



Getting started overview

NetApp Solutions

NetApp
May 17, 2024

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

Table of Contents

- Getting started overview 1
- On-premises 1
- AWS public cloud 1
- Getting started on premises 1
- Getting Started with AWS public cloud 54

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 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.

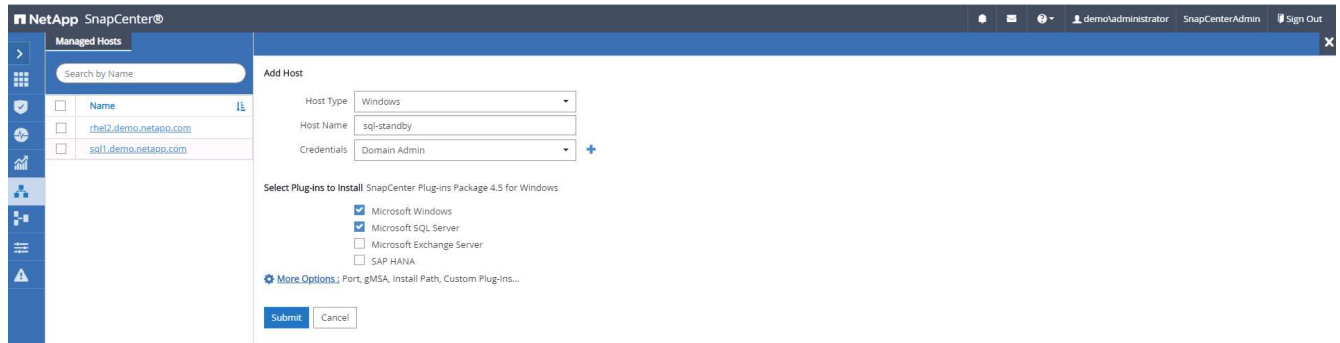
Configure SnapCenter VMware global settings

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

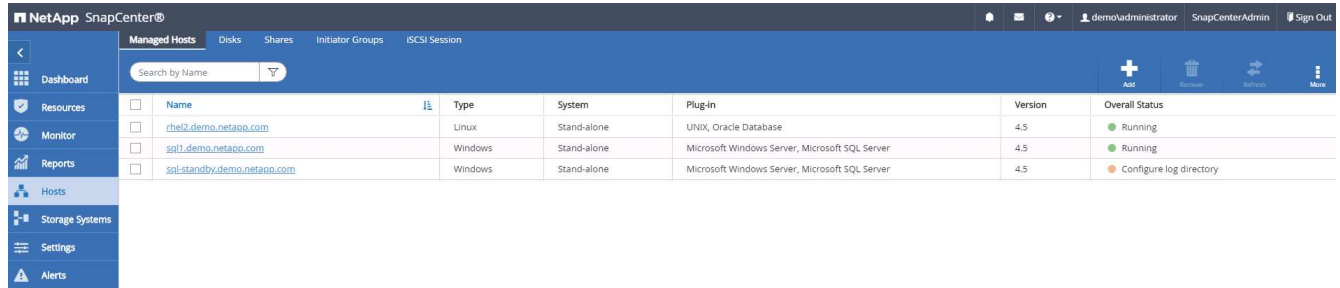


Add Windows host and installation of plugin on the host

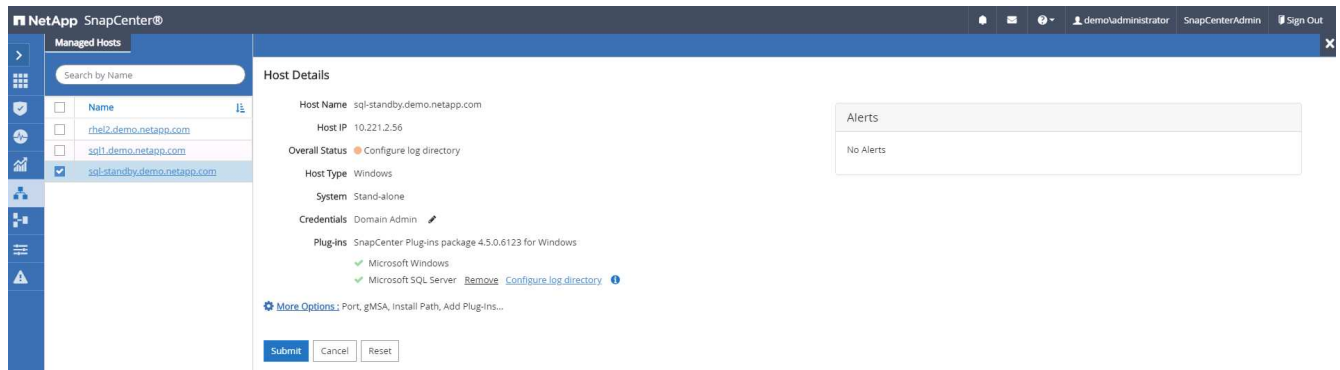
1. Log into SnapCenter with a user ID with SnapCenterAdmin privileges.
2. Click the Hosts tab from the left-hand menu, and then click Add to open the Add Host workflow.
3. Choose Windows for Host Type; the Host Name can be either a host name or an IP address. The host name must be resolved to the correct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



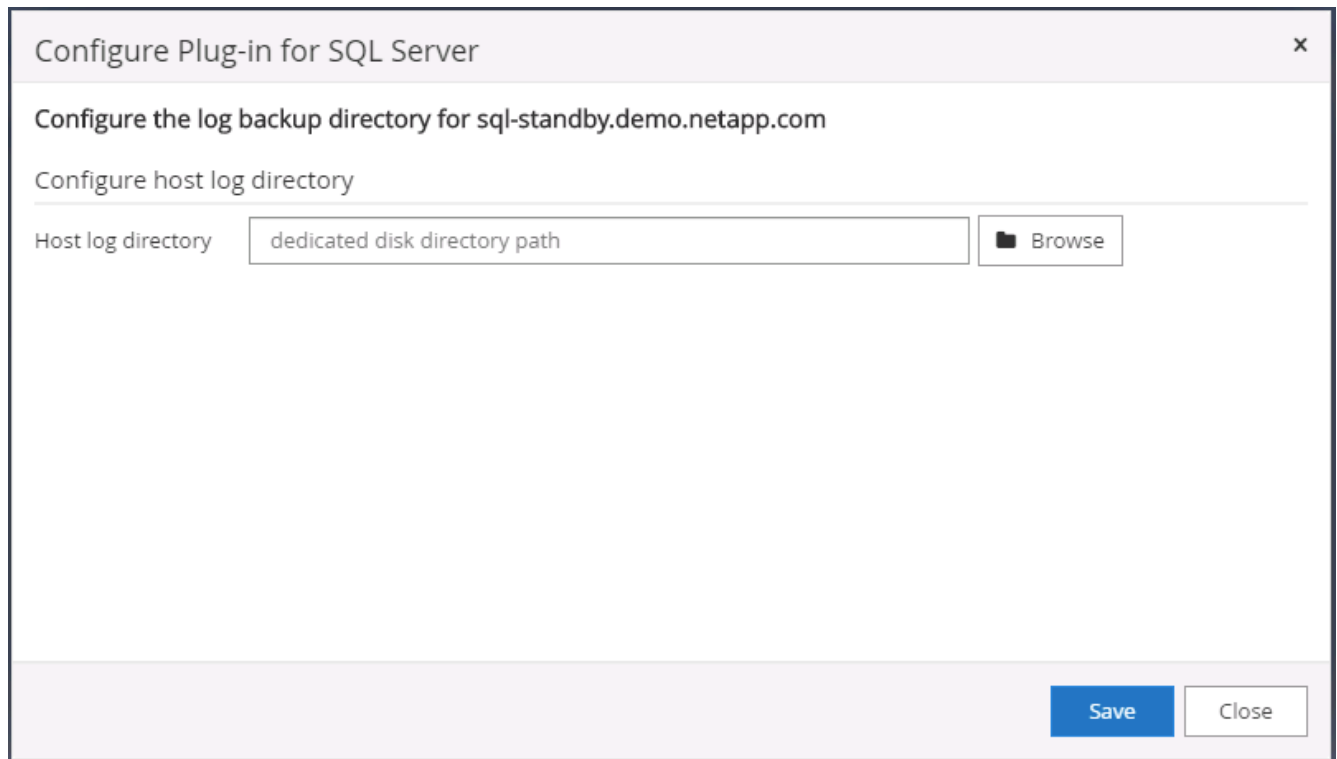
4. After the plugin is installed on a Windows host, its Overall Status is shown as "Configure log directory."



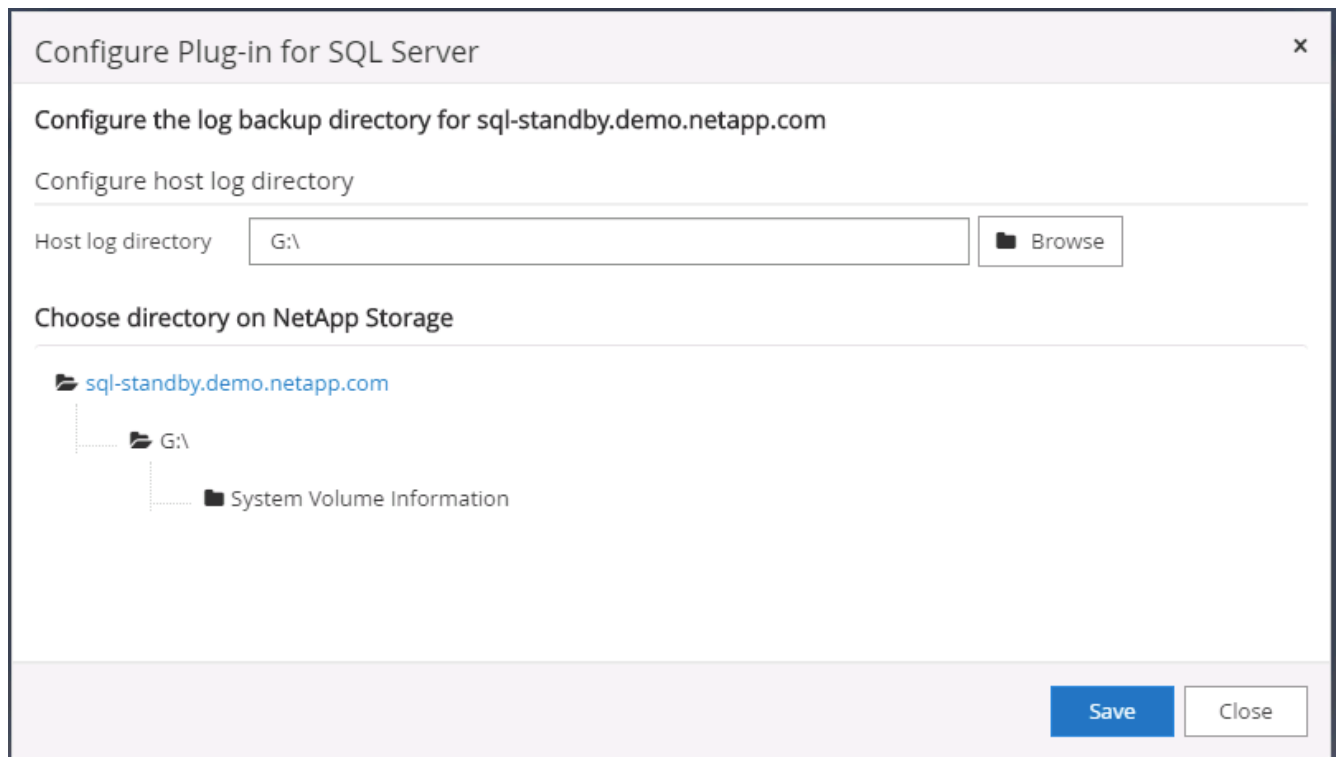
5. Click the Host Name to open the SQL Server log directory configuration.



