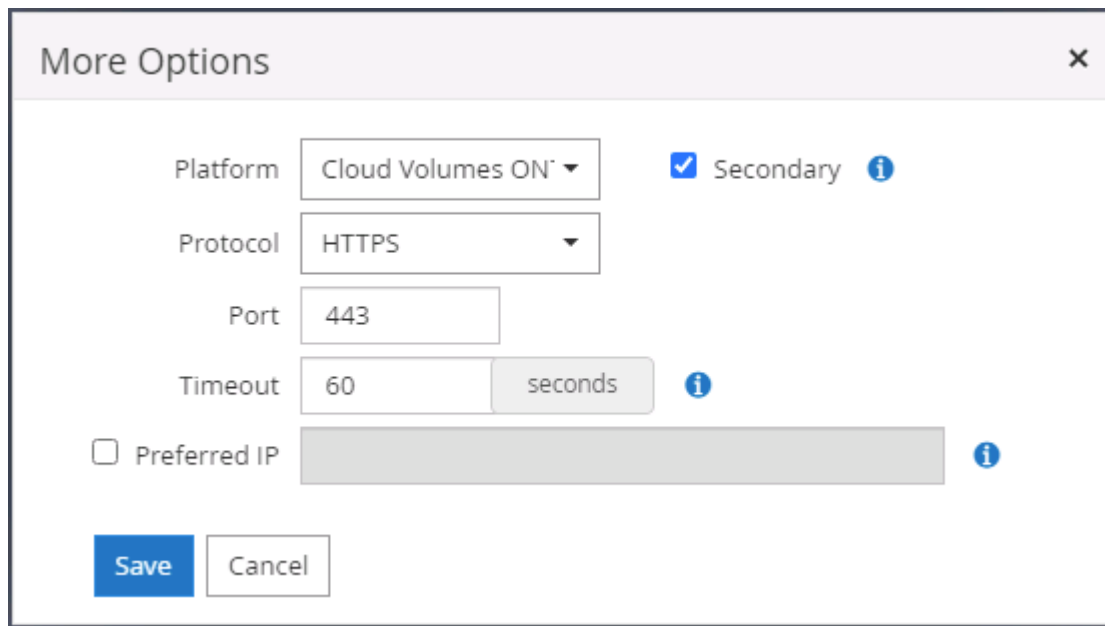
Source	Destination	Protection Policy	Relationship Health	Relationship Status	Lag
svm_onPremorhel2_u03	svm_hybridcvoorhel2_u03_dr	MirrorAllSnapshots	Healthy	Mirrored	12 seconds

## 6. Add CVO database storage SVM to SnapCenter


- 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.



- Click More Options to open additional storage configuration options. In the Platform field, select Cloud Volumes ONTAP, check Secondary, and then click Save.



- Assign the storage systems to SnapCenter database management user IDs as shown in 3. [SnapCenter host plugin installation](#).




## 7. Setup database backup policy in SnapCenter

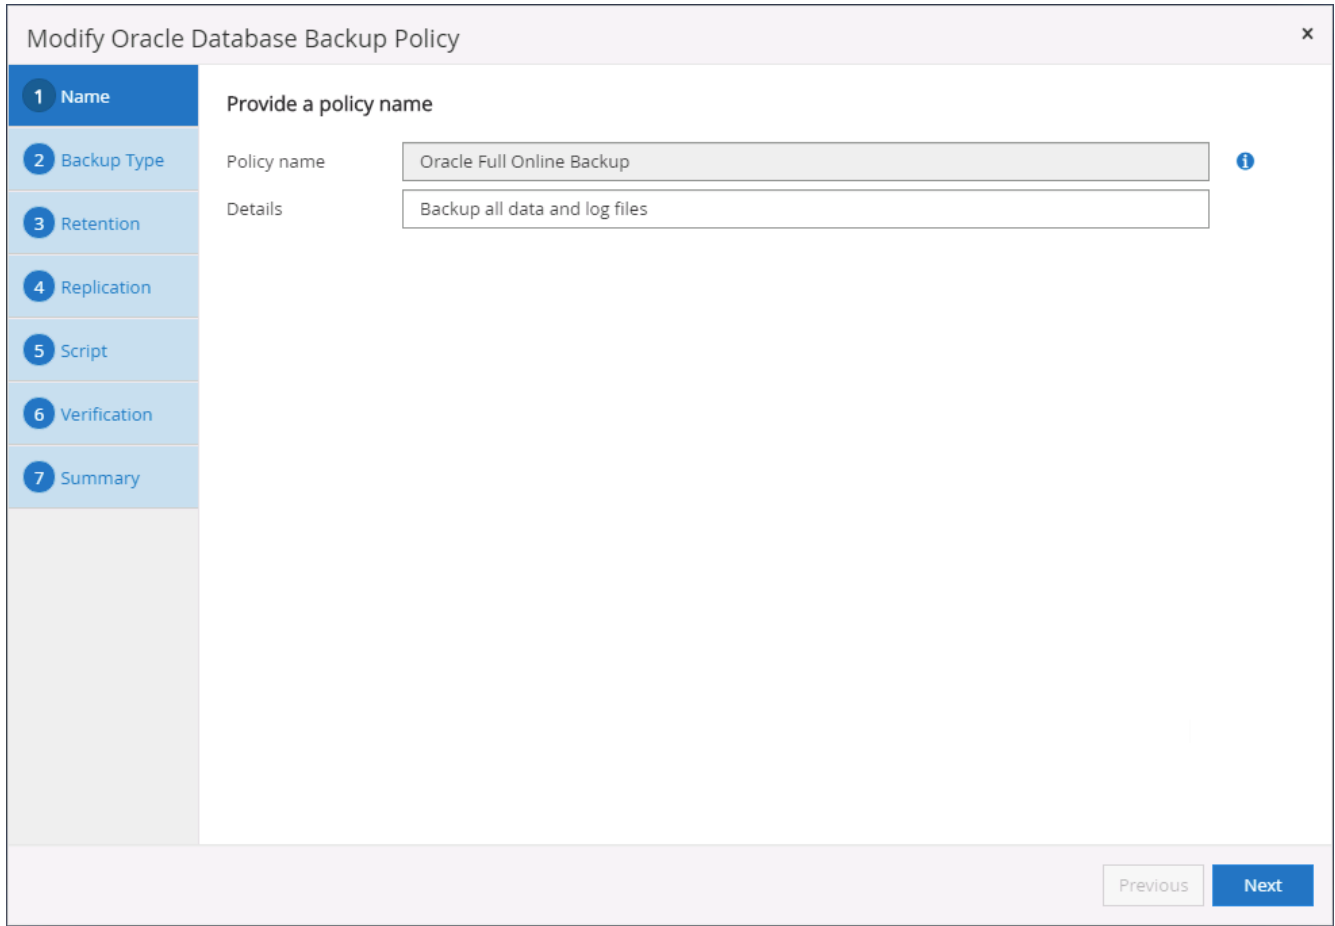
The following procedures demonstrates how to create a full database or log file backup policy. The policy can then be implemented to protect databases resources. The recovery point objective (RPO) or recovery time objective (RTO) dictates the frequency of database and/or log backups.

Create a full database backup policy for Oracle

- 1. Log into SnapCenter as a database management user ID, click Settings, and then click Policies.



- 2. Click New to launch a new backup policy creation workflow or choose an existing policy for modification.



- 3. Select the backup type and schedule frequency.

Modify Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select Oracle database backup options

Choose backup type

☒ Online backup

☒ Datafiles, control files, and archive logs

☐ Datafiles and control files

☐ Archive logs

☐ Offline backup 

i

☒ Mount

☐ Shutdown

☐ Save state of PDBs 

i

Choose schedule frequency

Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling you to stagger your start times.

☐ On demand

☐ Hourly

☒ Daily

Previous

Next

4. Set the backup retention setting. This defines how many full database backup copies to keep.

130



Modify Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Retention settings

Daily retention settings

Data backup retention settings

Total Snapshot copies to keep

7

Keep Snapshot copies for

14

days

Archive Log backup retention settings

Total Snapshot copies to keep

7

Keep Snapshot copies for

14

days

Previous

Next

5. Select the secondary replication options to push local primary snapshots backups to be replicated to a secondary location in cloud.

Modify Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select secondary replication options

☒ Update SnapMirror after creating a local Snapshot copy.

☐ Update SnapVault after creating a local Snapshot copy.

Secondary policy label

Daily

Error retry count

3

Previous

Next

6. Specify any optional script to run before and after a backup run.

Modify Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Specify optional scripts to run before and after performing a backup job

Prescript full path

/var/opt/snapcenter/spl/scripts/

Enter Prescript path

Prescript arguments

Postscript full path

/var/opt/snapcenter/spl/scripts/

Enter Postscript path

Postscript arguments

Script timeout

60

secs

Previous

Next

7. Run backup verification if desired.

133

Modify Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select the options to run backup verification

Run Verifications for following backup schedules

Select how often you want the schedules to occur in the policy. The specific verification times are set at backup job creation enabling you to stagger your verification start times.

☐ Daily

Verification script commands

Script timeout

60

secs

Prescript full path

/var/opt/snapcenter/spl/scripts/

Enter Prescript path

Prescript arguments

Choose optional arguments...

Postscript full path

/var/opt/snapcenter/spl/scripts/

Enter Postscript path

Postscript arguments

Choose optional arguments...

Previous

Next

8. Summary.

1

Name

2

Backup Type

3

Retention

4

Replication

5

Script

6

Verification

7

Summary

Summary

Policy name	Oracle Full Online Backup
Details	Backup all data and log files
Backup type	Online backup
Schedule type	Daily
RMAN catalog backup	Disabled
Archive log pruning	None
On demand data backup retention	None
On demand archive log backup retention	None
Hourly data backup retention	None
Hourly archive log backup retention	None
Daily data backup retention	Delete Snapshot copies older than : 14 days
Daily archive log backup retention	Delete Snapshot copies older than : 14 days
Weekly data backup retention	None
Weekly archive log backup retention	None
Monthly data backup retention	None
Monthly archive log backup retention	None
Replication	SnapMirror enabled , Secondary policy label: Daily , Error retry count: 3

Previous

Finish

## Create a database log backup policy for Oracle

1. Log into SnapCenter with a database management user ID, click Settings, and then click Policies.
2. Click New to launch a new backup policy creation workflow, or choose an existing policy for modification.

New Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Provide a policy name

Policy name

Details

Oracle Archive Log Backup

Backup Oracle archive logs

Previous

Next

3. Select the backup type and schedule frequency.

136

New Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select Oracle database backup options

Choose backup type

☒ Online backup

☐ Datafiles, control files, and archive logs

☐ Datafiles and control files

☒ Archive logs

☐ Offline backup 

i

☒ Mount

☐ Shutdown

☐ Save state of PDBs 

i

Choose schedule frequency

Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling you to stagger your start times.

☐ On demand

☒ Hourly

☐ Daily

Previous

Next

4. Set the log retention period.

137

New Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Retention settings

Hourly retention settings

Data backup retention settings

Total Snapshot copies to keep

7

Keep Snapshot copies for

14 days

Archive Log backup retention settings

Total Snapshot copies to keep

7

Keep Snapshot copies for

7 days

Previous

Next

5. Enable replication to a secondary location in the public cloud.

138



New Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select secondary replication options

☒ Update SnapMirror after creating a local Snapshot copy.

☐ Update SnapVault after creating a local Snapshot copy.

Secondary policy label

Hourly

Error retry count

3

Previous

Next

6. Specify any optional scripts to run before and after log backup.

New Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Specify optional scripts to run before and after performing a backup job

Prescript full path

/var/opt/snapcenter/spl/scripts/

Enter Prescript path

Prescript arguments

Postscript full path

/var/opt/snapcenter/spl/scripts/

Enter Postscript path

Postscript arguments

Script timeout

60

secs

Previous

Next

7. Specify any backup verification scripts.

140

New Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select the options to run backup verification

Run Verifications for following backup schedules

Select how often you want the schedules to occur in the policy. The specific verification times are set at backup job creation enabling you to stagger your verification start times.

Verification script commands

Script timeout

60secs

Prescript full path

/var/opt/snapcenter/spl/scripts/

Enter Prescript path

Prescript arguments

Choose optional arguments...

Postscript full path

/var/opt/snapcenter/spl/scripts/

Enter Postscript path

Postscript arguments

Choose optional arguments...

Previous

Next

8. Summary.

1

Name

2

Backup Type

3

Retention

4

Replication

5

Script

6

Verification

7

Summary

Summary


Policy name	Oracle Archive Log Backup
Details	Backup Oracle archive logs
Backup type	Online backup
Schedule type	Hourly
RMAN catalog backup	Disabled
Archive log pruning	None
On demand data backup retention	None
On demand archive log backup retention	None
Hourly data backup retention	None
Hourly archive log backup retention	Delete Snapshot copies older than : 7 days
Daily data backup retention	None
Daily archive log backup retention	None
Weekly data backup retention	None
Weekly archive log backup retention	None
Monthly data backup retention	None
Monthly archive log backup retention	None
Replication	SnapMirror enabled , Secondary policy label: Hourly , Error retry count: 3

Previous

Finish

## Create a full database backup policy for SQL

1. Log into SnapCenter with a database management user ID, click Settings, and then click Policies.



2. Click New to launch a new backup policy creation workflow, or choose an existing policy for modification.

New SQL Server Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Provide a policy name

Policy name

SQL Server Full Backup

Details

Backup all data and log files

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 Server Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select SQL server backup options

Choose backup type

☒ Full backup and log backup

☐ Full backup

☐ Log backup

☐ Copy only backup

Maximum databases backed up per Snapshot copy:

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.

144

New SQL Server Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Retention settings

Retention settings for up-to-the-minute restore operation ⓘ

☒ Keep log backups applicable to last

7

full backups

☐ Keep log backups applicable to last

14

days

Full backup retention settings ⓘ

Daily

☒ Total Snapshot copies to keep

7

☐ Keep Snapshot copies for

14

days

Previous

Next

5. Enable backup copy replication to a secondary location in cloud.

New SQL Server Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select secondary replication options

☒ Update SnapMirror after creating a local Snapshot copy.

☐ Update SnapVault after creating a local Snapshot copy.

Secondary policy label

Daily

Error retry count

3

Previous

Next

6. Specify any optional scripts to run before or after a backup job.



New SQL Server Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Specify optional scripts to run before performing a backup job

Prescript full path

Prescript arguments

Choose optional arguments...

Specify optional scripts to run after performing a backup job

Postscript full path

Postscript arguments

Choose optional arguments...

Script timeout

60

secs

Previous

Next

7. Specify the options to run backup verification.

New SQL Server Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select the options to run backup verification

Run verifications for the following backup schedules

Select how often you want the schedules to occur in the policy. The specific verification times are set at backup job creation enabling you to stagger your verification start times.

☐ Daily

Database consistency checks options

☒ Limit the integrity structure to physical structure of the database (PHYSICAL\_ONLY)

☒ Suppress all information message (NO\_INFOMSGS)

☐ Display all reported error messages per object (ALL\_ERRORMSGSGS)

☐ Do not check non-clustered indexes (NOINDEX)

☐ Limit the checks and obtain the locks instead of using an internal database Snapshot copy (TABLOCK)

Log backup

☐ Verify log backup.

Verification script settings

Script timeout  secs

Previous

Next

8. Summary.

New SQL Server Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Summary

Policy name	SQL Server Full Backup
Details	Backup all data and log files
Backup type	Full backup and log backup
Availability group settings	Backup only on preferred backup replica
Schedule Type	Daily
UTM retention	Total backup copies to retain : 7
Daily Full backup retention	Total backup copies to retain : 7
Replication	SnapMirror enabled , Secondary policy label: Daily , Error retry count: 3
Backup prescript settings	undefined Prescript arguments:
Backup postscript settings	undefined Postscript arguments:
Verification for backup schedule type	none
Verification prescript settings	undefined Prescript arguments:
Verification postscript settings	undefined Postscript arguments:

Previous

Finish

## Create a database log backup policy for SQL.

1. Log into SnapCenter with a database management user ID, click Settings > Policies, and then New to launch a new policy creation workflow.

New SQL Server Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Provide a policy name

Policy name

Details

SQL Server Log Backup

Backup SQL server log

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.

150

New SQL Server Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select SQL server backup options

Choose backup type

☐ Full backup and log backup

☐ Full backup

☒ Log backup

☐ Copy only backup

Maximum databases backed up per Snapshot copy:

100

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

3. SQL server data backup policy defines the log backup retention; accept the defaults here.

151

New SQL Server Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Log backup retention settings

Up-to-the-minute (UTM) retention settings retains log backups created as part of full backup and full and log backup operations. UTM retention settings also decides for how many full backups the log backups are to be retained. For example, if UTM retention settings is configured to retain log backups of the last 5 full backups, then the log backups of the last 5 full backups are retained and the rest are deleted.

Previous

Next

4. Enable log backup replication to secondary in the cloud.

New SQL Server Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select secondary replication options

☒ Update SnapMirror after creating a local Snapshot copy.

☐ Update SnapVault after creating a local Snapshot copy.

Secondary policy label

Hourly

Error retry count

3

Previous

Next

5. Specify any optional scripts to run before or after a backup job.

New SQL Server Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Specify optional scripts to run before performing a backup job

Prescript full path

Prescript arguments

Specify optional scripts to run after performing a backup job

Postscript full path

Postscript arguments

Script timeout

Choose optional arguments...

Choose optional arguments...

60secs

Previous

Next

6. Summary.



1

Name

2

Backup Type

3

Retention

4

Replication

5

Script

6

Verification

7

Summary

Summary

Policy name	SQL Server Log Backup
Details	Backup SQL server log
Backup type	Log transaction backup
Availability group settings	Backup only on preferred backup replica
Schedule Type	Hourly
Replication	SnapMirror enabled , Secondary policy label: Hourly , Error retry count: 3
Backup prescript settings	undefined Prescript arguments:
Backup postscript settings	undefined Postscript arguments:
Verification for backup schedule type	none
Verification prescript settings	undefined Prescript arguments:
Verification postscript settings	undefined Postscript arguments:

Previous

Finish

## 8. Implement backup policy to protect database

SnapCenter uses a resource group to backup a database in a logical grouping of database resources, such as multiple databases hosted on a server, a database sharing the same storage volumes, multiple databases supporting a business application, and so on. Protecting a single database creates a resource group of its own. The following procedures demonstrate how to implement a backup policy created in section 7 to protect Oracle and SQL Server databases.

### Create a resource group for full backup of Oracle

- 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 SnapCenter®							
Oracle Database							
View Database Search databases							
Dashboard	Resources	Monitor	Reports	Hosts	Storage Systems	Settings	Alerts
	Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Overall Status
	cdb2	Single Instance (Multitenant)	rhe12.demo.netapp.com				Not protected

- 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.

NetApp SnapCenter®

Oracle Database

Search databases

Name
cdb2

1 Name 2 Resources 3 Policies 4 Verification 5 Notification 6 Summary

Provide a name and tags for the resource group

Name: rhe12\_cdb2

Tags: orafulbkup

☒ Use custom name format for Snapshot copy

\$CustomText: rhe12\_cdb2

Backup settings

Exclude archive log destinations from backup: [ ]

3. Add database resources to the resource group.

NetApp SnapCenter®

Oracle Database

Search databases

Name
cdb2

1 Name 2 Resources 3 Policies 4 Verification 5 Notification 6 Summary

Add resources to Resource Group

Host: All

Available Resources

search available resources

Selected Resources

cdb2 (rhe12.demonetapp.com)

4. Select a full backup policy created in section 7 from the drop-down list.

NetApp SnapCenter®

Oracle Database

Search databases

Name
cdb2

1 Name 2 Resources 3 Policies 4 Verification 5 Notification 6 Summary

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

5. Click the (+) sign to configure the desired backup schedule.



## 8. Summary.

## Create a resource group for log backup of Oracle

1. Log into SnapCenter with a database management user ID, and navigate to the Resources tab. In the View drop-down list, choose either Database or Resource Group to launch the resource group creation workflow.

Name	Resources	Tags	Policies	Last Backup	Overall Status
rhei2_cdb2	1	orafullbkup	Oracle Full Online Backup		

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.

NetApp SnapCenter®

Oracle Database

Search resource groups

Name

rhel2\_cdb2

Provide a name and tags for the resource group

Name: rhel2\_cdb2\_log

Tags: oragbkup

☒ Use custom name format for Snapshot copy

\$CustomText: rhel2\_cdb2\_log

Backup settings

Exclude archive log destinations from backup: ☒

3. Add database resources to the resource group.

NetApp SnapCenter®

Oracle Database

Search resource groups

Name

rhel2\_cdb2

Add resources to Resource Group

Host: All

Available Resources

search available resources

Selected Resources

cdb2 (rhel2.demo.netapp.com)

Total 1

Previous Next

4. Select a log backup policy created in section 7 from the drop-down list.

NetApp SnapCenter®

Oracle Database

Search resource groups

Name

rhel2\_cdb2

Select one or more policies and configure schedules

Oracle Archive Log Backup

Oracle Full Online Backup

Oracle Archive Log Backup

Policy: Oracle Archive Log Backup

Applied Schedules: None

Configure Schedules: +

Total 1

Previous 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

☒ Expires on

12/31/2021 3:00 PM

Repeat every

1

hours

0

mins

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

Cancel

OK

6. If backup verification is configured, it displays here.

NetApp SnapCenter®
demolora@ba
App Backup and Clone Admin
Sign Out

Oracle Database

Search resource groups

Name

rhel2\_cdb2

Total 1

New Resource Group

1

2

3

4

5

6

Name

Resources

Policies

Verification

Notification

Summary

Configure verification schedules

Policy

Schedule Type

Applied Schedules

Configure Schedules

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

Total 0

Previous

Next

7. Configure an SMTP server for email notification if desired.

NetApp SnapCenter®

Oracle Database

Search resource groups

Name

rhel2\_cdb2

Total 1

New Resource Group

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 SMTP server.

1 Name 2 Resources 3 Policies 4 Verification 5 Notification 6 Summary

Provide email settings ⓘ

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

Email preference: Never

From: From email

To: Email to

Subject: Notification

☐ Attach job report

Previous Next

## 8. Summary.

NetApp SnapCenter®

Oracle Database

Search resource groups

Name

rhel2\_cdb2

Total 1

New Resource Group

1 Name 2 Resources 3 Policies 4 Verification 5 Notification 6 Summary

Resource group name: rhel2\_cdb2\_log

Tags: oralogbkup

Policy: Oracle Archive Log Backup: Hourly

Plug-in: SnapCenter Plug-in for Oracle Database


Verification enabled for policy: None

Send email: No


Previous Finish

## Create a resource group for full backup of SQL Server

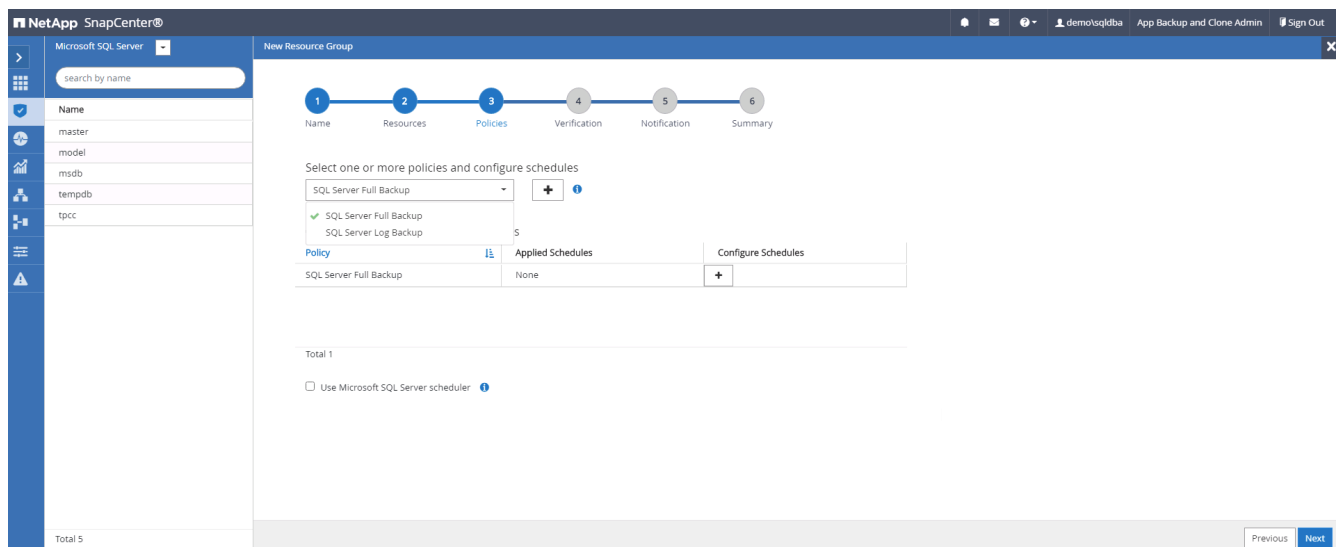
1. Log into SnapCenter with a database management user ID, and navigate to the Resources tab. In the View drop-down list, choose either a Database or Resource Group to launch the resource group creation workflow. Provide a name and tags for the resource group. You can define a naming format for the Snapshot copy.



2. Select the database resources to be backed up.



3. Select a full SQL backup policy created in section 7.





4. Add exact timing for backups as well as the frequency.

Add schedules for policy SQL Server Full Backup

Daily

Start date 09/10/2021 6:20 PM

☒ Expires on 12/31/2021 6:20 PM

Repeat every 1 days

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

Cancel OK

5. Choose the verification server for the backup on secondary if backup verification is to be performed. Click Load Locator to populate the secondary storage location.

NetApp SnapCenter

Microsoft SQL Server

New Resource Group

1 Name 2 Resources 3 Policies 4 Verification 5 Notification 6 Summary

Select the verification servers

Verification server Select one or more servers

Load secondary locators to verify backups on secondary Load locators

Secondary storage location: SnapVault or SnapMirror

Source Volume	Destination Volume
svm_onPremsql1_data	svm_hybridcvosql1_data_dr
svm_onPremsql1_log	svm_hybridcvosql1_log_dr

Configure verification schedules

Policy Schedule Type Applied Schedules Configure Schedules

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

Previous Next

6. Configure the SMTP server for email notification if desired.

NetApp SnapCenter®

Microsoft SQL Server

search by name

Name

master

model

msdb

tempdb

tpcc

Total 5

New Resource Group

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 SMTP server.

1 Name 2 Resources 3 Policies 4 Verification 5 Notification 6 Summary

Provide email settings

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

Email preference: Never

From: From email

To: Email to

Subject: Notification

☐ Attach job report

Previous Next

## 7. Summary.

NetApp SnapCenter®

Microsoft SQL Server

search by name

Name

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

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.

New Resource Group


1 Name 2 Resources 3 Policies 4 Verification 5 Notification 6 Summary

Resource group name	sql1_tpcc
Tags	sqlfullbkup
Policy	SQL Server Full Backup: Daily
Plug-in	SnapCenter Plug-in for Microsoft SQL Server
Verification Server	None
Verification enabled for policy	None
Send email	No


Previous 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.



2. Select the database resources to be backed up.



3. Select a SQL log backup policy created in section 7.



4. Add exact timing for the backup as well as the frequency.

The screenshot shows the NetApp SnapCenter console for a 'New Resource Group'. The left sidebar lists 'Microsoft SQL Server' and 'sql1\_tpcc'. The main area is titled 'New Resource Group' and shows a progress bar with steps: 1 Name, 2 Resources, 3 Policies (active), 4 Verification, 5 Notification, and 6 Summary. Below the progress bar, there is a section 'Select one or more policies and configure schedules' with a dropdown menu showing 'SQL Server Log Backup'. Below this is a table 'Configure schedules for selected policies' with columns 'Policy', 'Applied Schedules', and 'Configure Schedules'. The table has one row: 'SQL Server Log Backup' with 'Hourly: Repeat every 1 hours' in the 'Applied Schedules' column. At the bottom, there is a 'Total 1' summary and a checkbox 'Use Microsoft SQL Server scheduler' which is unchecked. Navigation buttons 'Previous' and 'Next' are at the bottom right.

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.

The screenshot shows the NetApp SnapCenter console for a 'New Resource Group'. The left sidebar lists 'Microsoft SQL Server' and 'sql1\_tpcc'. The main area is titled 'New Resource Group' and shows a progress bar with steps: 1 Name, 2 Resources, 3 Policies, 4 Verification (active), 5 Notification, and 6 Summary. Below the progress bar, there is a section 'Select the verification servers' with a dropdown menu 'Verification server: Select one or more servers'. Below this is a section 'Load secondary locators to verify backups on secondary' with a 'Load locators' button. Below this is a section 'Secondary storage location: SnapVault or SnapMirror' with a table 'Source Volume' and 'Destination Volume'. The table has two rows: 'svm\_onPrem:sql1\_data' with 'svm\_hybridv:sql1\_data\_dr' in the 'Destination Volume' column, and 'svm\_onPrem:sql1\_log' with 'svm\_hybridv:sql1\_log\_dr' in the 'Destination Volume' column. Below the table is a section 'Configure verification schedules' with tabs 'Policy', 'Schedule Type' (active), 'Applied Schedules', and 'Configure Schedules'. The 'Schedule Type' tab shows 'There is no match for your search or data is not available.' Navigation buttons 'Previous' and 'Next' are at the bottom right.