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

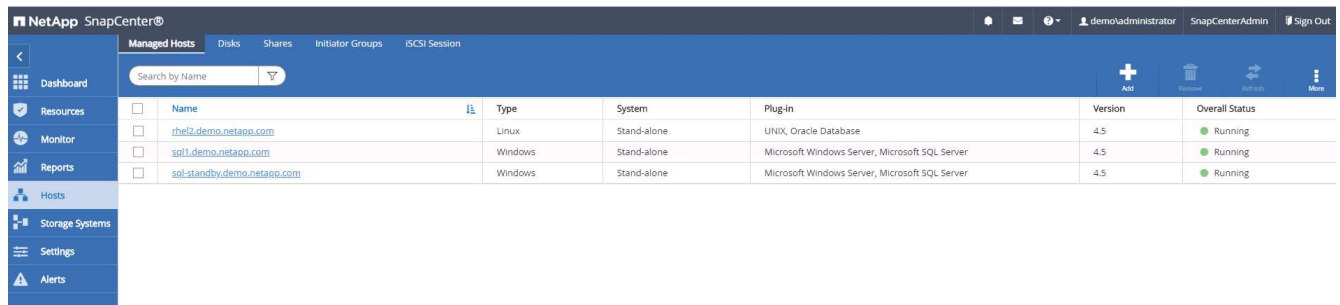


7. Click Browse to discover NetApp storage so that a log directory can be set; SnapCenter uses this log directory to roll up the SQL server transaction log files. Then click Save.

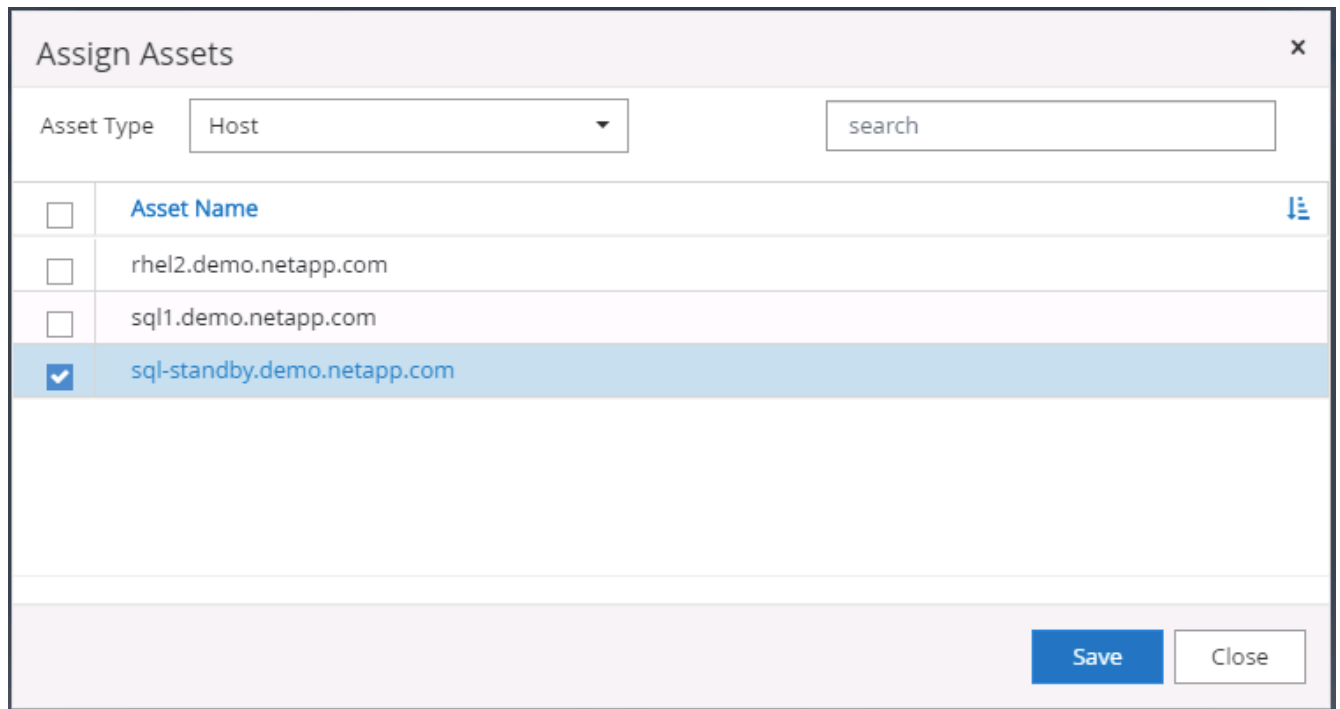
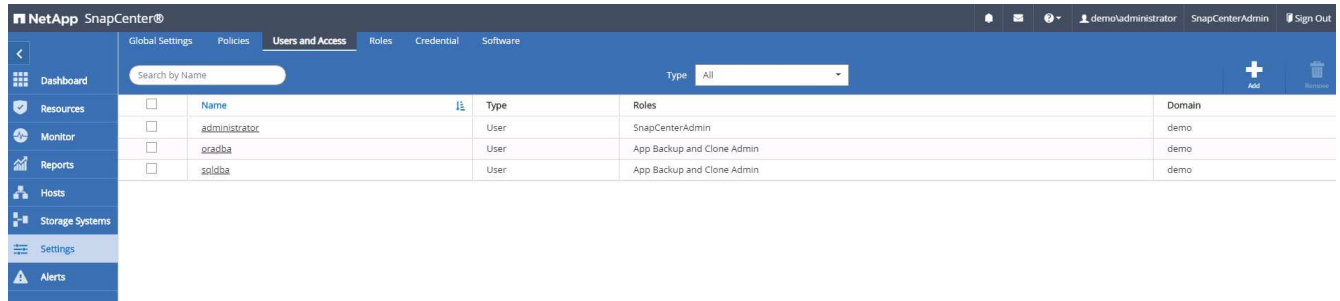


For NetApp storage provisioned to a DB host to be discovered, the storage (on-prem or CVO) must be added to SnapCenter, as illustrated in step 6 for CVO as an example.

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



- To assign the host to the database management user ID, navigate to the Access tab under Settings and Users, click the database management user ID (in our case the sqldba that the host needs to be assigned to), and click Save to complete host resource assignment.



Add Unix host and installation of plugin on the host

- Log into SnapCenter with a user ID with SnapCenterAdmin privileges.
- Click the Hosts tab from left-hand menu, and click Add to open the Add Host workflow.
- 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 plug-ins.