6. Configure the SMTP server for email notification if desired.

NetApp SnapCenter®

Microsoft SQL Server

search by name

sql1\_tpcc

Total 1

New Resource Group

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 SMTP server.

1 Name 2 Resources 3 Policies 4 Verification 5 Notification 6 Summary

Provide email settings

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

Email preference: Never

From: From email

To: Email to

Subject: Notification

☐ Attach job report

Previous Next

## 7. Summary.

NetApp SnapCenter®

Microsoft SQL Server

search by name

sql1\_tpcc

Total 1

New Resource Group

1 Name 2 Resources 3 Policies 4 Verification 5 Notification 6 Summary

Resource group name: sql1\_tpcc\_log

Tags: sqllogbkup

Policy: SQL Server Log Backup: Hourly

Plug-in: SnapCenter Plug-in for Microsoft SQL Server

Verification Server: None

Verification enabled for policy: None

Send email: No

Previous 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 SnapCenter®

Jobs Schedules Events Logs

search by name

Jobs - Filter

ID	Status	Name	Start date	End date	Owner
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 PM	demo:sqldba
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 PM	demo: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 PM	demo:sqldba
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 PM	demo:sqldba
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 PM	demo:sqldba
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 PM	demo: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 PM	demo: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 PM	demo:sqldba

Go to the Resources tab, click the database name to view details of database backup, and toggle between Local copies and mirror copies to verify that Snapshot backups are replicated to a secondary location in the

public cloud.

The screenshot shows the NetApp SnapCenter web interface. On the left is a sidebar with navigation icons. The main area is titled 'cdb2 Topology' and 'Manage Copies'. It displays a diagram with 'Local copies' (197 Backups, 0 Clones) and 'Mirror copies' (197 Backups, 3 Clones). A 'Summary Card' on the right shows: 394 Backups, 28 Data Backups, 366 Log Backups, and 3 Clones. Below this is a table of 'Primary Backup(s)' with columns: Backup Name, Count, Type, End Date, Verified, Mounted, RMAN Cataloged, and SCN. The table lists five backup entries with their respective details.

Backup Name	Count	Type	End Date	Verified	Mounted	RMAN Cataloged	SCN
rhel2_cdb2_09-23-2021_14.35.03.3242_1	1	Log	09/23/2021 2:35:45 PM	Not Applicable	False	Not Cataloged	6872761
rhel2_cdb2_09-23-2021_14.35.03.3242_0	1	Data	09/23/2021 2:35:30 PM	Unverified	False	Not Cataloged	6872715
rhel2_cdb2_09-22-2021_14.35.02.0014_1	1	Log	09/22/2021 2:35:24 PM	Not Applicable	False	Not Cataloged	6737479
rhel2_cdb2_09-22-2021_14.35.02.0014_0	1	Data	09/22/2021 2:35:14 PM	Unverified	False	Not Cataloged	6737395
rhel2_cdb2_09-21-2021_14.35.02.1884_1	1	Log	09/21/2021 2:35:35 PM	Not Applicable	False	Not Cataloged	6598735

At this point, database backup copies in the cloud are ready to clone to run dev/test processes or for disaster recovery in the event of a primary failure.

## Getting Started with AWS public cloud

This section describes the process of deploying Cloud Manager and Cloud Volumes ONTAP in AWS.

### AWS public cloud



To make things easier to follow, we have created this document based on a deployment in AWS. However, the process is very similar for Azure and GCP.

## 1. Pre-flight check

Before deployment, make sure that the infrastructure is in place to allow for the deployment in the next stage. This includes the following:

- ☐ AWS account
- ☐ VPC in your region of choice
- ☐ Subnet with access to the public internet
- ☐ Permissions to add IAM roles into your AWS account
- ☐ A secret key and access key for your AWS user

## 2. Steps to deploy Cloud Manager and Cloud Volumes ONTAP in AWS



There are many methods for deploying Cloud Manager and Cloud Volumes ONTAP; this method is the simplest but requires the most permissions. If this method is not appropriate for your AWS environment, please consult the [NetApp Cloud Documentation](#).

### Deploy the Cloud Manager connector

1. Navigate to [NetApp Cloud Central](#) and log in or sign up.



[Continue to Cloud Manager](#)

## Log In to NetApp Cloud Central


---

Don't have an account yet? [Sign Up](#)


**LOGIN**

[Forgot your password?](#)

2. After you log in, you should be taken to the Canvas.




3. Click "Add Working Environment" and choose Cloud Volumes ONTAP in AWS. Here, you also choose whether you want to deploy a single node system or a high availability pair. I have chosen to deploy a high availability pair.




4. If no connector has been created, a pop-up appears asking you to create a connector.





5. Click Lets Start, and then choose AWS.



6. Enter your secret key and access key. Make sure that your user has the correct permissions outlined on the [NetApp policies page](#).

Cloud Manager

Add Connector

Need Help? X

Get Ready AWS Credentials Details Network Security Group Review

### AWS Credentials

AWS Access Key

AWS Access Key is required

AWS Secret Key

Region

us-east-1 | US East (N. Virginia)

Want to launch an instance without AWS Credentials?

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 Help? X

Get Ready AWS Credentials Details Network Security Group Review

### Details

Connector Instance Name

awscloudmanager

Connector Role

Create Role Select an existing Role

Role Name

Cloud-Manager-Operator-IBHt24j

Add Tags to Connector Instance

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

**Add Connector**

Get Ready AWS Credentials Details **Network** Security Group Review

**Connectivity**

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 Configuration (Optional)**

HTTP Proxy: Example: http://172.16.254.1:8080

Define Credentials for this Proxy

Upload a root certificate

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.

**Add Connector**

Get Ready AWS Credentials Details Network **Security Group** Review

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)

Source Type: My IP

Source (CIDR): 216.240.31.145/32

**HTTPS** (Port 443)

Source Type: My IP

Source (CIDR): 216.240.31.145/32

**SSH** (Port 22)

Source Type: My IP

Source (CIDR): 216.240.31.145/32

Previous Next

10. Review the information on the summary page and click Add to deploy the connector.

**Add Connector**

Get Ready AWS Credentials Details Network Security Group **Review**

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

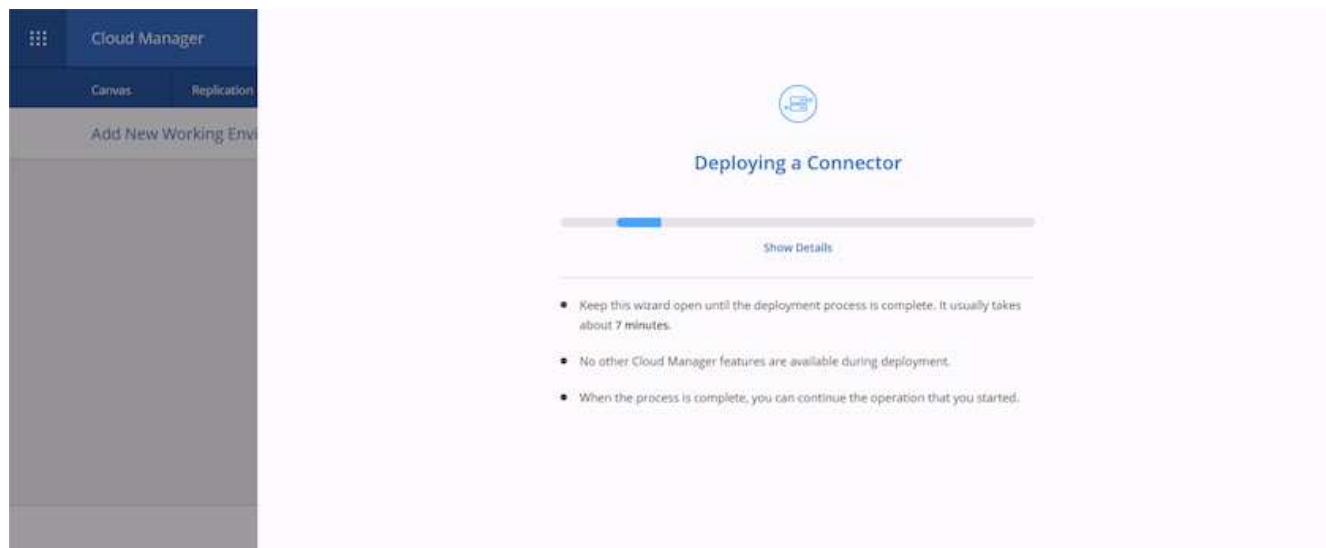
Public IP: Enable

Proxy: None


Security Group: HTTP: 216.240.31.145/32, HTTPS: 216.240.31.145/32, SSH: 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.

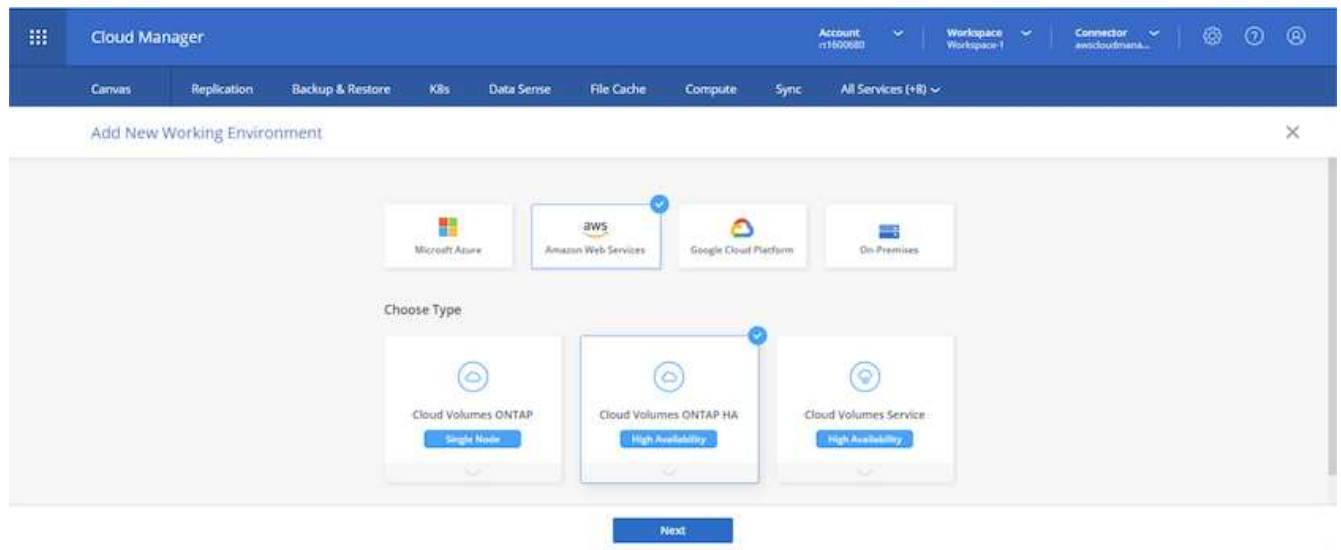


12. When the deployment is complete, a success page appears.

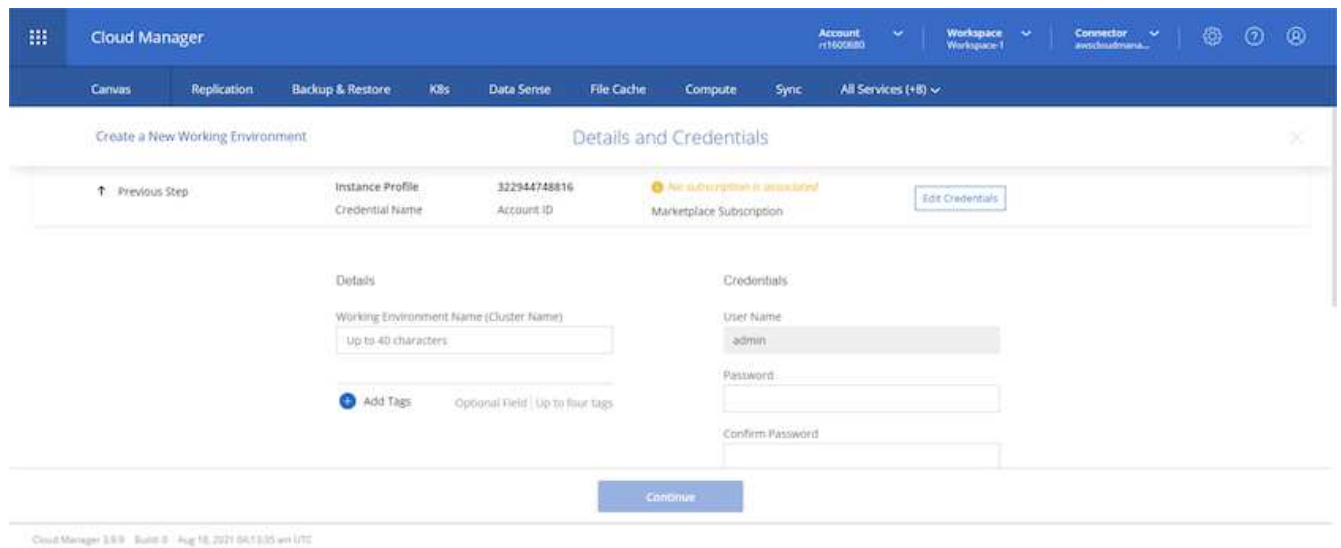


## Deploy Cloud Volumes ONTAP

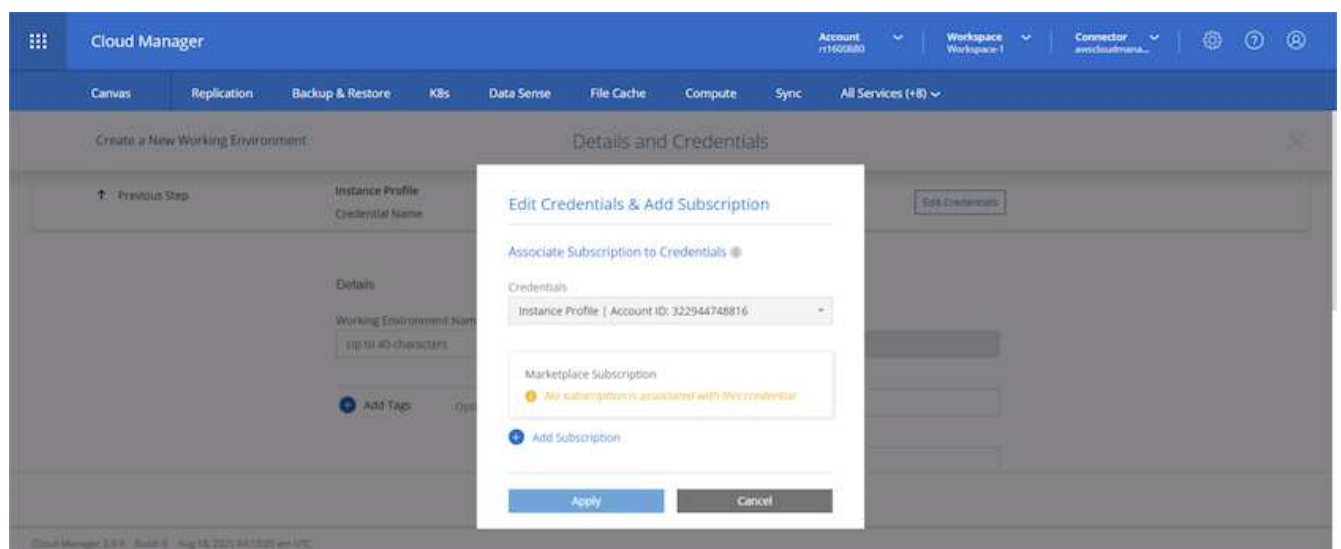
1. Select AWS and the type of deployment based on your requirements.