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

- 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

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

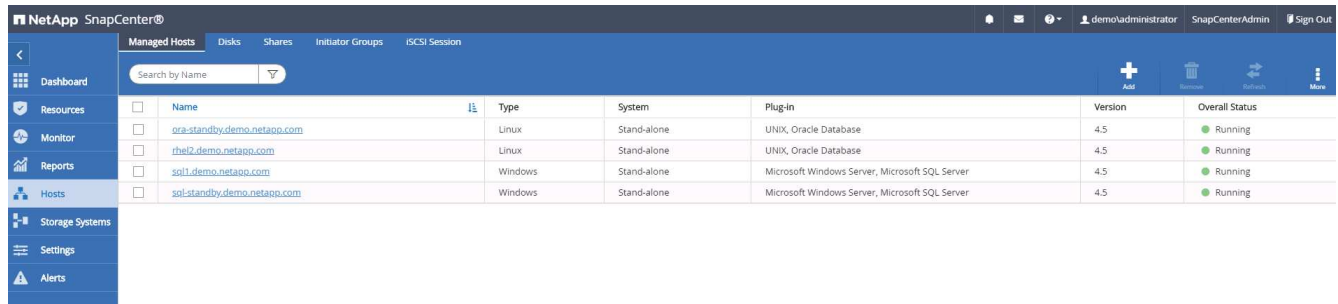
Confirm Fingerprint

Authenticity of the host cannot be determined

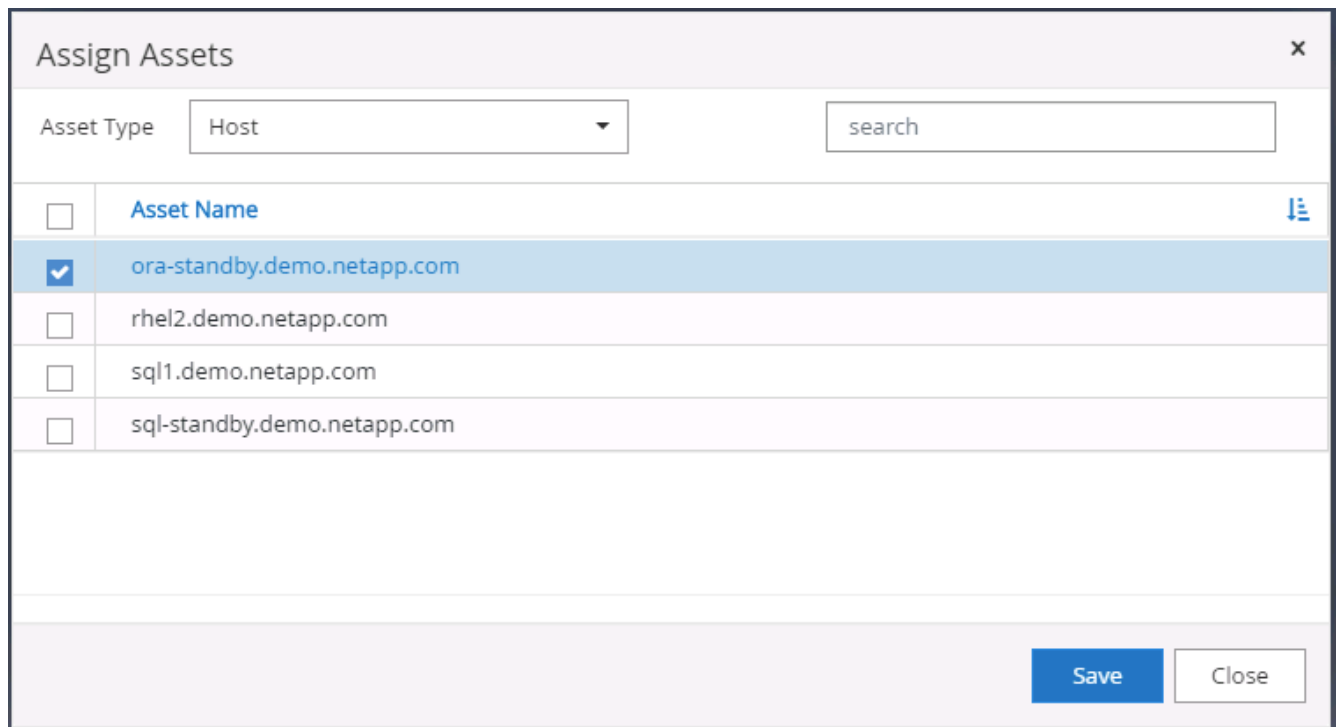
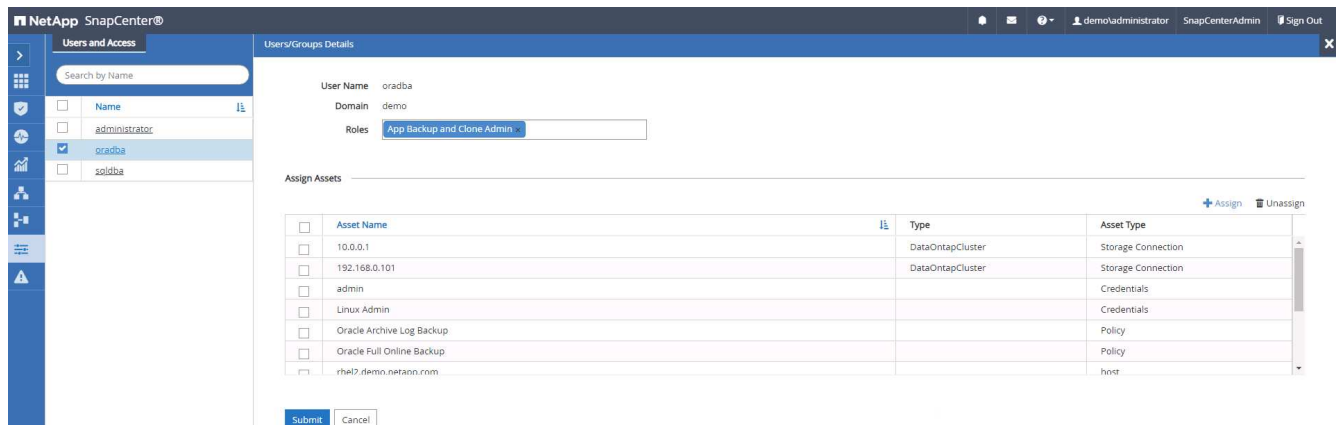
Host name	Fingerprint	Valid
ora-standby.demo.netapp.com	ssh-rsa 3072 5C:02:EF:6B:63:54:59:10:84:DF:4D:6B:AB:FB:61:67	

Confirm and Submit Close

- 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.



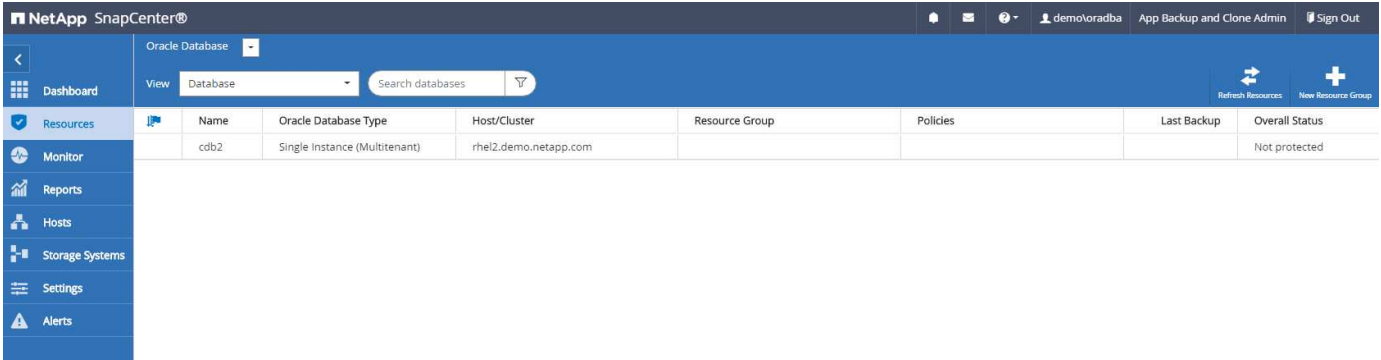
- Assign the newly added host to the proper database management user ID (in our case, oradba).



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.

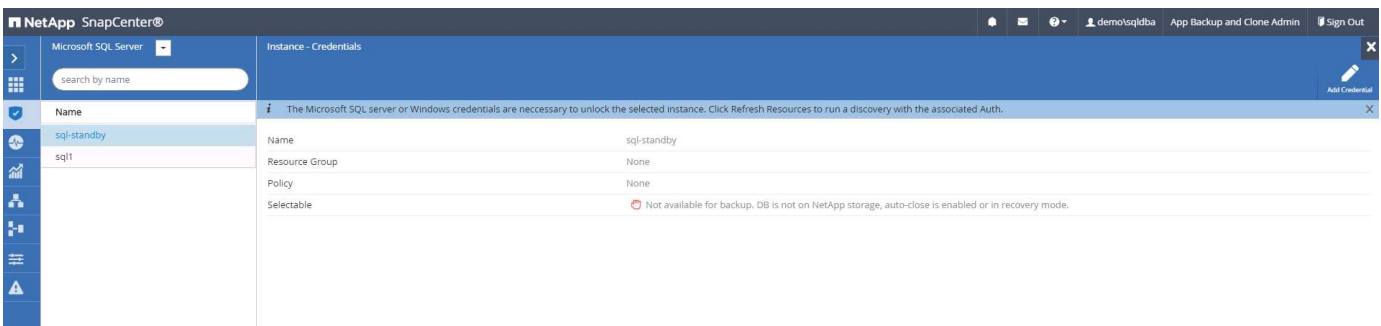


When the database is initially discovered, the Overall Status is shown as "Not protected." The previous screenshot shows an Oracle database not protected yet by a backup policy.

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.



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.



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.

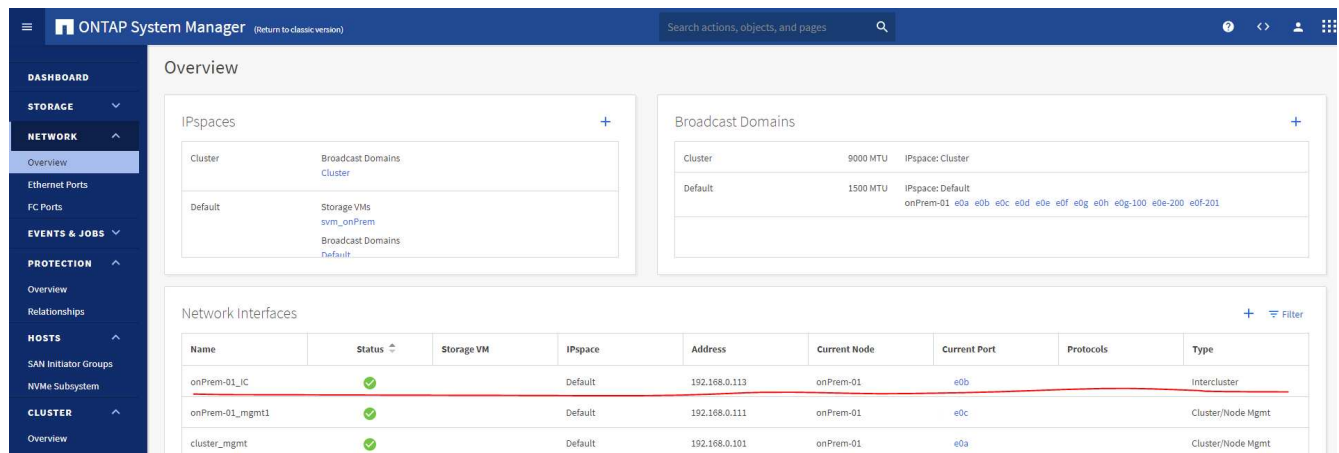


5. Setup storage cluster peering and DB volumes replication

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:



Target CVO cluster:

The screenshot shows the ONTAP System Manager Overview page. It features three main sections: IPspaces, Broadcast Domains, and Network Interfaces.

IPspaces:

Cluster	Broadcast Domains
Cluster	Cluster
Default	Storage VMS svm_hybridcvo Broadcast Domains Default

Broadcast Domains:

Cluster	9000 MTU	IPspace: Cluster
Cluster	9000 MTU	hybridcvo-01 e0b hybridcvo-02 e0b
Default	9001 MTU	IPspace: Default hybridcvo-01 e0a hybridcvo-02 e0a

Network Interfaces:

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

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

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

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

The screenshot shows the ONTAP System Manager Cluster Settings page. It features three main sections: UI Settings, Intercluster Settings, and Cluster Peers.

UI Settings:

- LOG LEVEL: DEBUG
- INACTIVITY TIMEOUT: 30 minutes

Intercluster Settings:

- Network Interfaces: IP ADDRESS 192.168.0.113 (✓)

Cluster Peers:

- PEERED CLUSTER NAME: hybridcvo (✓)
- Peer Cluster (highlighted)
- Generate Passphrase
- Manage Cluster Peers

Storage VM Peers:

- PEERED STORAGE VMS: 1 (✓)


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

The screenshot shows the ONTAP System Manager interface. The left sidebar contains navigation menus for Dashboard, Storage, Network, Events & Jobs, Protection, Hosts, and Cluster. The main area displays the 'Volumes' page for the 'rhel2_u03' volume. At the top, there are buttons for '+ Add', 'Delete', 'Protect' (circled in red), and 'More'. Below this is a table of volumes, with 'rhel2_u03' selected. To the right of the table, the 'Overview' tab is active, showing various properties of the volume. On the far right, there are two charts: 'Capacity' and 'Performance'.

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

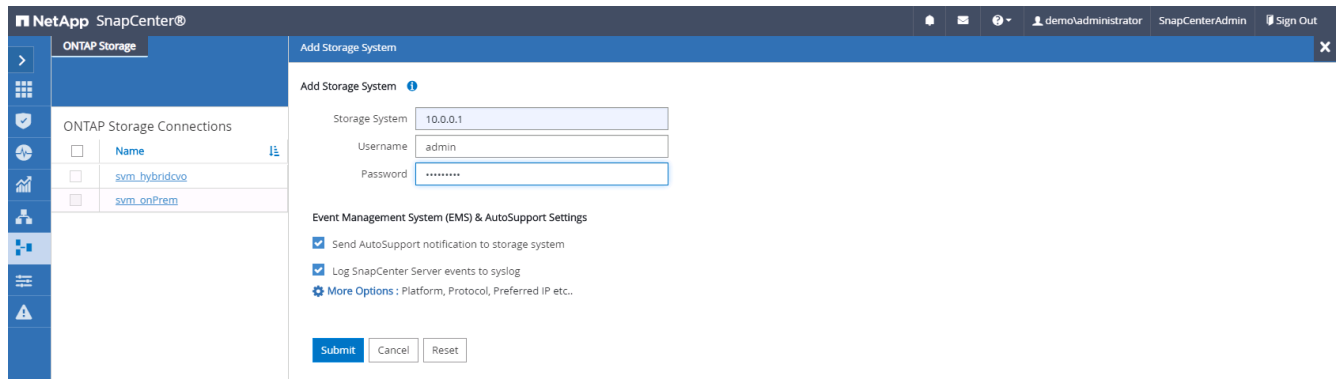
The screenshot shows the 'Protect Volumes' dialog box. The 'Protection Policy' dropdown is set to 'Asynchronous'. The 'Source' section shows 'Cluster: onPrem' and 'Storage VM: svm_onPrem'. The 'Destination' section shows 'Cluster: hybridcvo' and 'Storage VM: svm_hybridcvo'. Under 'Destination Settings', it indicates '2 matching labels' and shows 'Volume Name' fields with 'PREFIX: vol_' and 'SUFFIX: _dest'. The 'Initialize relationship' checkbox is checked, and 'Enable FabricPool' is unchecked. 'Save' and 'Cancel' buttons are at the bottom.

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

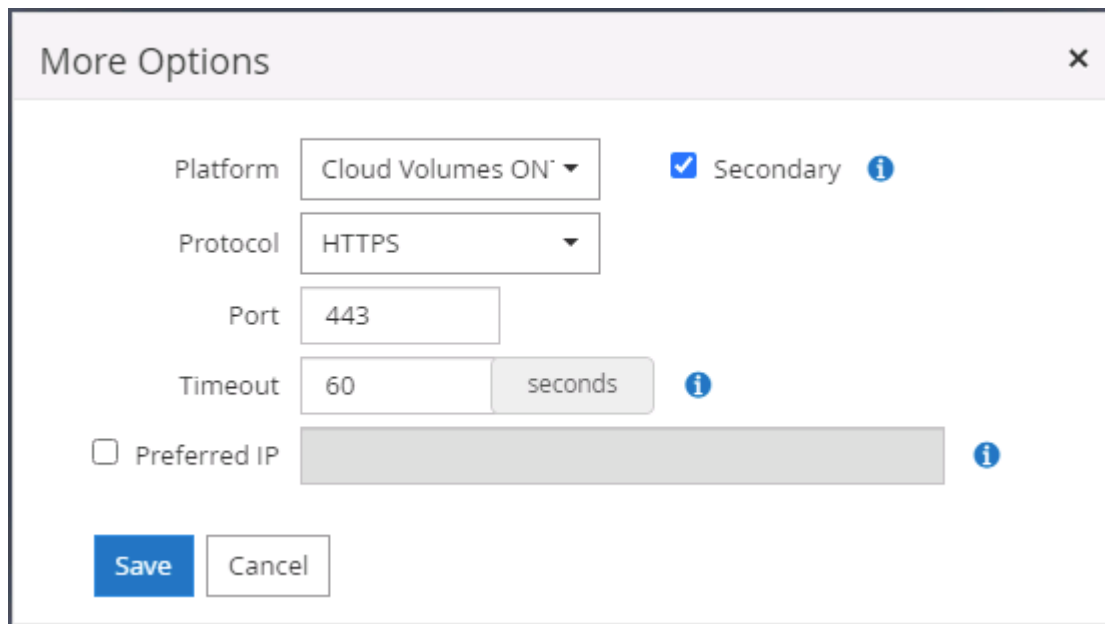


6. Add CVO database storage SVM to SnapCenter

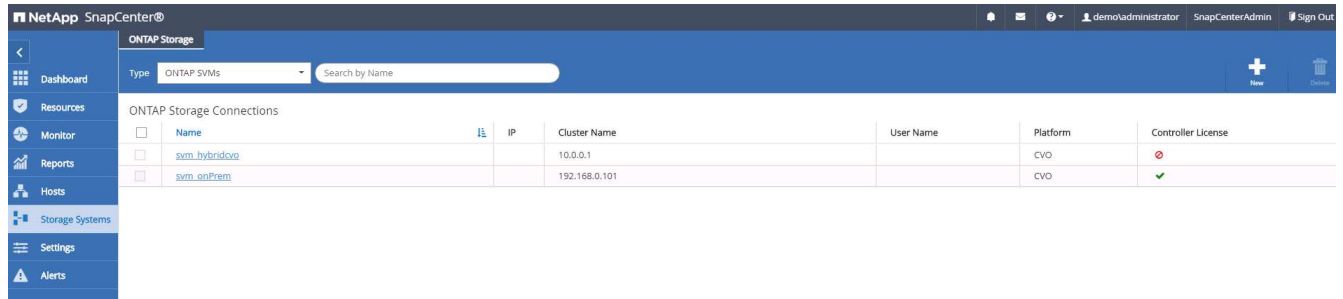
1. Log into SnapCenter with a user ID with SnapCenterAdmin privileges.
2. Click the Storage System tab from the menu, and then click New to add a CVO storage SVM that hosts replicated target database volumes to SnapCenter. Enter the cluster management IP in the Storage System field, and enter the appropriate username and password.



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



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

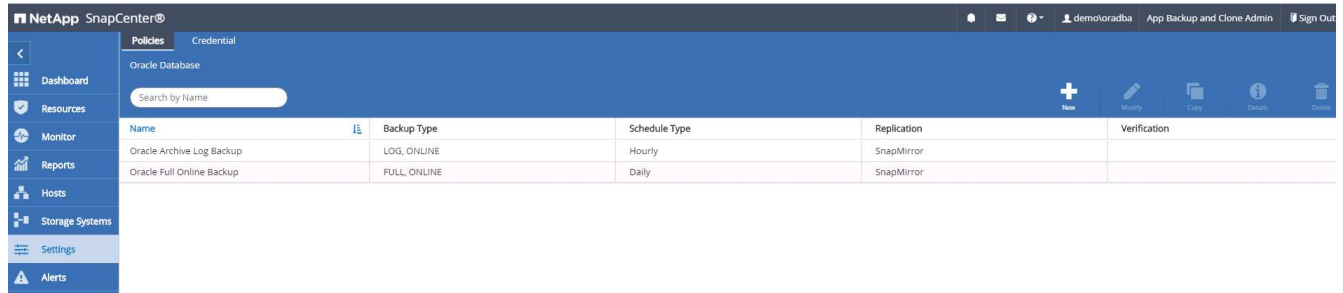


7. Setup database backup policy in SnapCenter

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

Create a full database backup policy for Oracle

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



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

Modify Oracle Database Backup Policy ×

- 1 Name**
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Provide a policy name

Policy name i

Details

3. Select the backup type and schedule frequency.

Modify Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select Oracle database backup options

Choose backup type

Online backup

- Datafiles, control files, and archive logs
- Datafiles and control files
- Archive logs

Offline backup ?

- Mount
- Shutdown
- Save state of PDBs ?

Choose schedule frequency

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

On demand

Hourly

Daily

Previous Next

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

Modify Oracle Database Backup Policy x

- 1 Name
- 2 Backup Type
- 3 Retention**
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Retention settings ⓘ

Daily retention settings

Data backup retention settings ⓘ

Total Snapshot copies to keep

Keep Snapshot copies for days

Archive Log backup retention settings

Total Snapshot copies to keep

Keep Snapshot copies for days

Previous Next

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

Modify Oracle Database Backup Policy ×

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 x

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script**
- 6 Verification
- 7 Summary

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

Prescript full path

Prescript arguments

Postscript full path

Postscript arguments

Script timeout

7. Run backup verification if desired.

Modify Oracle Database Backup Policy x

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script
- 6 Verification**
- 7 Summary

Select the options to run backup verification

Run Verifications for following backup schedules

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

Daily

Verification script commands

Script timeout secs

Prescript full path

Prescript arguments

Postscript full path

Postscript arguments

8. Summary.

Modify Oracle Database Backup Policy

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

Previous Finish

Create a database log backup policy for Oracle

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

New Oracle Database Backup Policy x

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Provide a policy name

Policy name i

Details

PreviousNext

3. Select the backup type and schedule frequency.

New Oracle Database Backup Policy

- 1 Name
- 2 Backup Type**
- 3 Retention
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Select Oracle database backup options

Choose backup type

- Online backup
 - Datafiles, control files, and archive logs
 - Datafiles and control files
 - Archive logs
- Offline backup i
 - Mount
 - Shutdown
 - Save state of PDBs i

Choose schedule frequency

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

- On demand
- Hourly
- Daily

Previous Next

4. Set the log retention period.

New Oracle Database Backup Policy ✕

- 1 Name
- 2 Backup Type
- 3 Retention**
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Retention settings ?

Hourly retention settings

Data backup retention settings ?

Total Snapshot copies to keep

Keep Snapshot copies for days

Archive Log backup retention settings

Total Snapshot copies to keep

Keep Snapshot copies for days

Previous Next

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

New Oracle Database Backup Policy ×

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select secondary replication options ⓘ

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label: Hourly ⓘ

Error retry count: 3 ⓘ

Previous Next

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

New Oracle Database Backup Policy x

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

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

Prescript full path

Prescript arguments

Postscript full path

Postscript arguments

Script timeout

7. Specify any backup verification scripts.

New Oracle Database Backup Policy ✕

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script
- 6 Verification**
- 7 Summary

Select the options to run backup verification

Run Verifications for following backup schedules

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

Verification script commands

Script timeout secs


Prescript full path

Prescript arguments

Postscript full path

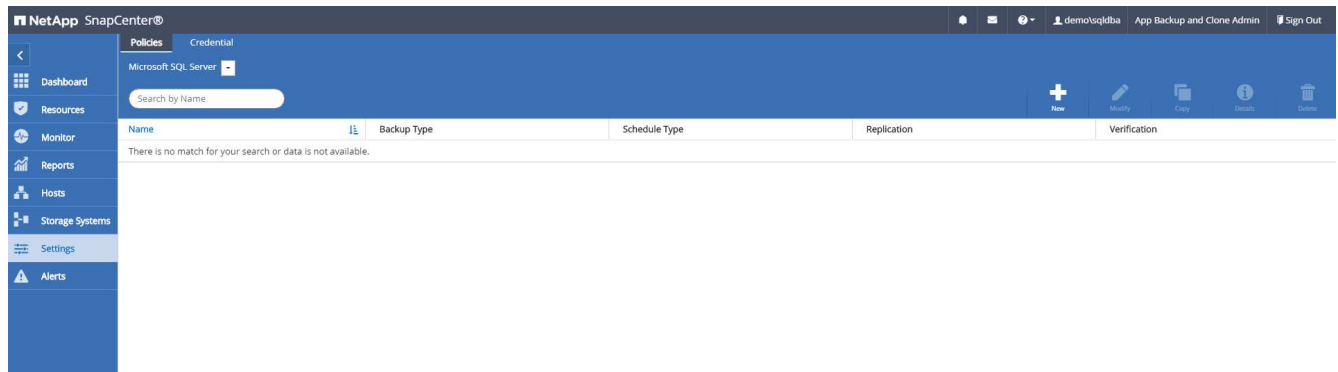
Postscript arguments

8. Summary.



Create a full database backup policy for SQL

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



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

New SQL Server Backup Policy x

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Provide a policy name

Policy name i

Details

Backup all data and log files

PreviousNext

3. Define the backup option and schedule frequency. For SQL Server configured with an availability group, a preferred backup replica can be set.