2. If no subscription has been assigned and you wish to purchase with PAYGO, choose Edit Credentials.




3. Choose Add Subscription.




4. Choose the type of contract that you wish to subscribe to. I chose Pay-as-you-go.




5. You are redirected to AWS; choose Continue to Subscribe.




6. Subscribe and you are redirected back to NetApp Cloud Central. If you have already subscribed and don't get redirected, choose the "Click here" link.




7. You are redirected to Cloud Central where you must name your subscription and assign it to your Cloud Central account.



8. When successful, a check mark page appears. Navigate back to your Cloud Manager tab.



9. The subscription now appears in Cloud Central. Click Apply to continue.



10. Enter the working environment details such as:

- a. Cluster name
- b. Cluster password
- c. AWS tags (Optional)



The screenshot shows the 'Details and Credentials' step in the NetApp Cloud Manager interface. The top navigation bar includes 'Cloud Manager' and various service tabs like 'Canvas', 'Replication', 'Backup & Restore', etc. The main content area is titled 'Create a New Working Environment' and 'Details and Credentials'. It features a 'Previous Step' button and a table with instance profile details. Below this, there are two sections: 'Details' and 'Credentials'. The 'Details' section has a text input for 'Working Environment Name (Cluster Name)' with the value 'hybridawsco'. The 'Credentials' section has inputs for 'User Name' (admin), 'Password' (masked), and 'Confirm Password' (masked). An 'Add Tags' button is also present. A 'Continue' button is at the bottom.

Instance Profile	322944748816	demo.netapp.com-cloud-vol...
Credential Name	Account ID	Marketplace Subscription

Details

Working Environment Name (Cluster Name)

hybridawsco

+ Add Tags Optional Field | Up to four tags

Credentials

User Name

admin

Password

\*\*\*\*\*

Confirm Password

\*\*\*\*\*

Continue

Cloud Manager 3.9.9 Built 0 Aug 18, 2021 06:13:35 am UTC

- Choose which additional services you would like to deploy. To discover more about these services, visit the [NetApp Cloud Homepage](#).

The screenshot shows the 'Services' step in the NetApp Cloud Manager interface. The top navigation bar is the same as the previous screenshot. The main content area is titled 'Create a New Working Environment' and 'Services'. It features a 'Previous Step' button and a list of services with toggle switches and dropdown menus. The services listed are 'Data Sense & Compliance', 'Backup to Cloud', and 'Monitoring'. All three services have their toggle switches turned on. A 'Continue' button is at the bottom.

Previous Step

Data Sense & Compliance


Backup to Cloud

Monitoring

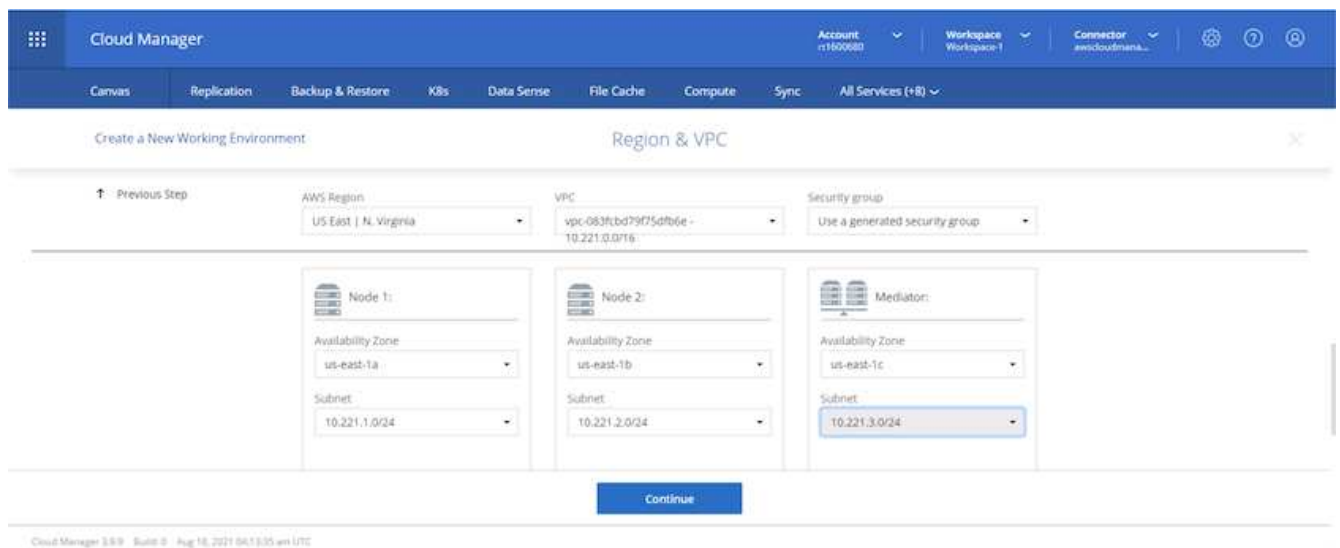
Continue

Cloud Manager 3.9.9 Built 0 Aug 18, 2021 06:13:35 am UTC

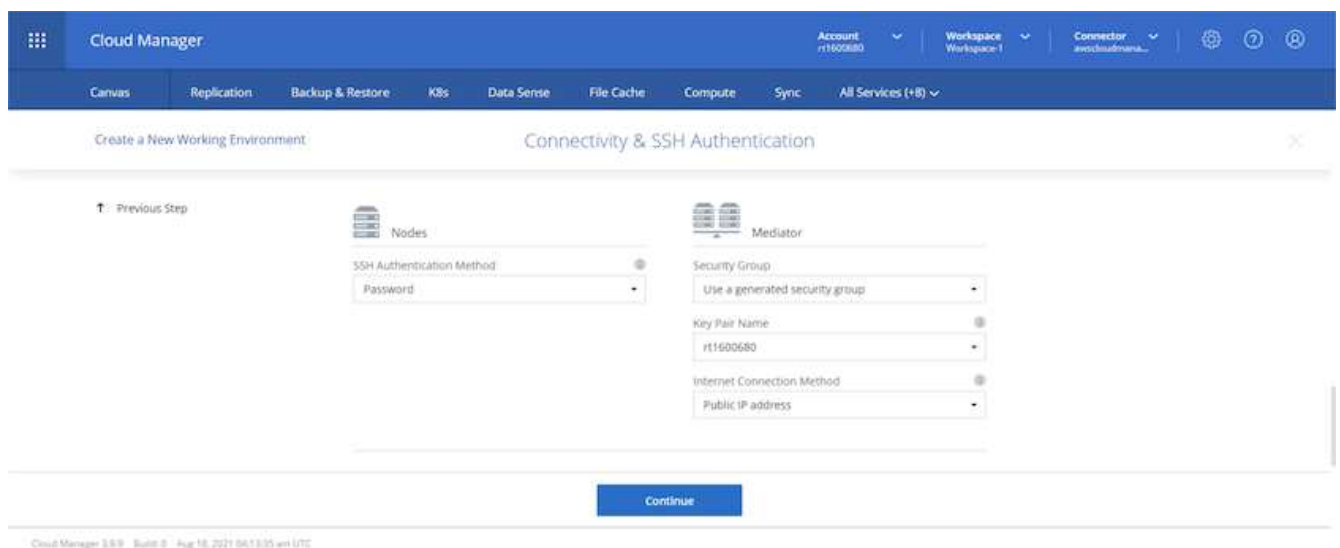
- Choose whether to deploy in multiple availability zones (requires three subnets, each in a different AZ), or a single availability zone. I chose multiple AZs.



13. Choose the region, VPC, and security group for the cluster to be deployed into. In this section, you also assign the availability zones per node (and mediator) as well as the subnets that they occupy.



14. Choose the connection methods for the nodes as well as the mediator.





The mediator requires communication with the AWS APIs. A public IP address is not required so long as the APIs are reachable after the mediator EC2 instance has been deployed.

1. Floating IP addresses are used to allow access to the various IP addresses that Cloud Volumes ONTAP uses, including cluster management and data serving IPs. These must be addresses that are not already routable within your network and are added to route tables in your AWS environment. These are required to enable consistent IP addresses for an HA pair during failover. More information about floating IP addresses can be found in the [NetApp Cloud Documentation](#).

The screenshot shows the 'Floating IPs' configuration step in the Cloud Manager console. The page title is 'Floating IPs'. Below the title, there is a 'Previous Step' link. The main content area contains instructions: 'Floating IP addresses are required for cluster and SVM access and for NFS and CIFS data access. These floating IPs can migrate between HA nodes if failures occur. To access the data from outside the VPC, you can set up an AWS transit gateway. You must specify IP addresses that are outside of the CIDR blocks for all VPCs in the selected AWS region.' There are four input fields for IP addresses: 'Floating IP address for cluster management' (10.222.0.200), 'Floating IP address 1 for NFS and CIFS data' (10.222.0.201), 'Floating IP address 2 for NFS and CIFS data' (10.222.0.202), and 'Floating IP address for SVM management (Optional)' (Enter Floating IP Address). A 'Continue' button is at the bottom.


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.

The screenshot shows the 'Route Tables' configuration step in the Cloud Manager console. The page title is 'Route Tables'. Below the title, there is a 'Previous Step' link. The main content area contains instructions: 'Select the route tables that should include routes to the floating IP addresses. This enables client access to the Cloud Volumes ONTAP HA pair. If you leave a route table unselected, clients that are associated with the route table cannot access the HA pair.' There is a table with the following data:


Name	Main	ID	Associate with Subnet	Tags
private_rt_11600680	No	rtb-08b4cb88f5c826a5	3 Subnets	1 Tags
public_rt_11600680	Yes	rtb-0e46720d0da10c593	1 Subnets	1 Tags

Below the table, it says '2 Route Tables | The main route table is the default for the VPC'. A 'Continue' button is at the bottom.


3. Choose whether to enable AWS managed encryption or AWS KMS to encrypt the ONTAP root, boot, and data disks.



4. Choose your licensing model. If you don't know which to choose, contact your NetApp representative.



5. Select which configuration best suits your use case. This is related to the sizing considerations covered in the prerequisites page.




6. Optionally, create a volume. This is not required, because the next steps use SnapMirror, which creates the volumes for us.

The screenshot shows the 'Create Volume' interface in the Cloud Manager console. The top navigation bar includes 'Cloud Manager', 'Account: r1160d880', 'Workspace: Workspace 1', and 'Connector: awscloudmana...'. The main menu has tabs for 'Canvas', 'Replication', 'Backup & Restore', 'K8s', 'Data Sense', 'File Cache', 'Compute', 'Sync', and 'All Services (+8)'. The breadcrumb trail is 'Create a New Working Environment' > 'Create Volume'. The interface is divided into two main sections: 'Details & Protection' and 'Protocol'. In 'Details & Protection', there is a 'Volume Name' field, a 'Size (GB)' field with a 'Volume size' button, a 'Snapshot Policy' dropdown set to 'default', and a 'Default Policy' button. In the 'Protocol' section, there are three tabs: 'NFS', 'CIFS', and 'iSCSI'. The 'NFS' tab is selected, showing 'Access Control' as 'Custom export policy' and a 'Custom export policy' field with the value '10.221.0.0/16'. There is also an 'Advanced options' dropdown. At the bottom, there are 'Continue' and 'Skip' buttons. The footer shows 'Cloud Manager 5.8.9 Build 9 Aug 18, 2021 04:13:35 am UTC'.

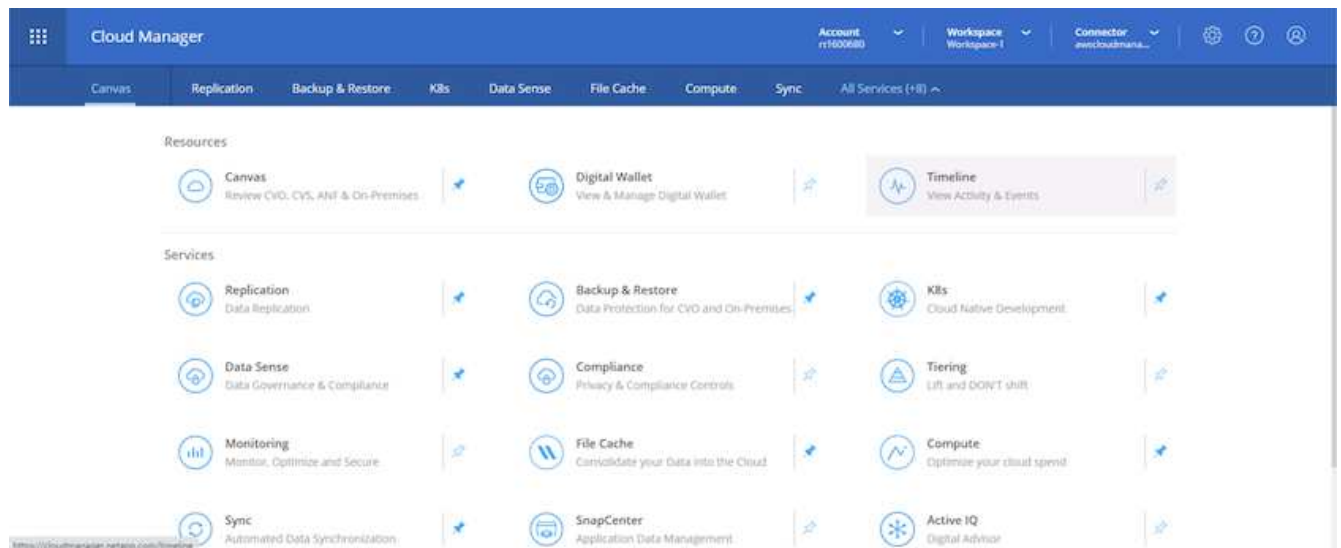
7. Review the selections made and tick the boxes to verify that you understand that Cloud Manager deploys resources into your AWS environment. When ready, click Go.