New SQL Server Backup Policy x

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select SQL server backup options

Choose backup type

Full backup and log backup

Full backup

Log backup

Copy only backup i

Maximum databases backed up per Snapshot copy: i

Availability Group Settings v

Schedule frequency

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

On demand

Hourly

Daily

Weekly

Monthly

4. Set the backup retention period.

New SQL Server Backup Policy x

- 1 Name
- 2 Backup Type
- 3 Retention**
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Retention settings

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

Keep log backups applicable to last full backups

Keep log backups applicable to last days

Full backup retention settings ⓘ

Daily

Total Snapshot copies to keep

Keep Snapshot copies for days

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

New SQL Server Backup Policy x

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select secondary replication options i

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label i

Error retry count i

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

New SQL Server Backup Policy x

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script**
- 6 Verification
- 7 Summary

Specify optional scripts to run before performing a backup job

Prescript full path

Prescript arguments

Specify optional scripts to run after performing a backup job

Postscript full path

Postscript arguments

Script timeout

7. Specify the options to run backup verification.

New SQL Server Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select the options to run backup verification

Run verifications for the following backup schedules

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

Daily

Database consistency checks options

Limit the integrity structure to physical structure of the database (PHYSICAL_ONLY)

Suppress all information message (NO_INFOMSGS)

Display all reported error messages per object (ALL_ERRORMSGs)

Do not check non-clustered indexes (NOINDEX)

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

Log backup

Verify log backup. i

Verification script settings

Script timeout secs

Previous Next

8. Summary.

New SQL Server Backup Policy
×

- 1 Name
- 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 x

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Provide a policy name

Policy name

Details

PreviousNext

2. Define the log backup option and schedule frequency. For SQL Server configured with a availability group, a preferred backup replica can be set.

New SQL Server Backup Policy x

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Select SQL server backup options

Choose backup type

Full backup and log backup

Full backup

Log backup

Copy only backup i

Maximum databases backed up per Snapshot copy: i

Availability Group Settings▼

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

PreviousNext

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

New SQL Server Backup Policy ×

- 1 Name
- 2 Backup Type
- 3 Retention**
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Log backup retention settings

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

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

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

New SQL Server Backup Policy ✕

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select secondary replication options ⓘ

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label: Hourly ⓘ

Error retry count: 3 ⓘ

Previous Next

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

New SQL Server Backup Policy ×

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script**
- 6 Verification
- 7 Summary

Specify optional scripts to run before performing a backup job

Prescript full path

Prescript arguments

Specify optional scripts to run after performing a backup job

Postscript full path

Postscript arguments

Script timeout

6. Summary.

New SQL Server Backup Policy

- 1 Name
- 2 Backup Type
- 3 Retention
- 4 Replication
- 5 Script
- 6 Verification
- 7 Summary

Summary

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

Previous
Finish

8. Implement backup policy to protect database

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


Create a resource group for full backup of Oracle

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

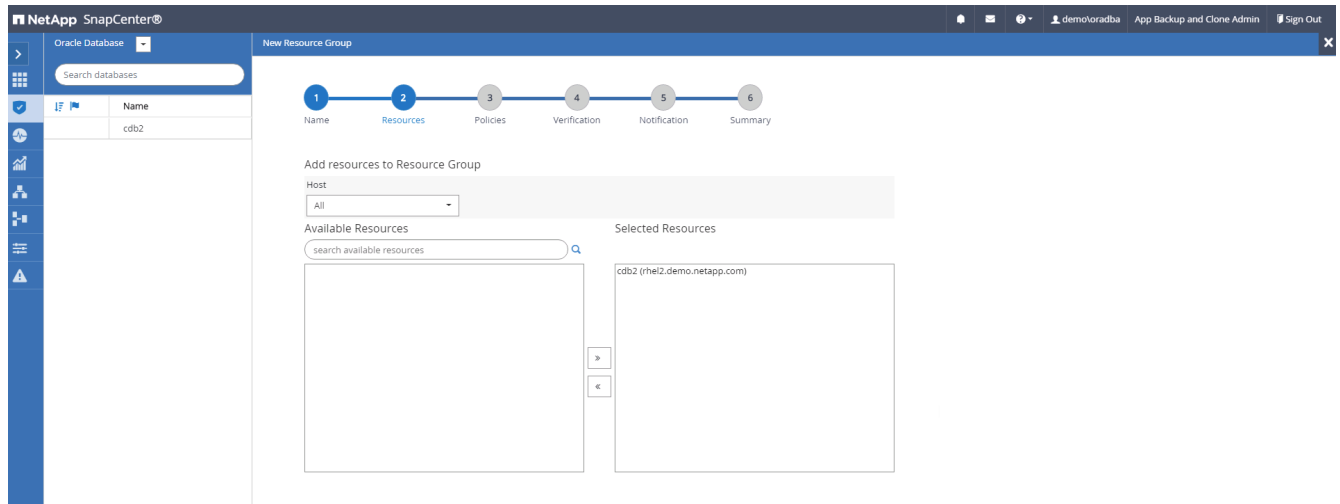
The screenshot shows the NetApp SnapCenter interface. The top navigation bar includes 'NetApp SnapCenter', user 'demo/oradba', and 'App Backup and Clone Admin'. The left sidebar has a 'Resources' tab selected. The main content area shows a table with the following data:

Name	Oracle Database Type	Host/Cluster	Resource Group	Policies	Last Backup	Overall Status
cdb2	Single Instance (Multitenant)	rhe2.demo.netapp.com				Not protected

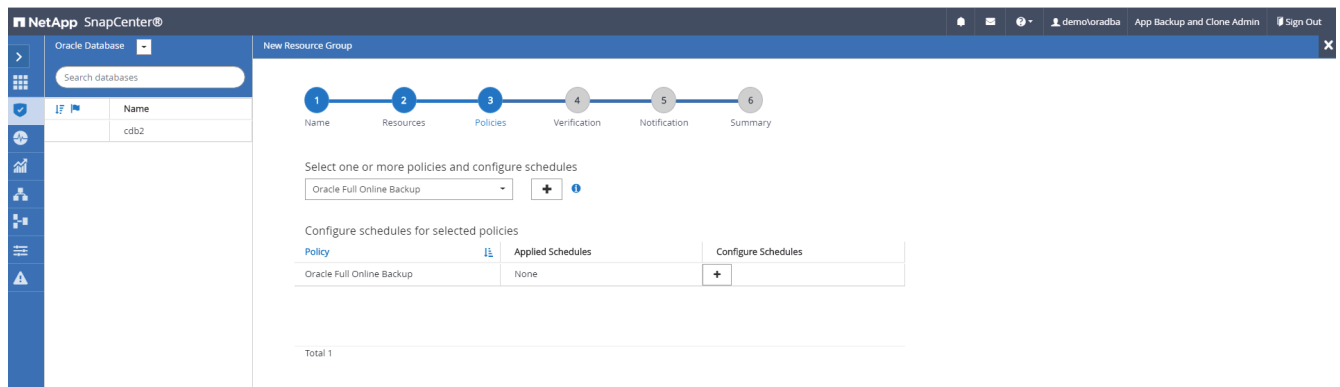
2. Provide a name and tags for the resource group. You can define a naming format for the Snapshot copy and bypass the redundant archive log destination if configured.




3. Add database resources to the resource group.




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




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



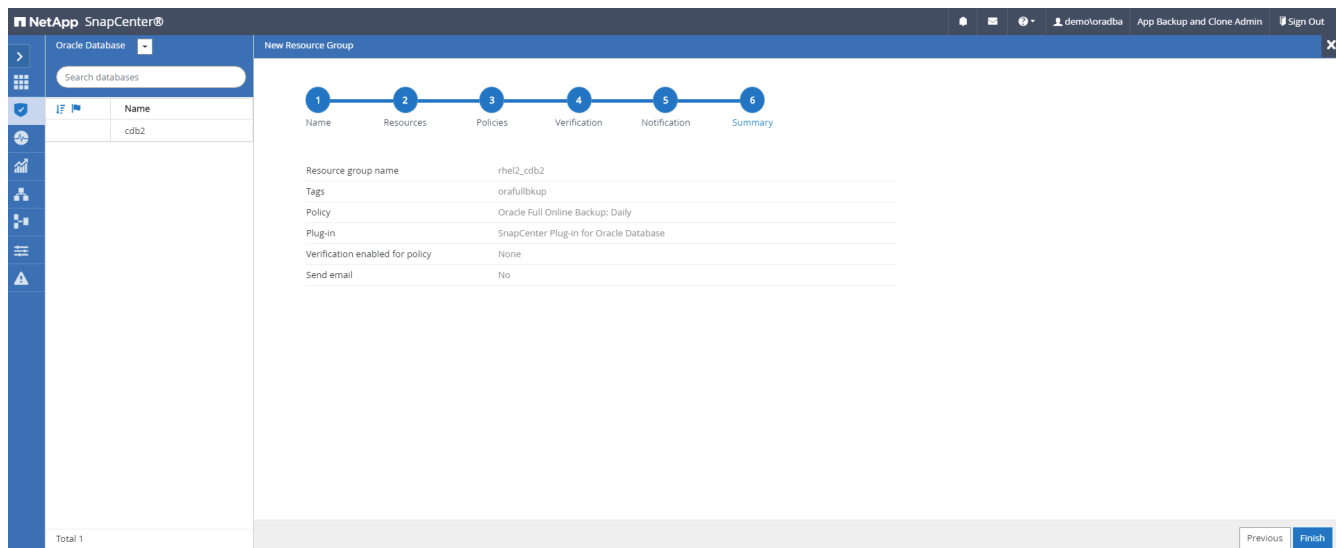
6. Click Load Locators to load the source and destination volume.