The screenshot shows the 'Review & Approve' interface in the Cloud Manager console. The top navigation bar is the same as the previous screenshot. The breadcrumb trail is 'Create a New Working Environment' > 'Review & Approve'. The interface shows the 'hybridawscvo' deployment with tabs for 'AWS', 'us-east-1', and 'HA'. There are two checkboxes with text: 'I understand that in order to activate support, I must first register Cloud Volumes ONTAP with NetApp. More information >' and 'I understand that Cloud Manager will allocate the appropriate AWS resources to comply with my above requirements. More information >'. Below these are three tabs: 'Overview', 'Networking', and 'Storage'. The 'Overview' tab is selected, showing a table of deployment details: 'Storage System: Cloud Volumes ONTAP HA', 'License Type: Cloud Volumes ONTAP Standard', 'Capacity Limit: 10TB', 'HA Deployment Model: Multiple Availability Zones', 'Encryption: AWS Managed', and 'Customer Master Key: aws/ebs'. At the bottom, there is a 'Go' button. The footer shows 'Cloud Manager 5.8.9 Build 9 Aug 18, 2021 04:13:35 am UTC'.

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.



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 Manager** Account: r1600880 Workspace: Workspace-1 Connector: awscloudmana...

Canvas | Replication | Backup & Restore | K8s | Data Sense | File Cache | Compute | Sync | All Services (+8)

Timeline

Filters: Time (1) Service Action Agent (1) Resource User Status Reset

Time	Action	Service	Agent	Resource	User	Status
Aug 18 2021, 9:42:32 pm	Check Connectivity	Cloud Manager	awscloudman...	hybridawsco	Full Name	Success
Aug 18 2021, 9:42:00 pm	Create Aws Ha Working Environment	Cloud Manager	awscloudma...	hybridawsco	Full Name	Pending
Aug 18 2021, 10:09:39 pm	Describe Operation Status					Success
Aug 18 2021, 10:00:01 pm	Describe Operation Status					Success

- 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.

**Cloud Manager** Account: r1600880 Workspace: Workspace-1 Connector: awscloudmana...

Canvas | Replication | Backup & Restore | K8s | Data Sense | File Cache | Compute | Sync | All Services (+8)

Canvas

Add Working Environment

Working environments


- 1 Cloud Volumes ONTAP (High-Availability)  
1 GB Allocated Capacity
- 1 Amazon S3  
0 Buckets

## 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](#).

- Click the source ONTAP system (on-premises) and either drag and drop it to the destination, select Replication > Enable, or select Replication > Menu > Replicate.



Select Enable.

#### SERVICES



Replication


■ Off

Enable

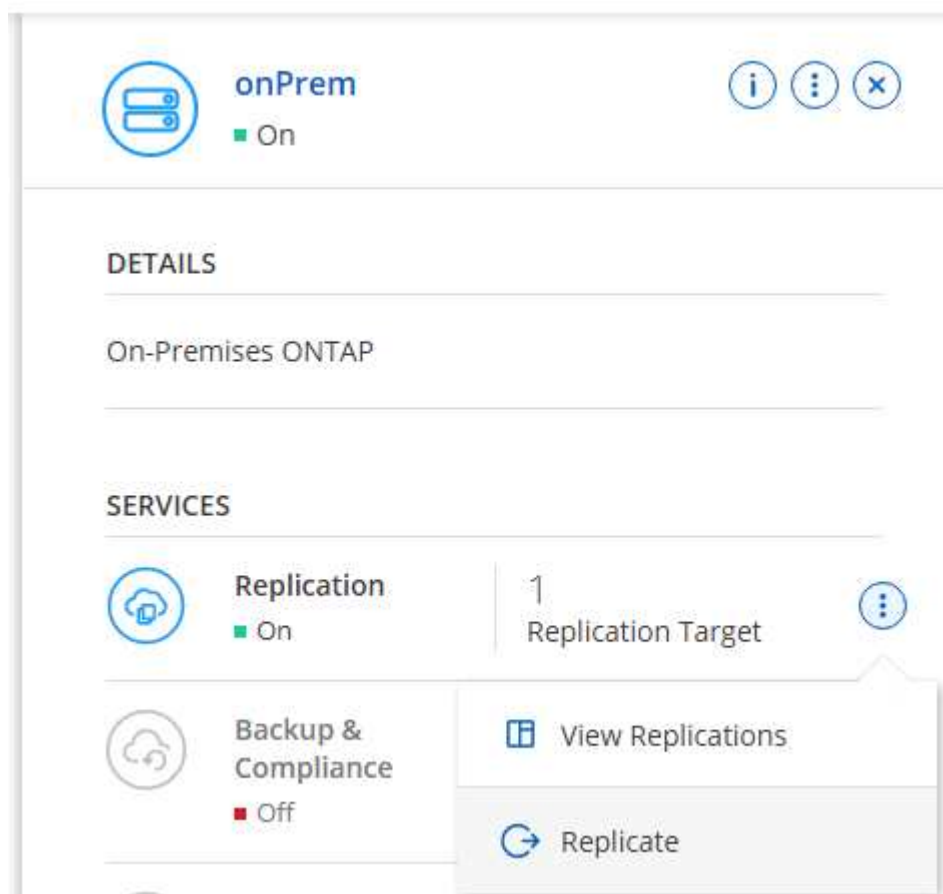


Or Options.





Replicate.



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 want 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 Setup Source Volume Selection

Volume Name	Storage VM Name	Tiering Policy	Volume Type	Capacity	Allocated	Disk Used
rhel2_u03	svm_onPrem	None	RW	100 GB	7.29 GB	35.83 MB
rhel2_u0309232119421203118	svm_onPrem	None	RW	100 GB	35.83 MB	35.83 MB
sql1_data	svm_onPrem	None	RW	53.37 GB	45.09 GB	45.09 GB
sql1_log	svm_onPrem	None	RW	21.35 GB	18.16 GB	21.23 GB
sql1_snapctr	svm_onPrem	None	RW	24.87 GB	21.23 GB	21.23 GB

Cloud Manager 3.9.10 Build: 2 Sep 12, 2021 06:47:41 am UTC

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

Destination Disk Type and Tiering

Previous Step

Destination Disk Type

General Purpose SSD

General Purpose SSD - Dynamic Performance

Throughput Optimized HDD

S3 Tiering

What are storage tiers?

☒ Enabled ☐ Disabled

Note: If you enable S3 tiering, thin provisioning must be enabled on volumes created in this aggregate.

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

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.


- ☒ Limited to:  MB/s
- ☐ Unlimited (recommended for DR only machines)

7. Define the replication policy. We chose a Mirror, which takes the most recent dataset and replicates that into the destination volume. You could also choose a different policy based on your requirements.

## Replication Policy


Default Policies

Additional Policies

 Mirror

Typically used for disaster recovery

More info

 Mirror and Backup (1 month retention)

Configures disaster recovery and long-term retention of backups on the same destination volume

More info

8. Choose the schedule for triggering replication. NetApp recommends setting a "daily" schedule of for the data volume and an "hourly" schedule for the log volumes, although this can be changed based on requirements.

Replication Setup Schedule

---

↑ Previous Step Select a replication schedule

One-time copy

No schedule

10min

Every hour  
Minutes: 0th, 10th, 20th, 3...

12-hourly

Every day  
Hours: 12 AM and 12 PM  
Minutes: 15th minute

5min

Every hour  
Minutes: 0th, 5th, 10th, 15t...

6-hourly

Every day  
Hours: 12 AM, 6 AM, 12 PM...  
Minutes: 15th minute

8hour

Every day  
Hours: 2 AM, 10 AM and 6 ...  
Minutes: 15th minute

daily

Every day  
Hours: 12 AM  
Minutes: 10th minute

hourly

Every hour  
Minutes: 5th minute

monthly

Every month  
Days: 2nd  
Hours: 12 AM  
Minutes: 20th minute

pg-15-minutely

Every hour

pg-6-hourly

Every day

pg-daily

Every day

pg-daily-set2

Every day

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 & Approve

---

↑ Previous Step Review your selection and start the replication process

Source

onPrem

sql1\_data

Destination

hybridcvo

sql1\_data\_copy

☒ I understand that Cloud Manager will allocate the appropriate AWS resources to comply with my above requirements.  
[More information >](#)

Source Volume Allocated Size:	53.37 GB	Destination Thin Provisioning:	Yes
Source Volume Used Size:	45.09 GB	Destination Aggregate:	aggr1 (Automatically s...
Source Thin Provisioning:	Yes	Destination Storage VM:	svm_hybridcvo
Destination Volume Allocated Size:	53.37 GB	Max Transfer Rate:	100 MB/s
Destination Volume Disk Type:	General Purpose SSD (...)	SnapMirror Policy:	Mirror
Capacity Tiering:	S3	Replication Schedule:	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.

Replication

7 Volume Relationships

153.32 GiB Replicated Capacity

0 Currently Transferring

7 Healthy

0 Failed

7 Volume Relationships

Health Status	Source Volume	Target Volume	Total Transfer Time	Status	Mirror State	Last Successful Transfer
	rhel2_u01 onPrem	rhel2_u01_dr hybridcvo	43 minutes 43 seconds	idle	snapmirrored	Sep 30, 2021, 12:12:50 AI 19.73 MiB
	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
	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
	sql1_data onPrem	sql1_data_dr hybridcvo	1 hour 6 minutes 50 seconds	idle	snapmirrored	Sep 30, 2021, 12:12:28 AI 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](#).

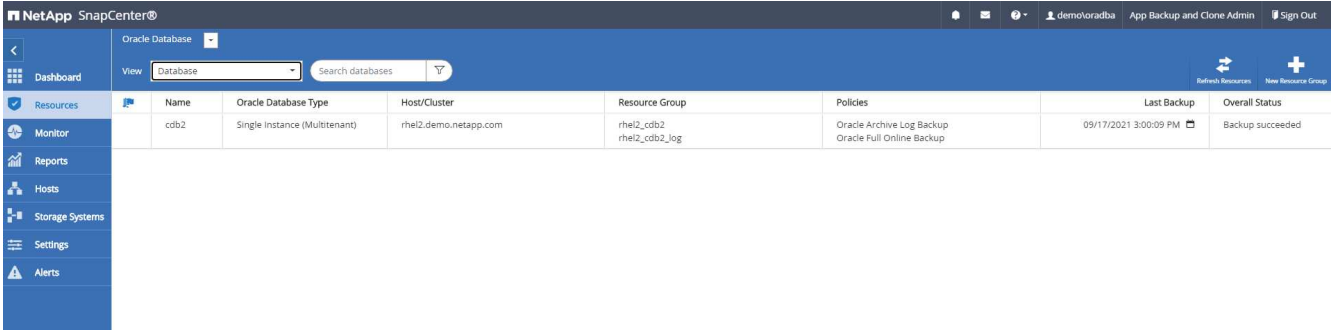
#### Workflow for dev/test bursting to cloud

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

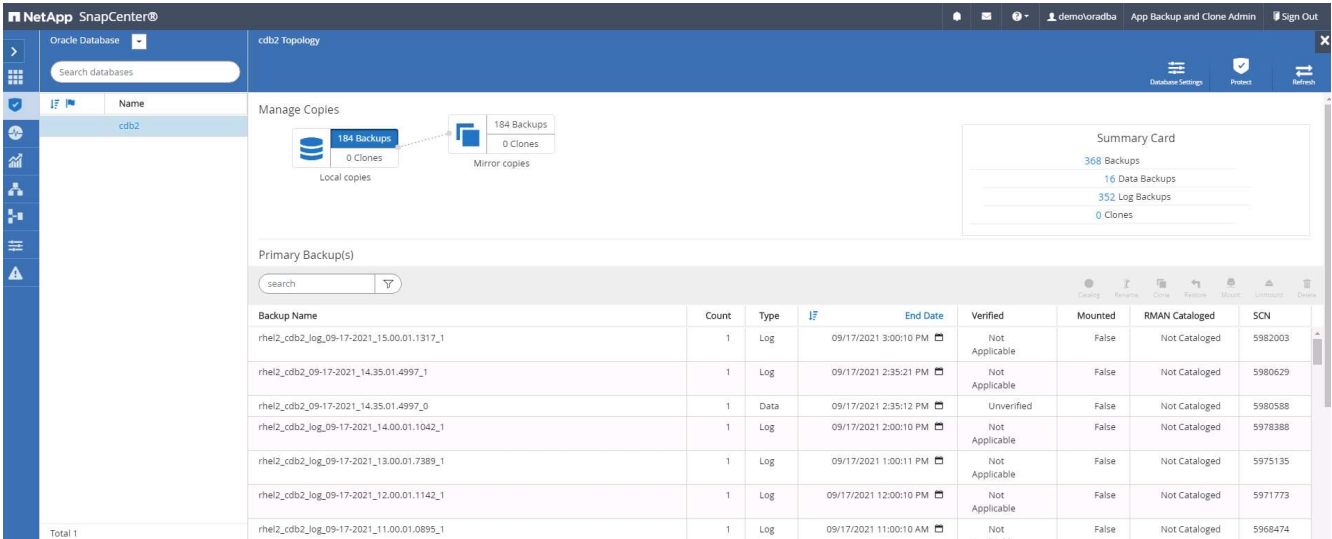
reality. SnapCenter can not only protect your production database on-premises, but can also it quickly clone a copy for application development or code testing in the public cloud while consuming very little extra storage. Following are details of the step-by-step processes for using this tool.

**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.



- 2. Click the intended on-premises database name for the backup topology and the detailed view. If a secondary replicated location is enabled, it shows linked mirror backups.



- 3. Toggled to the mirrored backups view by clicking mirrored backups. The secondary mirror backup(s) is then displayed.

NetApp SnapCenter®

Oracle Database

Search databases

cdb2

Manage Copies

184 Backups  
0 Clones  
Local copies

184 Backups  
0 Clones  
Mirror copies

Summary Card

368 Backups  
16 Data Backups  
352 Log Backups  
0 Clones

Secondary Mirror Backup(s)

search

Backup Name	Count	Type	IF	End Date	Verified	Mounted	RMAN Cataloged	SCN
rhel2_cdb2_log_09-17-2021_15.00.01.1317_1	1	Log		09/17/2021 3:00:10 PM	Not Applicable	False	Not Cataloged	5982003
rhel2_cdb2_log_09-17-2021_14.35.01.4997_1	1	Log		09/17/2021 2:35:21 PM	Not Applicable	False	Not Cataloged	5980629
rhel2_cdb2_log_09-17-2021_14.35.01.4997_0	1	Data		09/17/2021 2:35:12 PM	Unverified	False	Not Cataloged	5980588
rhel2_cdb2_log_09-17-2021_14.00.01.1042_1	1	Log		09/17/2021 2:00:10 PM	Not Applicable	False	Not Cataloged	5978388
rhel2_cdb2_log_09-17-2021_13.00.01.7389_1	1	Log		09/17/2021 1:00:11 PM	Not Applicable	False	Not Cataloged	5975135
rhel2_cdb2_log_09-17-2021_12.00.01.1142_1	1	Log		09/17/2021 12:00:10 PM	Not Applicable	False	Not Cataloged	5971773
rhel2_cdb2_log_09-17-2021_11.00.01.0895_1	1	Log		09/17/2021 11:00:10 AM	Not Applicable	False	Not Cataloged	5968474

Total 1

- Choose a mirrored secondary database backup copy to be cloned and determine a recovery point either by time and system change number or by SCN. Generally, the recovery point should be trailing the full database backup time or SCN to be cloned. After a recovery point is decided, the required log file backup must be mounted for recovery. The log file backup should be mounted to target DB server where the clone database is to be hosted.

Mount backups

Choose the host to mount the backup

ora-standby.demo.netapp.com

Mount path : /var/opt/snapcenter/sco/backup\_mount/rhel2\_cdb2\_09-17-2021\_14.35.01.4997\_1/cdb2

Secondary storage location : Snap Vault / Snap Mirror

Source Volume

svm\_onPrem:rhel2\_u03

Destination Volume

svm\_hybridcvo:rhel2\_u03\_dr

Mount Cancel



NetApp SnapCenter®

Oracle Database

Search databases

cdb2 Topology

Manage Copies

184 Backups  
0 Clones  
Local copies

184 Backups  
1 Clone  
Mirror copies

Summary Card

368 Backups  
16 Data Backups  
352 Log Backups  
1 Clone

Secondary Mirror Backup(s)

search

Backup Name	Count	Type	IF	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	1	Log		09/17/2021 3:00:10 PM	Not Applicable	False	Not Cataloged	5982003
rhel2_cdb2_log_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_log_09-17-2021_14.35.01.4997_0	1	Data		09/17/2021 2:35:12 PM	Unverified	False	Not Cataloged	5980588
rhel2_cdb2_log_09-17-2021_14.00.01.1042_1	1	Log		09/17/2021 2:00:10 PM	Not Applicable	False	Not Cataloged	5978388



If log pruning is enabled and the recovery point is extended beyond the last log pruning, multiple archive log backups might need to be mounted.

- Highlight the full database backup copy to be cloned, and then click the clone button to start the DB clone Workflow.

cdb2 Topology

search

Clone

Backup Name	Count	Type	IF	End Date	Verified	Mounted	RMAN Cataloged	SCN
rhel2_cdb2_log_09-17-2021_16.00.01.2156_1	1	Log		09/17/2021 4:00:10 PM	Not Applicable	False	Not Cataloged	5985272
rhel2_cdb2_log_09-17-2021_15.00.01.1317_1	1	Log		09/17/2021 3:00:10 PM	Not Applicable	False	Not Cataloged	5982003
rhel2_cdb2_log_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_log_09-17-2021_14.35.01.4997_0	1	Data		09/17/2021 2:35:12 PM	Unverified	False	Not Cataloged	5980588
rhel2_cdb2_log_09-17-2021_14.00.01.1042_1	1	Log		09/17/2021 2:00:10 PM	Not Applicable	False	Not Cataloged	5978388

- Choose a proper clone DB SID for a complete container database or CDB clone.

Clone from cdb2

×

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

☒ Complete Database Clone

Clone SID

cdb2test

Exclude PDBs

Type to find PDBs

☐ PDB Clone

Secondary storage location : Snap Vault / Snap Mirror

Data

Source Volume

svm\_onPrem:rhel2\_u02

Destination Volume

svm\_hybridcvo:rhel2\_u02\_dr

Logs

Source Volume

svm\_onPrem:rhel2\_u03

Destination Volume

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 cdb2

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

Select the host to create a clone

Clone host
ora-standby.demo.netapp.com

Datafile locations ⓘ

/u02\_cdb2test
Reset

Control files ⓘ

/u02\_cdb2test/cdb2test/control/control01.ctl
/u02\_cdb2test/cdb2test/control/control02.ctl
Reset

Redo logs ⓘ

Group	Size	Unit	Number of files
<div> RedoGroup 1 </div>	200	MB	1
/u02_cdb2test/cdb2test/redolog/redo03.log			
<div> RedoGroup 2 </div>	200	MB	1

Previous
Next

- The None credential name is used for OS-based authentication, which renders the database port irrelevant. Fill in the proper Oracle Home, Oracle OS User, and Oracle OS Group as configured in the target clone DB server.

Clone from cdb2

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Database Credentials for the clone

Credential name for sys user

None

+

i

Database port

1521

Oracle Home Settings

i

Oracle Home

/u01/app/oracle/product/19800/cdb2

Oracle OS User

oracle

Oracle OS Group

oinstall

Previous

Next

9. Specify the scripts to run before clone operation. More importantly, the database instance parameter can be adjusted or defined here.

Clone from cdb2

1

Name

2

Locations

3

Credentials

4

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

processes	320	×	<div>+</div> <div>Reset</div>
remote_login_passwordfile	EXCLUSIVE	×	
sga_target	4311744512	×	
undo_tablespace	UNDOTBS1	×	

Previous

Next

- Specify the recovery point either by the date and time or SCN. Until Cancel recovers the database up to the available archive logs. Specify the external archive log location from the target host where the archive log volume is mounted. If target server Oracle owner is different from the on-premises production server, verify that the archive log directory is readable by the target server Oracle owner.

Clone from cdb2

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

☒ Recover Database

☐ Until Cancel

☐ Date and Time
 

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

☒ Until SCN (System Change Number)
 

5980629

Specify external archive log locations

/var/opt/snapcenter/sco/backup\_mount/rhel2\_cdb2\_09-17-2021\_14.35.01.4997\_1/cdb2/1/orareco/CDB2/archivelog/

☒ Create new DBID
 ☒ Create tempfile for temporary tablespace
 

Enter SQL queries to apply when clone is created

Enter scripts to run after clone operation

Previous

Next

```

oracle@ora-standby/tmp
[oracle@ora-standby tmp]$ ls /var/opt/snapcenter/sco/backup_mount/rhel2_cdb2_09-17-2021_14.35.01.4997_1/cdb2/1/orareco/CDB2/archivelog/
2021_08_26 2021_08_28 2021_08_30 2021_09_01 2021_09_03 2021_09_05 2021_09_07 2021_09_09 2021_09_11 2021_09_13 2021_09_15 2021_09_17
2021_08_27 2021_08_29 2021_08_31 2021_09_02 2021_09_04 2021_09_06 2021_09_08 2021_09_10 2021_09_12 2021_09_14 2021_09_16
[oracle@ora-standby tmp]$
  
```

11. Configure the SMTP server for email notification if desired.

Clone from cdb2

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Provide email settings ⓘ

Email preference

Never

From

From email

To

Email to

Subject

Notification

☐ Attach job report

⚠

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

✕

Previous

Next

12. Clone summary.

Clone from cdb2

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Summary

Clone from backup	rhel2_cdb2_09-17-2021_14.35.01.4997_0
Clone SID	cdb2test
Clone server	ora-standby.demo.netapp.com
Exclude PDBs	none
Oracle home	/u01/app/oracle/product/19800/cdb2
Oracle OS user	oracle
Oracle OS group	oinstall
Datafile mountpaths	/u02_cdb2test
Control files	/u02_cdb2test/cdb2test/control/control01.ctl /u02_cdb2test/cdb2test/control/control02.ctl
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
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.

```

oracle@ora-standby/tmp
[oracle@ora-standby tmp]$ export ORACLE_SID=cdb2test
[oracle@ora-standby tmp]$ export ORACLE_HOME=/u01/app/oracle/product/19800/cdb2
[oracle@ora-standby tmp]$ export PATH=$PATH:$ORACLE_HOME/bin
[oracle@ora-standby tmp]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Fri Sep 17 17:49:29 2021
Version 19.3.0.0.0

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

Connected to:
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
Version 19.3.0.0.0

SQL> select name, log_mode from v$database;

NAME          LOG MODE
-----
CDB2TEST      ARCHIVELOG

SQL> select instance_name, host_name from v$instance;

INSTANCE_NAME
-----
HOST_NAME
-----
cdb2test
ora-standby.demo.netapp.com

SQL> show pdbs

  CON_ID CON_NAME              OPEN MODE RESTRICTED
  -
2 PDB$SEED                  READ ONLY NO
3 CDB2_PDB1                  READ WRITE NO
4 CDB2_PDB2                  READ WRITE NO
5 CDB2_PDB3                  READ WRITE NO

SQL>

```



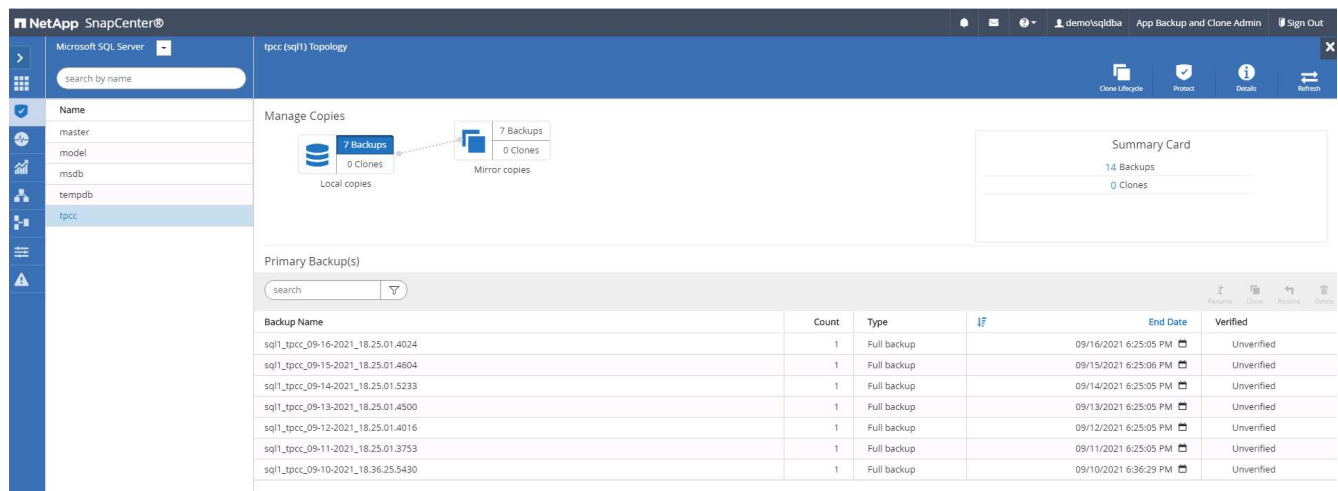
## Clone a SQL database for dev/test from a replicated Snapshot backup

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.



Name	Instance	Host	Last Backup	Overall Status	Type
master	sql1	sql1.demo.netapp.com		Not available for backup	System database
model	sql1	sql1.demo.netapp.com		Not available for backup	System database
msdb	sql1	sql1.demo.netapp.com		Not available for backup	System database
tempdb	sql1	sql1.demo.netapp.com		Not available for backup	System database
tpcc	sql1	sql1.demo.netapp.com	09/16/2021 7:35:05 PM	Backup succeeded	User database
master	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database
model	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database
msdb	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database
tempdb	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database

2. Click on the intended on-premises SQL Server user database name for the backups topology and detailed view. If a secondary replicated location is enabled, it shows linked mirror backups.



tpcc (sql1) Topology

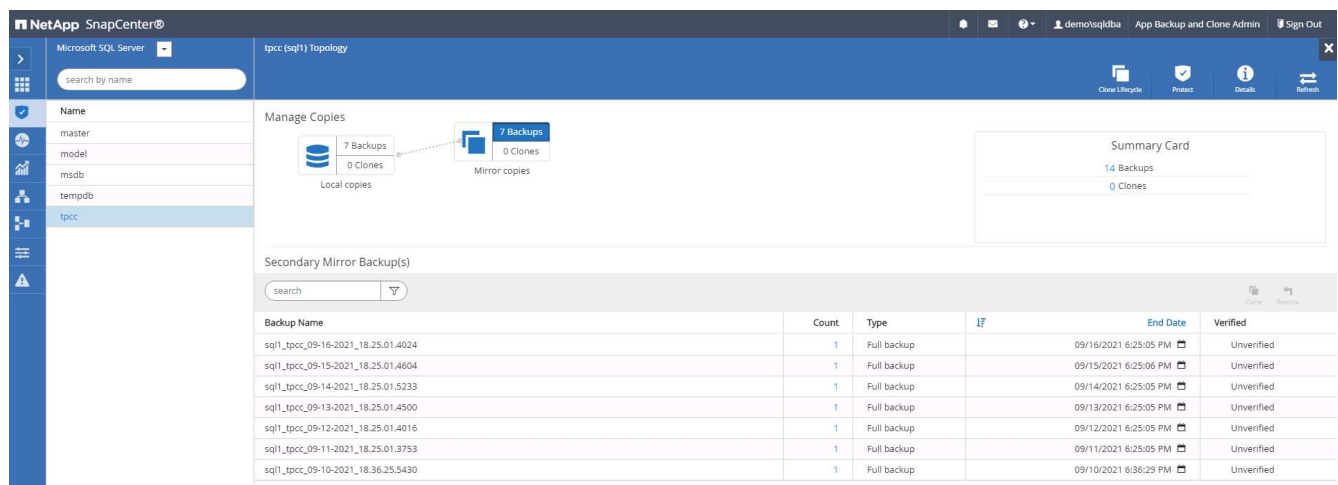
Manage Copies: 7 Backups, 0 Clones, 0 Clones, 7 Backups, 0 Clones, Mirror copies

Summary Card: 14 Backups, 0 Clones

Primary Backup(s)

Backup Name	Count	Type	IF	End Date	Verified
sql1_tpcc_09-16-2021_18.25.01.4024	1	Full backup		09/16/2021 6:25:05 PM	Unverified
sql1_tpcc_09-15-2021_18.25.01.4604	1	Full backup		09/15/2021 6:25:06 PM	Unverified
sql1_tpcc_09-14-2021_18.25.01.5233	1	Full backup		09/14/2021 6:25:05 PM	Unverified
sql1_tpcc_09-13-2021_18.25.01.4500	1	Full backup		09/13/2021 6:25:05 PM	Unverified
sql1_tpcc_09-12-2021_18.25.01.4016	1	Full backup		09/12/2021 6:25:05 PM	Unverified
sql1_tpcc_09-11-2021_18.25.01.3753	1	Full backup		09/11/2021 6:25:05 PM	Unverified
sql1_tpcc_09-10-2021_18.36.25.5430	1	Full backup		09/10/2021 6:36:29 PM	Unverified