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

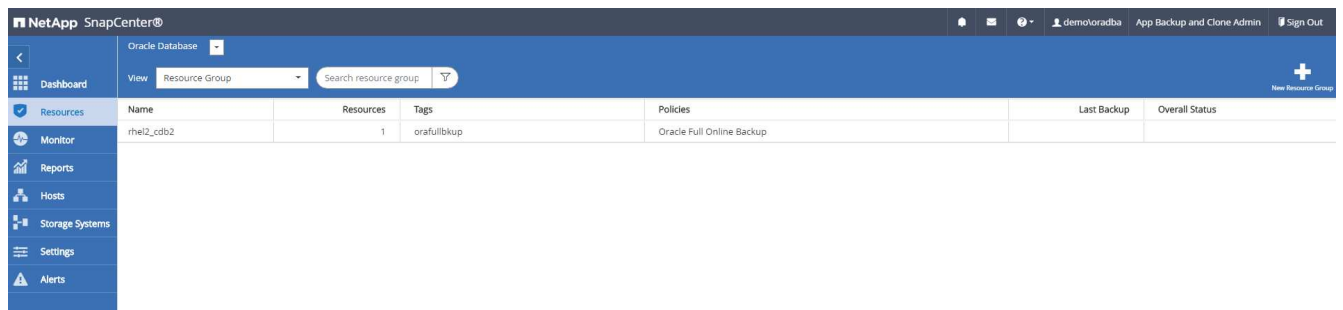


8. Summary.




Create a resource group for log backup of Oracle


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




2. Provide a name and tags for the resource group. You can define a naming format for the Snapshot copy and bypass the redundant archive log destination if configured.



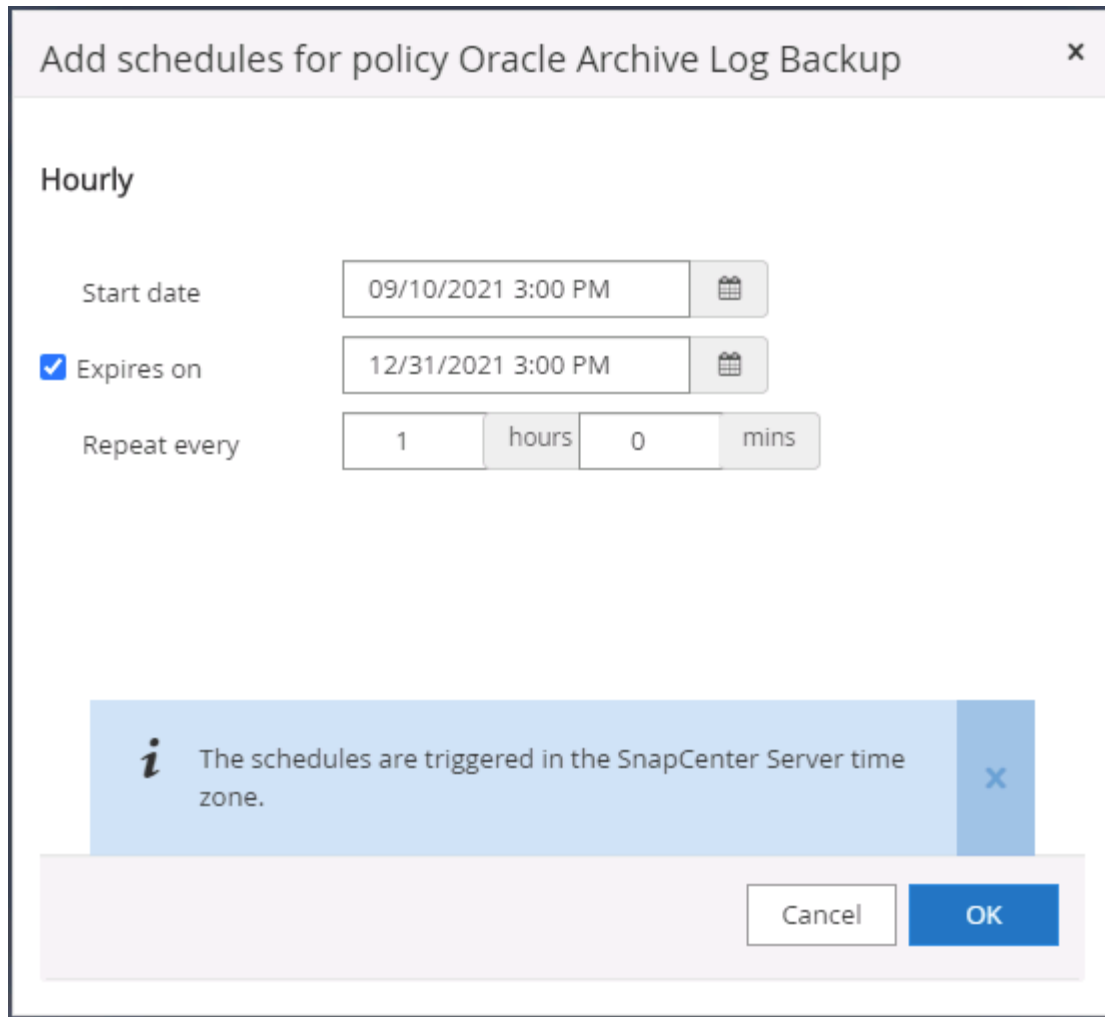
3. Add database resources to the resource group.



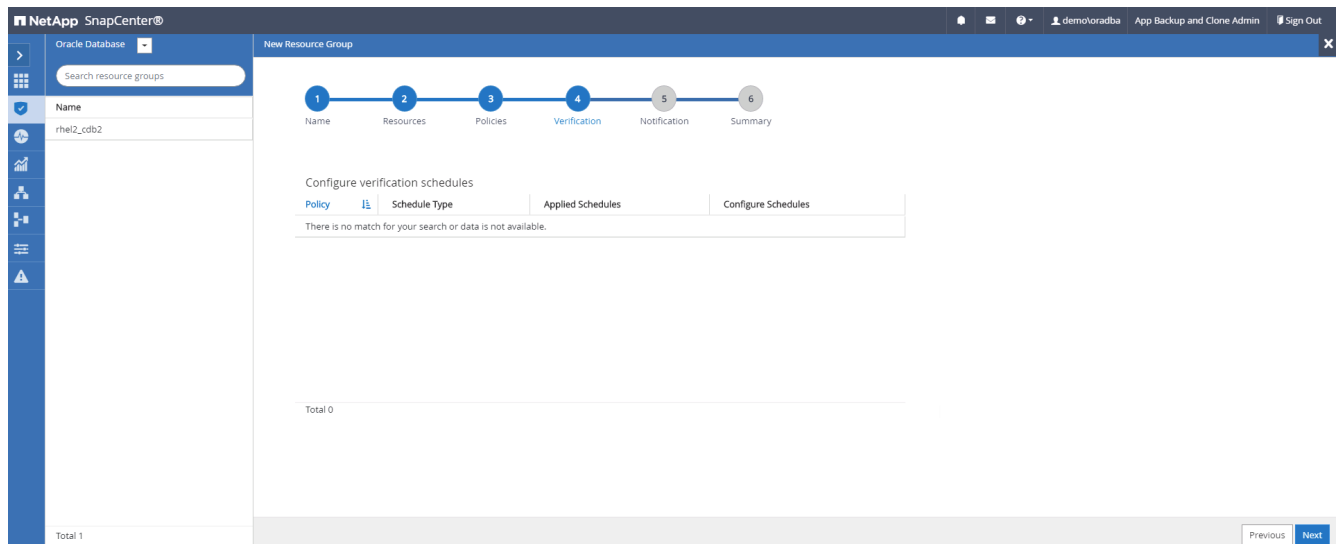
4. Select a log backup policy created in section 7 from the drop-down list.



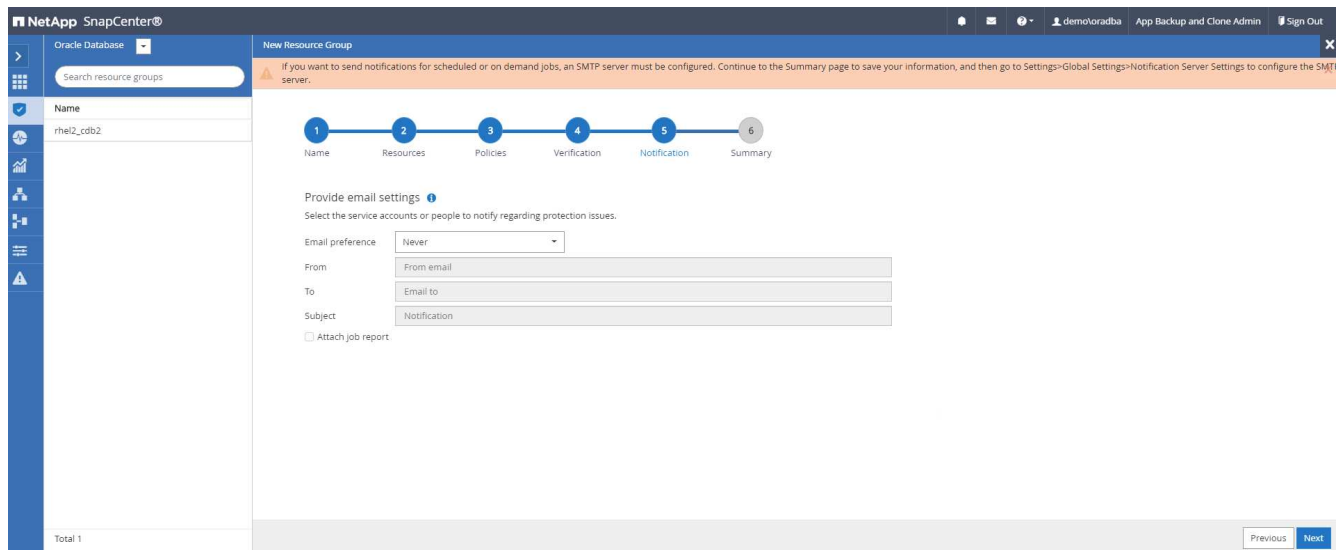
5. Click on the (+) sign to configure the desired backup schedule.