3. Toggle to the Mirrored Backups view by clicking Mirrored Backups. Secondary Mirror Backup(s) are then displayed. Because SnapCenter backs up the SQL Server transaction log to a dedicated drive for recovery, only full database backups are displayed here.



tpcc (sql1) Topology

Manage Copies: 7 Backups, 0 Clones, 0 Clones, 7 Backups, 0 Clones, Mirror copies

Summary Card: 14 Backups, 0 Clones

Secondary Mirror Backup(s)

Backup Name	Count	Type	IF	End Date	Verified
sql1_tpcc_09-16-2021_18.25.01.4024	1	Full backup		09/16/2021 6:25:05 PM	Unverified
sql1_tpcc_09-15-2021_18.25.01.4604	1	Full backup		09/15/2021 6:25:06 PM	Unverified
sql1_tpcc_09-14-2021_18.25.01.5233	1	Full backup		09/14/2021 6:25:05 PM	Unverified
sql1_tpcc_09-13-2021_18.25.01.4500	1	Full backup		09/13/2021 6:25:05 PM	Unverified
sql1_tpcc_09-12-2021_18.25.01.4016	1	Full backup		09/12/2021 6:25:05 PM	Unverified
sql1_tpcc_09-11-2021_18.25.01.3753	1	Full backup		09/11/2021 6:25:05 PM	Unverified
sql1_tpcc_09-10-2021_18.36.25.5430	1	Full backup		09/10/2021 6:36:29 PM	Unverified

4. Choose a backup copy, and then click the Clone button to launch the Clone from Backup workflow.

The screenshot shows the NetApp SnapCenter interface for the 'tpcc (sql1) Topology'. The left sidebar lists various components: master, model, msdb, tempdb, tpcc, master, model, msdb, tempdb, and tpcc\_clone. The main area displays 'Manage Copies' with a diagram showing '7 Backups' and '1 Clone'. A 'Summary Card' on the right shows '14 Backups' and '1 Clone'. Below this, a table lists 'Secondary Mirror Backup(s)' with columns for Backup Name, Count, Type, End Date, and Verified.

Backup Name	Count	Type	End Date	Verified
sql1_tpcc_09-19-2021_18.25.01.4134	1	Full backup	09/19/2021 6:25:05 PM	Unverified
sql1_tpcc_09-18-2021_18.25.01.3963	1	Full backup	09/18/2021 6:25:05 PM	Unverified
sql1_tpcc_09-17-2021_18.25.01.4218	1	Full backup	09/17/2021 6:25:05 PM	Unverified
sql1_tpcc_09-16-2021_18.25.01.4024	1	Full backup	09/16/2021 6:25:05 PM	Unverified
sql1_tpcc_09-15-2021_18.25.01.4604	1	Full backup	09/15/2021 6:25:06 PM	Unverified
sql1_tpcc_09-14-2021_18.25.01.5233	1	Full backup	09/14/2021 6:25:05 PM	Unverified
sql1_tpcc_09-13-2021_18.25.01.4500	1	Full backup	09/13/2021 6:25:05 PM	Unverified

The screenshot shows the 'Clone from backup' workflow in NetApp SnapCenter. The left sidebar has five steps: 1 Clone Options, 2 Logs, 3 Script, 4 Notification, and 5 Summary. The main area is titled 'Clone settings' and includes fields for 'Clone server' (Choose), 'Clone instance' (Nothing selected), and 'Clone name' (tpcc). Below these, there are options for 'Choose mount option': 'Auto assign mount point' (selected) and 'Auto assign volume mount point under path' (with a 'full file path' input field). The 'Secondary storage location : Snap Vault / Snap Mirror' section shows a table with 'Source Volume' and 'Destination Volume'.

Source Volume	Destination Volume
svm_onPrem:sql1_data	svm_hybridcvo:sql1_data_dr
svm_onPrem:sql1_log	svm_hybridcvo:sql1_log_dr

At the bottom right, there are 'Previous' and 'Next' buttons.

5. Select a cloud server as the target clone server, clone instance name, and clone database name. Choose either an auto-assign mount point or a user-defined mount point path.

Clone from backup

1 Clone Options

2 Logs

3 Script

4 Notification

5 Summary

Clone settings

Clone server

sql-standby.demo.netapp.com

Clone instance

sql-standby

Clone name

tpcc\_clone

Choose mount option

☒ Auto assign mount point

☐ Auto assign volume mount point under path

full file path

Secondary storage location : Snap Vault / Snap Mirror

Source Volume	Destination Volume
svm_onPrem:sql1_data	svm_hybridcvo:sql1_data_dr
svm_onPrem:sql1_log	svm_hybridcvo:sql1_log_dr

Previous

Next

6. Determine a recovery point either by a log backup time or by a specific date and time.

Clone from backup

1 Clone Options

2 Logs

3 Script

4 Notification

5 Summary

Choose logs

☐ All log backups

☒ By log backups until

9/17/2021 6:25:10 PM

☐ By specific date until

09/17/2021 6:25:05 PM

☐ None

Previous

Next

7. Specify optional scripts to run before and after the cloning operation.

Clone from backup

1 Clone Options

2 Logs

3 Script

4 Notification

5 Summary

Specify optional scripts to run before and after performing a clone from backup job

Prescript full path

Prescript arguments

Choose optional arguments...

Postscript full path

Postscript arguments

Choose optional arguments...

Script timeout

60

secs

Previous

Next

8. Configure an SMTP server if email notification is desired.

Clone from backup

1 Clone Options

2 Logs

3 Script

4 Notification

5 Summary

Provide email settings ⓘ

Email preference

Never

From

From email

To

Email to

Subject

Notification

☐ Attach Job Report

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

✕

Previous

Next

9. Clone Summary.

Clone from backup

1 Clone Options

2 Logs

3 Script

4 Notification

5 Summary

Summary

Clone server

sql-standby.demo.netapp.com

Clone instance

sql-standby

Clone name

tpcc\_dev

Mount option

Auto assign volume mount point under custom path

Prescript full path

None

Prescript arguments

Postscript full path

None

Postscript arguments

Send email

No

Previous

Finish

- 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 SnapCenter®						
Jobs - Filter						
	ID	Status	Name	Start date	End date	Owner
	766	✓	Clone from backup 'sql1_tpcc_09-16-2021_18.25.01.4024'	09/16/2021 8:05:25 PM	09/16/2021 8:06:17 PM	demo:sqldba
	763	✓	Discover resources for all hosts	09/16/2021 7:56:49 PM	09/16/2021 7:56:54 PM	demo:sqldba
	761	✓	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/16/2021 7:35:00 PM	09/16/2021 7:37:08 PM	demo:sqldba
	760	⚠	Discover resources for all hosts	09/16/2021 7:19:05 PM	09/16/2021 7:19:09 PM	demo:sqldba
	759	⚠	Discover resources for all hosts	09/16/2021 7:18:43 PM	09/16/2021 7:18:48 PM	demo:sqldba
	756	⚠	Discover resources for all hosts	09/16/2021 6:59:51 PM	09/16/2021 6:59:56 PM	demo:sqldba
	753	✓	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/16/2021 6:35:00 PM	09/16/2021 6:37:07 PM	demo:sqldba
	750	✓	Backup of Resource Group 'sql1_tpcc' with policy 'SQL Server Full Backup'	09/16/2021 6:25:01 PM	09/16/2021 6:27:14 PM	demo:sqldba
	749	✓	Discover resources for host 'sql-standby.demo.netapp.com'	09/16/2021 6:19:00 PM	09/16/2021 6:19:05 PM	DemoAdministrator
	745	✓	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/16/2021 5:35:00 PM	09/16/2021 5:37:08 PM	demo:sqldba

## Post-clone configuration

- 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.
- Configure an Oracle listener, or register the newly cloned DB with an existing listener for user access.
- For SQL Server, change the log mode from Full to Easy so that the SQL Server dev/test log file can be readily shrunk when it is filling up the log volume.

### Refresh clone database

1. Drop cloned databases and clean up the cloud DB server environment. Then follow the previous procedures to clone a new DB with fresh data. It only takes few minutes to clone a new database.
2. Shutdown the clone database, run a clone refresh command by using the CLI. See the following SnapCenter documentation for details: [Refresh a clone](#).

### Where to go for help?

If you need help with this solution and use cases, join the [NetApp Solution Automation community support Slack channel](#) and look for the solution-automation channel to post your questions or inquires.

### Disaster recovery workflow

Enterprises have embraced the public cloud as a viable resource and destination for disaster recovery. SnapCenter makes this process as seamless as possible. This disaster recovery workflow is very similar to the clone workflow, but database recovery runs through the last available log that was replicated to cloud to recover all the business transactions possible. However, there are additional pre-configuration and post-configuration steps specific to disaster recovery.

### Clone an on-premises Oracle production DB to cloud for DR

1. To validate that the clone recovery runs through last available log, we created a small test table and inserted a row. The test data would be recovered after a full recovery to last available log.

```
oracle@rhel2~$
SQL> create table dr_test(
  2 id integer,
  3 event varchar(200),
  4 dt timestamp);

Table created.

SQL> insert into dr_test values(1, 'testing DB clone for DR and roll forward DB to last available log', sysdate);

1 row created.

SQL> select * from dr_test;

      ID
-----
EVENT
-----
DT
-----
1
testing DB clone for DR and roll forward DB to last available log
17-SEP-21 02.12.13.000000 PM

SQL> commit;

Commit complete.



SQL>
```

2. Log into SnapCenter as a database management user ID for Oracle. Navigate to the Resources tab, which shows the Oracle databases being protected by SnapCenter.

NetApp SnapCenter®					
Oracle Database					
View Resource Group Search resource group					
	Name	Resources	Tags	Policies	Last Backup Overall Status
	rhel2_cdb2	1	orafullbup	Oracle Full Online Backup	09/17/2021 2:38:16 PM Completed
	rhel2_cdb2_log	1	oralogbup	Oracle Archive Log Backup	09/17/2021 6:02:13 PM Completed



3. Select the Oracle log resource group and click Backup Now to manually run an Oracle log backup to flush the latest transaction to the destination in the cloud. In a real DR scenario, the last transaction recoverable depends on the database log volume replication frequency to the cloud, which in turn depends on the RTO or RPO policy of the company.



Asynchronous SnapMirror loses data that has not made it to the cloud destination in the database log backup interval in a disaster recovery scenario. To minimize data loss, more frequent log backup can be scheduled. However there is a limit to the log backup frequency that is technically achievable.

4. Select the last log backup on the Secondary Mirror Backup(s), and mount the log backup.

NetApp SnapCenter®

Oracle Database

Search databases

cdb2 Topology

Manage Copies

Local copies: 185 Backups, 0 Clones

Mirror copies: 185 Backups, 2 Clones

Summary Card

- 370 Backups
- 16 Data Backups
- 354 Log Backups
- 2 Clones

Secondary Mirror Backup(s)

Backup Name	Count	Type	LF	End Date	Verified	Mounted	RMAN Cataloged	SCN
rhel2_cdb2_log_09-17-2021_18.20.04.1177_1	1	Log		09/17/2021 6:20:13 PM	Not Applicable	False	Not Cataloged	5994710
rhel2_cdb2_log_09-17-2021_18.00.01.2424_1	1	Log		09/17/2021 6:00:09 PM	Not Applicable	False	Not Cataloged	5992079
rhel2_cdb2_log_09-17-2021_17.00.01.1566_1	1	Log		09/17/2021 5:00:20 PM	Not Applicable	False	Not Cataloged	5988842

### Mount backups

Choose the host to mount the backup:

Mount path : /var/opt/snapcenter/sco/backup\_mount/rhel2\_cdb2\_log\_09-17-2021\_18.20.04.1177\_1/cdb2

Secondary storage location : Snap Vault / Snap Mirror

Source Volume: svm\_onPrem:rhel2\_u03

Destination Volume:

5. Select the last full database backup and click Clone to initiate the clone workflow.

NetApp SnapCenter®

Oracle Database

Search databases

cdb2 Topology

Manage Copies

Local copies: 185 Backups, 0 Clones

Mirror copies: 185 Backups, 2 Clones

Summary Card

- 370 Backups
- 16 Data Backups
- 354 Log Backups
- 2 Clones

Secondary Mirror Backup(s)

Backup Name	Count	Type	LF	End Date	Verified	Mounted	RMAN Cataloged	SCN
rhel2_cdb2_log_09-17-2021_18.20.04.1177_1	1	Log		09/17/2021 6:20:13 PM	Not Applicable	True	Not Cataloged	5994710
rhel2_cdb2_log_09-17-2021_18.00.01.2424_1	1	Log		09/17/2021 6:00:09 PM	Not Applicable	False	Not Cataloged	5992079
rhel2_cdb2_log_09-17-2021_17.00.01.1566_1	1	Log		09/17/2021 5:00:20 PM	Not Applicable	False	Not Cataloged	5988842
rhel2_cdb2_log_09-17-2021_16.00.01.2156_1	1	Log		09/17/2021 4:00:10 PM	Not Applicable	False	Not Cataloged	5985272
rhel2_cdb2_log_09-17-2021_15.00.01.1317_1	1	Log		09/17/2021 3:00:10 PM	Not Applicable	False	Not Cataloged	5982003
rhel2_cdb2_log_09-17-2021_14.35.01.4997_1	1	Log		09/17/2021 2:35:21 PM	Not Applicable	False	Not Cataloged	5980629
rhel2_cdb2_log_09-17-2021_14.35.01.4997_0	1	Data		09/17/2021 2:35:12 PM	Unverified	False	Not Cataloged	5980588

Total 3

6. Select a unique clone DB ID on the host.

Clone from cdb2

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

☒ Complete Database Clone

Clone SID:

Exclude PDBs:

☐ PDB Clone

Secondary storage location : Snap Vault / Snap Mirror

☒ Data

Source Volume: svm\_onPrem:rhel2\_u02

Destination Volume:

☒ Logs

Source Volume: svm\_onPrem:rhel2\_u03

Destination Volume:

Previous Next

7. Provision a log volume and mount it to the target DR server for the Oracle flash recovery area and online logs.

ONTAP System Manager

Search actions, objects, and pages

DASHBOARD

STORAGE

Overview

Applications

Volumes

LUNs

Shares

Qtrees

Quotas

Storage VMs

Tiers

NETWORK

EVENTS & JOBS

PROTECTION

HOSTS

Volumes

+ Add More

	Name	Storage VM	Status	Capacity
<input type="checkbox"/>	ora_standby_u01	svm_hybridcvo	Online	12.3 GB used 17.7 GB available 31.6 GB
<input checked="" type="checkbox"/>	rhel2_u01_dr	svm_hybridcvo	Online	
<input checked="" type="checkbox"/>	rhel2_u02_dr	svm_hybridcvo	Online	
<input checked="" type="checkbox"/>	rhel2_u02_dr09172116081193_60	svm_hybridcvo	Online	
<input checked="" type="checkbox"/>	rhel2_u02_dr09172117035348_63	svm_hybridcvo	Online	
<input checked="" type="checkbox"/>	rhel2_u03_dr	svm_hybridcvo	Online	
<input checked="" type="checkbox"/>	rhel2_u03_dr09172118245747_75	svm_hybridcvo	Online	

Add Volume

NAME:

CAPACITY:

More Options Cancel Save

```
ec2-user@ora-standby:tmp
[ec2-user@ora-standby tmp]$ sudo mkdir /u03_cdb2dr
[ec2-user@ora-standby tmp]$ chown oracle:oinstall /u03_cdb2dr
chown: changing ownership of '/u03_cdb2dr': Operation not permitted
[ec2-user@ora-standby tmp]$ sudo chown oracle:oinstall /u03_cdb2dr
[ec2-user@ora-standby tmp]$ sudo mount -t nfs 10.221.1.6:/ora_standby_u03 /u03_cdb2dr
[ec2-user@ora-standby tmp]$ df -h
Filesystem                                Size  Used Avail Use% Mounted on
devtmpfs                                  7.6G   0  7.6G   0% /dev
tmpfs                                     7.6G   0  7.6G   0% /dev/shm
tmpfs                                     7.6G  17M  7.6G   1% /run
tmpfs                                     7.6G   0  7.6G   0% /sys/fs/cgroup
/dev/nvme0n1p2                           10G   9.0G  1.1G  90% /
10.221.1.6:/ora_standby_u01              21G   13G   8G   62% /u01
tmpfs                                     1.6G   0  1.6G   0% /run/user/1000
10.221.1.6:/Sc28182452-3fa8-448c-9e4a-c5a9e465f353 1.6G   0  1.6G   0% /run/user/54321
10.221.1.6:/Sc39c05df8-4b00-4b3a-853c-9d6d338e5df7 100G   3.7G  97G   4% /u02_cdb2test
10.221.1.6:/Sccf886a5c-3273-475e-ad97-472b2a8dccee 100G   3.8G  97G   4% /var/opt/snapcenter/sco/backup_mount/rhel2_cdb2_log_09-17-2021_18.20.04.1177_1/cdb2/1
10.221.1.6:/ora_standby_u03             21G  320K  20G   1% /u03_cdb2dr
[ec2-user@ora-standby tmp]$
```



The Oracle clone procedure does not create a log volume, which needs to be provisioned on the DR server before cloning.

8. Select the target clone host and location to place the data files, control files, and redo logs.

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Select the host to create a clone

Clone host 

ora-standby.demo.netapp.com

Datafile locations

/u02\_cdb2dr

Reset

Control files

/u02\_cdb2dr/cdb2dr/control/control01.ctl

X

/u03\_cdb2dr/cdb2dr/control/control02.ctl

X

Reset

Redo logs

Group	Size	Unit	Number of files
RedoGroup 1	200	MB	1
<div><div>/u03_cdb2dr/cdb2dr/redolog/redo03.log</div><div>X</div></div>			
RedoGroup 2	200	MB	1

Reset

Previous

Next

9. Select the credentials for the clone. Fill in the details of the Oracle home configuration on the target server.

Clone from cdb2

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Database Credentials for the clone

Credential name for sys user

None

+

i

Database port

1521

Oracle Home Settings

i

Oracle Home

/u01/app/oracle/product/19800/cdb2

Oracle OS User

oracle

Oracle OS Group

oinstall

Previous

Next

10. Specify the scripts to run before cloning. Database parameters can be adjusted if needed.

Clone from cdb2

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Specify scripts to run before clone operation

Prescript full path

/var/opt/snapcenter/spl/scripts/

Enter Prescript path

Arguments

Script timeout

60

secs

Database Parameter settings

audit_file_dest	/u01/app/oracle/admin/cdb2dr/adump	X
audit_trail	DB	X
open_cursors	300	X
pga_aggregate_target	1432354816	X

+

Reset

Previous

Next

11. Select Until Cancel as the recovery option so that the recovery runs through all available archive logs to recoup the last transaction replicated to the secondary cloud location.



Clone from cdb2

1 Name

2 Locations

3 Credentials

4 PreOps

5 PostOps

6 Notification

7 Summary

Provide email settings ⓘ

Email preference

Never

From

From email

To

Email to

Subject

Notification

☐ Attach job report

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

✕

Previous

Next

13. DR clone summary.



Clone from cdb2

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

### Summary

Clone from backup	rhel2_cdb2_09-17-2021_14.35.01.4997_0
Clone SID	cdb2dr
Clone server	ora-standby.demo.netapp.com
Exclude PDBs	none
Oracle home	/u01/app/oracle/product/19800/cdb2
Oracle OS user	oracle
Oracle OS group	oinstall
Datafile mountpaths	/u02_cdb2dr
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
Finish

- Cloned DBs are registered with SnapCenter immediately after clone completion and are then available for backup protection.

NetApp SnapCenter®							
<div> <div> Dashboard Resources Monitor Reports Hosts Storage Systems Settings Alerts </div> <div> Oracle Database View Database Search databases </div> <div> demo/oradb App Backup and Clone Admin Sign Out </div> </div>							
	Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Overall Status
	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
	cdb2dev	Single Instance (Multitenant)	ora-standby.demo.netapp.com				Not protected
	cdb2dr	Single Instance (Multitenant)	ora-standby.demo.netapp.com				Not protected
	cdb2test	Single Instance (Multitenant)	ora-standby.demo.netapp.com				Not protected

## Post DR clone validation and configuration for Oracle

- Validate the last test transaction that has been flushed, replicated, and recovered at the DR location in the cloud.

```

oracle@ora-standby:/u01/app/oracle/product/19800/cdb2/dbs
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
Version 19.3.0.0.0

SQL> set lin 200
SQL> select instance_name, host_name from v$instance;

INSTANCE_NAME      HOST_NAME
-----
cdb2dr             ora-standby.demo.netapp.com

SQL> alter pluggable database cdb2_pdb1 open;

Pluggable database altered.

SQL> alter session set container=cdb2_pdb1;

Session altered.

SQL> select * from pdbadmin.dr_test;

      ID
-----
EVENT
-----
DT
-----
1
testing DB clone for DR and roll forward DB to last available log
17-SEP-21 02.12.13.000000 PM

SQL>

```

## 2. Configure the flash recovery area.

```

oracle@ora-standby:/u01/app/oracle/product/19800/cdb2/dbs
[oracle@ora-standby:dba]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Fri Sep 17 22:07:11 2021
Version 19.3.0.0.0

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

Connected to:
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
Version 19.3.0.0.0

SQL> show parameter db_recovery_file_dest

NAME                                 TYPE      VALUE
-----
db_recovery_file_dest                string    /u03_cdb2dr/cdb2dr
db_recovery_file_dest_size           big integer 17208M
SQL> alter system set db_recovery_file_dest='/u03_cdb2dr/cdb2dr' scope=both;

System altered.

SQL> show parameter db_recovery_file_dest

NAME                                 TYPE      VALUE
-----
db_recovery_file_dest                string    /u03_cdb2dr/cdb2dr
db_recovery_file_dest_size           big integer 17208M
SQL>

```

3. Configure the Oracle listener for user access.
4. Split the cloned volume off of the replicated source volume.
5. Reverse replication from the cloud to on-premises and rebuild the failed on-premises database server.




Clone split may incur temporary storage space utilization that is much higher than normal operation. However, after the on-premises DB server is rebuilt, extra space can be released.

## Clone an on-premises SQL production DB to cloud for DR


1. Similarly, to validate that the SQL clone recovery ran through last available log, we created a small test table and inserted a row. The test data would be recovered after a full recovery to the last available log.

```
Administrator Command Prompt - sqlcmd - SQLCMD
C:\Users\administrator.DEMO>sqlcmd
1> select host_name()
2> go
-----
SQL1
(1 rows affected)
1> use tpcc
2> go
Changed database context to 'tpcc'.
1> insert into snap_sync values ('test snap mirror DR for SQL', getdate())
2> go
(1 rows affected)
1> select * from snap_sync
2> go
event                                         dt
-----
test snap mirror DR for SQL                 2021-09-20 14:23:04.533
(1 rows affected)
1>
```

2. Log into SnapCenter with a database management user ID for SQL Server. Navigate to the Resources tab, which shows the SQL Server protection resources group.



3. Manually run a log backup to flush the last transaction to be replicated to secondary storage in the public cloud.



4. Select the last full SQL Server backup for the clone.

NetApp SnapCenter®

Microsoft SQL Server

tpcc (sql1) Topology

Manage Copies

Local copies: 7 Backups, 0 Clones

Mirror copies: 7 Backups, 2 Clones

Summary Card

14 Backups

2 Clones

Secondary Mirror Backup(s)

Backup Name	Count	Type	if	End Date	Verified
sql1_tpcc_09-19-2021_18.25.01.4134	1	Full backup		09/19/2021 6:25:05 PM	Unverified
sql1_tpcc_09-18-2021_18.25.01.3963	1	Full backup		09/18/2021 6:25:05 PM	Unverified
sql1_tpcc_09-17-2021_18.25.01.4218	1	Full backup		09/17/2021 6:25:05 PM	Unverified

- 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 backup

1 Clone Options

2 Logs

3 Script

4 Notification

5 Summary

Clone settings

Clone server: sql-standby.demo.netapp.com

Clone instance: sql-standby

Clone name: tpcc\_dr

Choose mount option

☒ Auto assign mount point

☐ Auto assign volume mount point under path: full file path

Secondary storage location : Snap Vault / Snap Mirror

Source Volume	Destination Volume
svm_onPrem:sql1_data	svm_hybridcvo:sql1_data_dr
svm_onPrem:sql1_log	svm_hybridcvo:sql1_log_dr

Previous Next

- Select all log backups to be applied.

Clone from backup

1 Clone Options

2 Logs

3 Script

4 Notification

5 Summary

Choose logs

☒ All log backups

☐ By log backups until

9/19/2021 6:25:10 PM

☐ By specific date until

09/19/2021 6:25:05 PM

☐ None

Previous

Next

7. Specify any optional scripts to run before or after cloning.

Clone from backup

1 Clone Options

2 Logs

3 Script

4 Notification

5 Summary

Specify optional scripts to run before and after performing a clone from backup job

Prescript full path

Prescript arguments

Choose optional arguments...

Postscript full path

Postscript arguments

Choose optional arguments...

Script timeout

60

secs

Previous

Next

8. Specify an SMTP server if email notification is desired.

Clone from backup

1 Clone Options

2 Logs

3 Script

4 Notification

5 Summary

Provide email settings ⓘ

Email preference

Never

From

From email

To

Email to

Subject

Notification

☐ Attach Job Report

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

✕

Previous

Next

9. DR clone summary. Cloned databases are immediately registered with SnapCenter and available for backup protection.

Clone from backup

1 Clone Options

2 Logs

3 Script

4 Notification

5 Summary

Summary

Clone server

sql-standby.demo.netapp.com

Clone instance

sql-standby

Clone name

tpcc\_dr

Mount option

Auto Mount

Prescript full path

None

Prescript arguments

Postscript full path

None

Postscript arguments

Send email

No

Previous

Finish

NetApp SnapCenter®							
Microsoft SQL Server							
View Database search by name							
Resources							
	Name	Instance	Host	Last Backup	Overall Status	Type	
	master	sql1	sql1.demo.netapp.com		Not available for backup	System database	
	model	sql1	sql1.demo.netapp.com		Not available for backup	System database	
	msdb	sql1	sql1.demo.netapp.com		Not available for backup	System database	
	tempdb	sql1	sql1.demo.netapp.com		Not available for backup	System database	
	tpcc	sql1	sql1.demo.netapp.com	09/22/2021 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	
	model	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database	
	msdb	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database	
	tempdb	sql-standby	sql-standby.demo.netapp.com		Not available for backup	System database	
	tpcc_clone	sql-standby	sql-standby.demo.netapp.com		Not protected	User database	
	tpcc_dev	sql-standby	sql-standby.demo.netapp.com		Not protected	User database	
	tpcc_dr	sql-standby	sql-standby.demo.netapp.com		Not protected	User database	

## Post DR clone validation and configuration for SQL

1. Monitor clone job status.

NetApp SnapCenter®						
Jobs Schedules Events Logs						
search by name						
Jobs - Filter						
	ID	Status	Name	Start date	End date	Owner
	1052	✓	Clone from backup 'sql1_tpcc_09-19-2021_18.25.01.4134'	09/20/2021 2:36:17 PM	09/20/2021 2:37:06 PM	demo/sqlqdba
	1047	✓	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/20/2021 2:35:01 PM	09/20/2021 2:37:08 PM	demo/sqlqdba
	1045	✓	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/20/2021 2:28:17 PM	09/20/2021 2:30:25 PM	demo/sqlqdba
	1044	✓	Clone from backup 'sql1_tpcc_09-17-2021_18.25.01.4218'	09/20/2021 1:39:24 PM	09/20/2021 1:40:09 PM	demo/sqlqdba
	1042	✓	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/20/2021 1:35:01 PM	09/20/2021 1:37:08 PM	demo/sqlqdba
	1040	✓	Backup of Resource Group 'sql1_tpcc_log' with policy 'SQL Server Log Backup'	09/20/2021 12:35:01 PM	09/20/2021 12:37:08 PM	demo/sqlqdba

2. Validate that last transaction has been replicated and recovered with all log file clones and recovery.



```
Administrator: Command Prompt - sqlcmd - SQLCMD
C:\Users\administrator.DEMO>sqlcmd
1> select host_name()
2> go

-----
SQL-STANDBY
(1 rows affected)
1> use tpcc_dr
2> go
Changed database context to 'tpcc_dr'.
1> select * from snap_sync
2> go
event                                     dt
-----
test snap mirror DR for SQL              2021-09-20 14:23:04.533
(1 rows affected)
1> select getdate()
2> go

-----
2021-09-20 14:39:19.937
(1 rows affected)
1> _
```

3. Configure a new SnapCenter log directory on the DR server for SQL Server log backup.
4. Split the cloned volume off of the replicated source volume.
5. Reverse replication from the cloud to on-premises and rebuild the failed on-premises database server.

### Where to go for help?

If you need help with this solution and use cases, please join the [NetApp Solution Automation community support Slack channel](#) and look for the solution-automation channel to post your questions or inquiries.

## 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.