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



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




8. Summary.

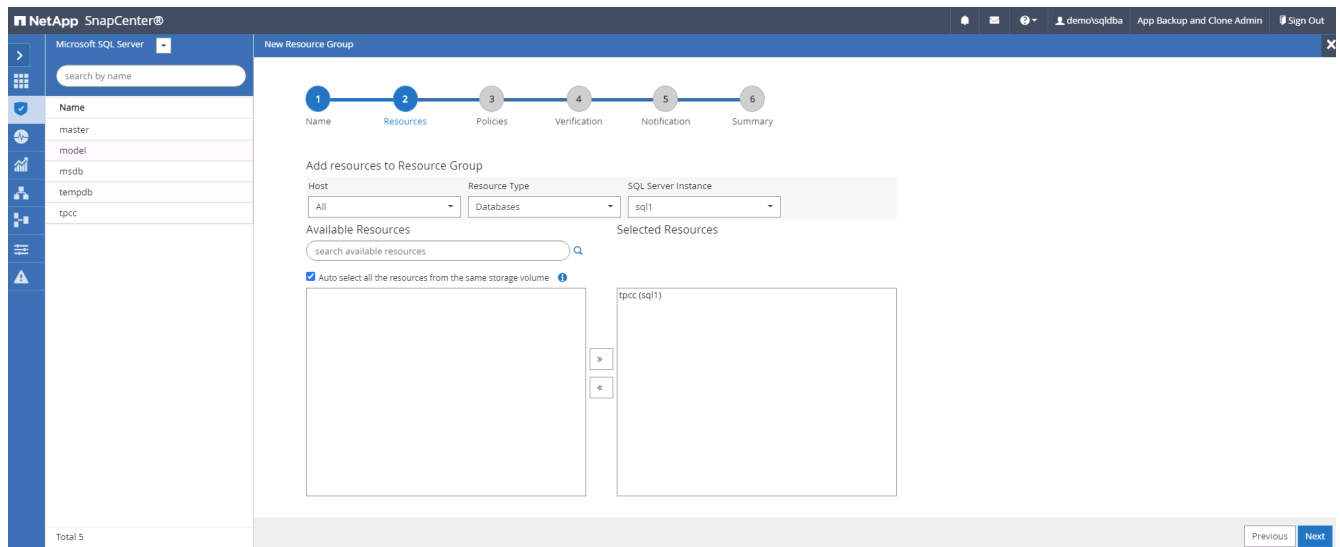


Create a resource group for full backup of SQL Server


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



2. Select the database resources to be backed up.



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



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

Add schedules for policy SQL Server Full Backup

Daily

Start date 09/10/2021 6:20 PM

Expires on 12/31/2021 6:20 PM

Repeat every 1 days

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

Cancel OK

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

NetApp SnapCenter

Microsoft SQL Server

New Resource Group

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

Select the verification servers

Verification server Select one or more servers

Load secondary locators to verify backups on secondary Load locators

Secondary storage location: SnapVault or SnapMirror

Source Volume	Destination Volume
svm_onPrem:sql1_data	svm_hybridcovsql1_data_dr
svm_onPrem:sql1_log	svm_hybridcovsql1_log_dr


Configure verification schedules

Policy Schedule Type Applied Schedules Configure Schedules


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

Previous Next

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




7. Summary.




Create a resource group for log backup of SQL Server

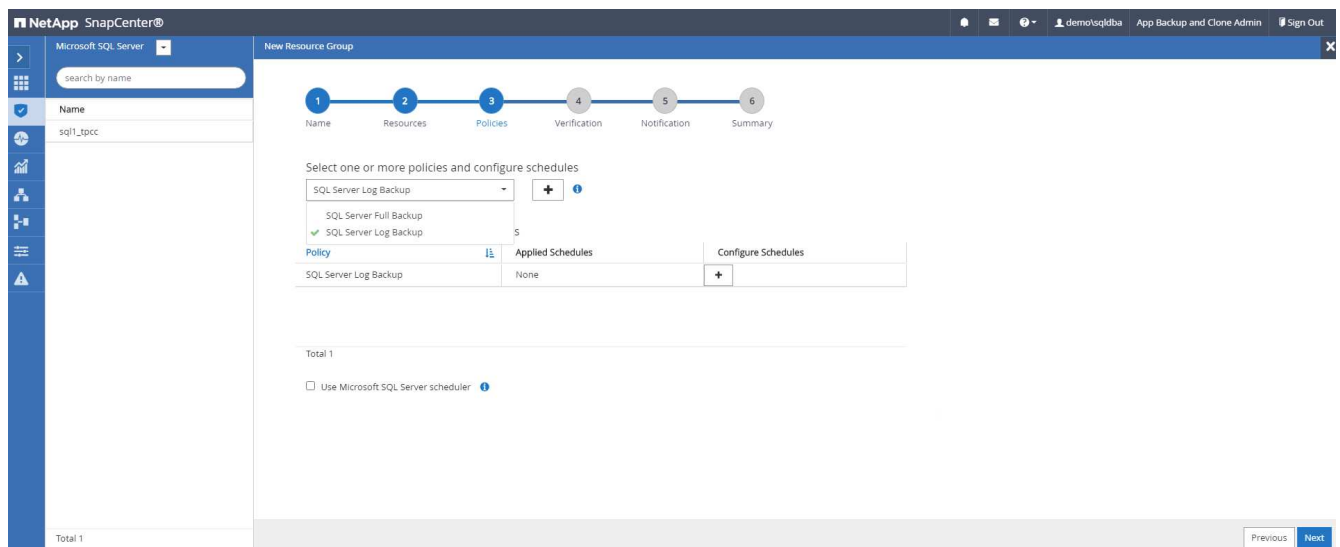
1. Log into SnapCenter with a database management user ID, and navigate to the Resources tab. In the View drop-down list, choose either a Database or Resource Group to launch the resource group creation workflow. Provide the name and tags for the resource group. You can define a naming format for the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




2. Select the database resources to be backed up.



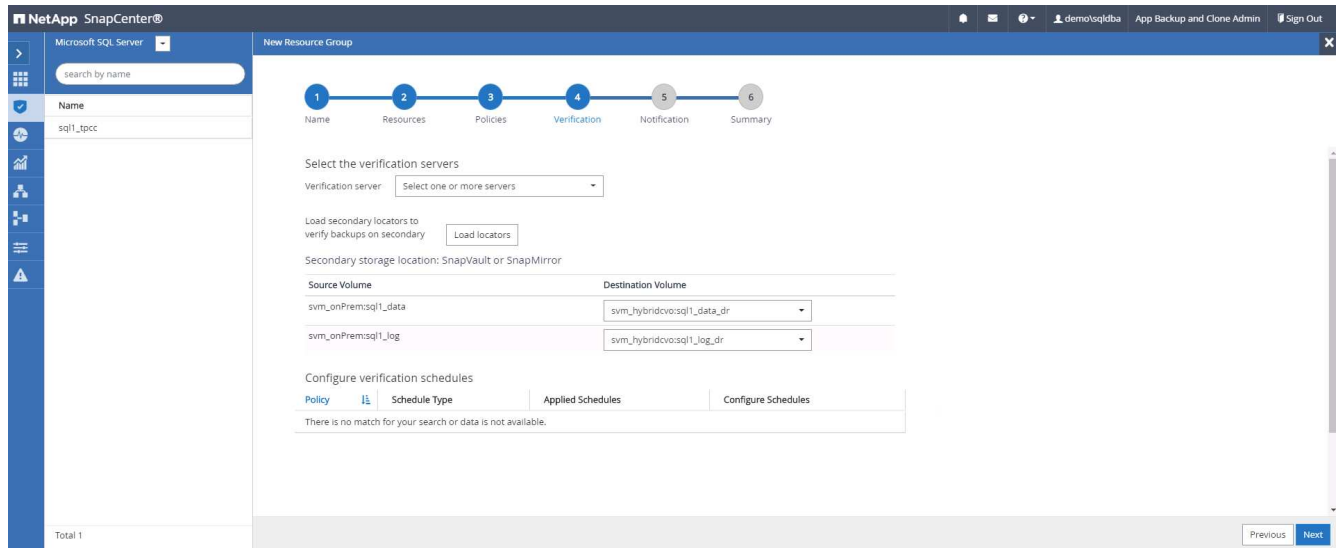
3. Select a SQL log backup policy created in section 7.



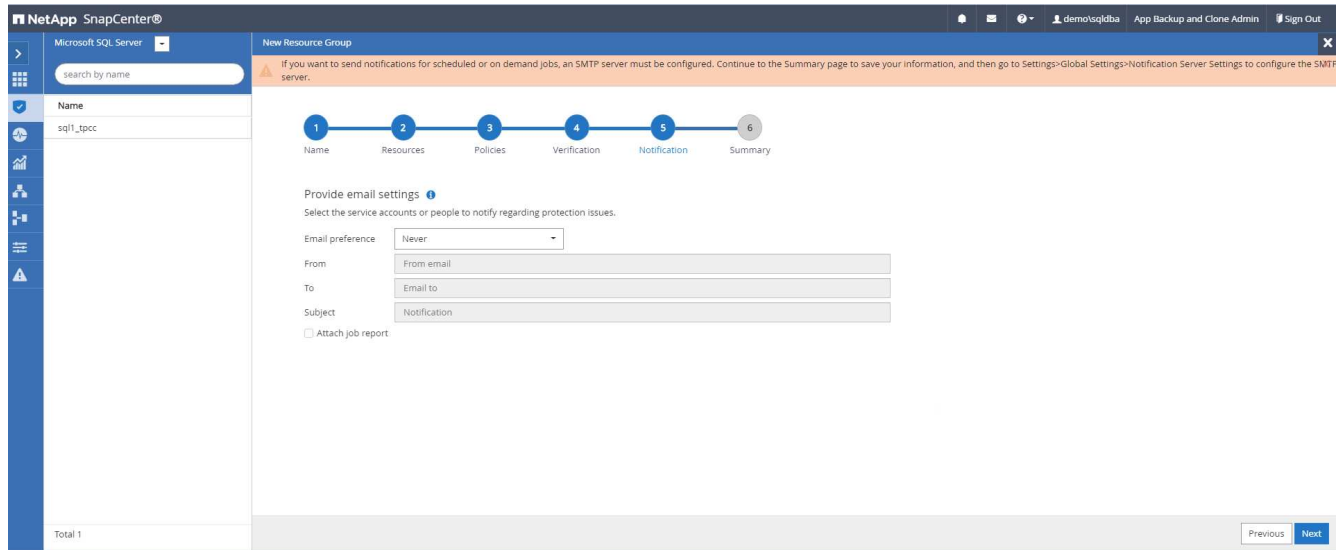
4. Add exact timing for the backup as well as the frequency.



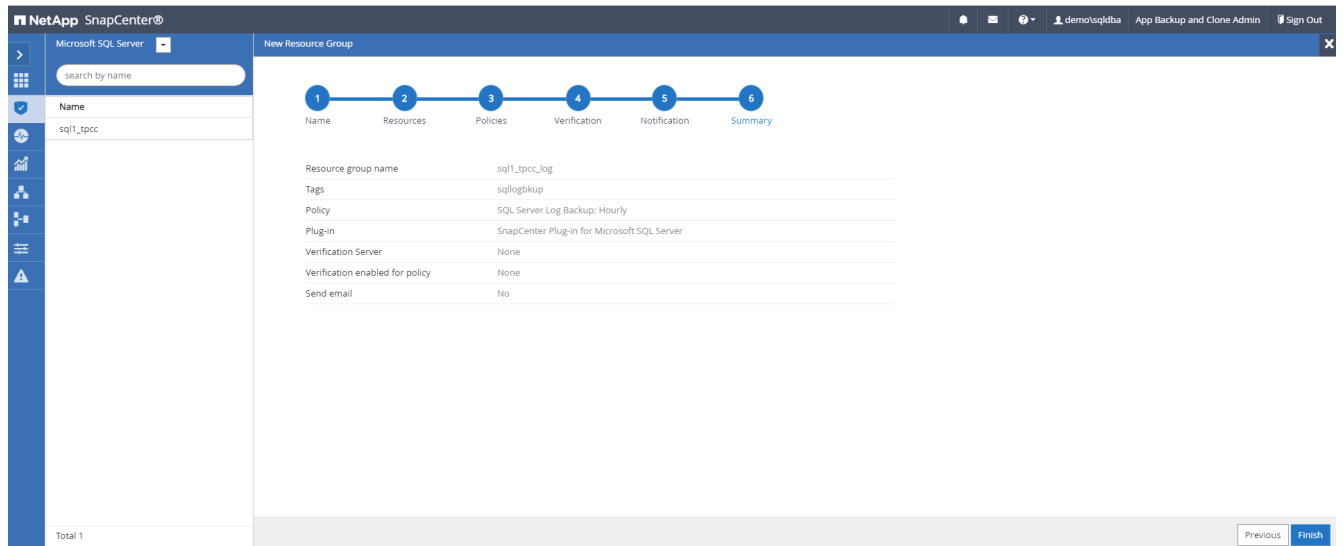
5. Choose the verification server for the backup on secondary if backup verification is to be performed. Click the Load Locator to populate the secondary storage location.



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



7. Summary.



9. Validate backup

After database backup resource groups are created to protect database resources, the backup jobs runs according to the predefined schedule. Check the job execution status under the Monitor tab.

ID	Status	Name	Start date	End date	Owner
532	Success	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'sqlqdba
528	Success	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'sqlqdba
524	Success	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'sqlqdba
521	Success	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'sqlqdba
517	Success	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'sqlqdba
513	Success	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'sqlqdba
509	Success	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'sqlqdba
503	Success	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'sqlqdba

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

public cloud.

Backup Name	Count	Type	End Date	Verified	Mounted	RMAN Cataloged	SCN
rhe12_cdb2_09-23-2021_14.35.03.3242_1	1	Log	09/23/2021 2:35:45 PM	Not Applicable	False	Not Cataloged	6872761
rhe12_cdb2_09-23-2021_14.35.03.3242_0	1	Data	09/23/2021 2:35:30 PM	Unverified	False	Not Cataloged	6872715
rhe12_cdb2_09-22-2021_14.35.02.0014_1	1	Log	09/22/2021 2:35:24 PM	Not Applicable	False	Not Cataloged	6737479
rhe12_cdb2_09-22-2021_14.35.02.0014_0	1	Data	09/22/2021 2:35:14 PM	Unverified	False	Not Cataloged	6737395
rhe12_cdb2_09-21-2021_14.35.02.1884_1	1	Log	09/21/2021 2:35:35 PM	Not Applicable	False	Not Cataloged	6598735

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

Getting Started with AWS public cloud

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

AWS public cloud



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

1. Pre-flight check

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

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

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



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

Deploy the Cloud Manager connector

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



[Continue to Cloud Manager](#)


Log In to NetApp Cloud Central

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


LOGIN

[Forgot your password?](#)


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



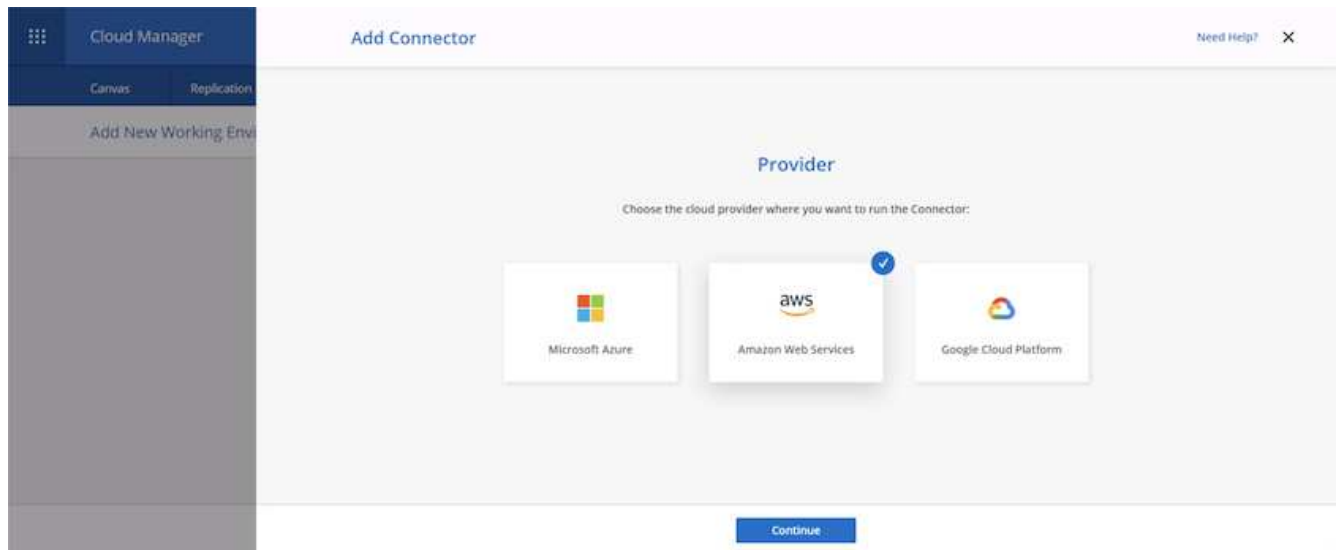
3. Click "Add Working Environment" and choose Cloud Volumes ONTAP in AWS. Here, you also choose whether you want to deploy a single node system or a high availability pair. I have chosen to deploy a high availability pair.



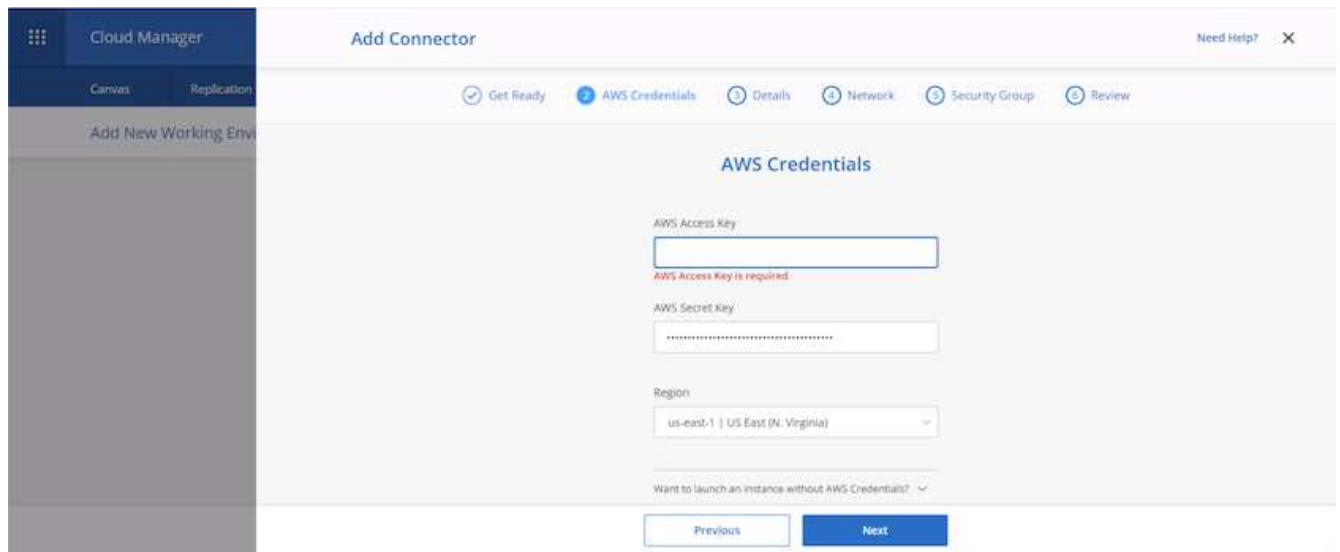
4. If no connector has been created, a pop-up appears asking you to create a connector.



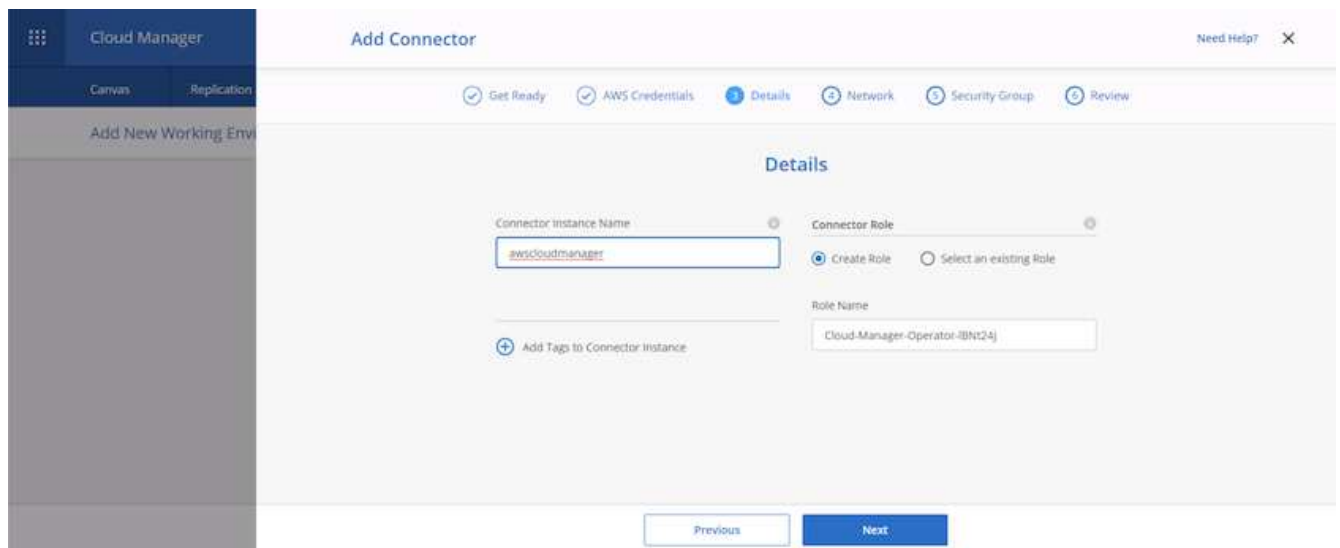
5. Click Lets Start, and then choose AWS.



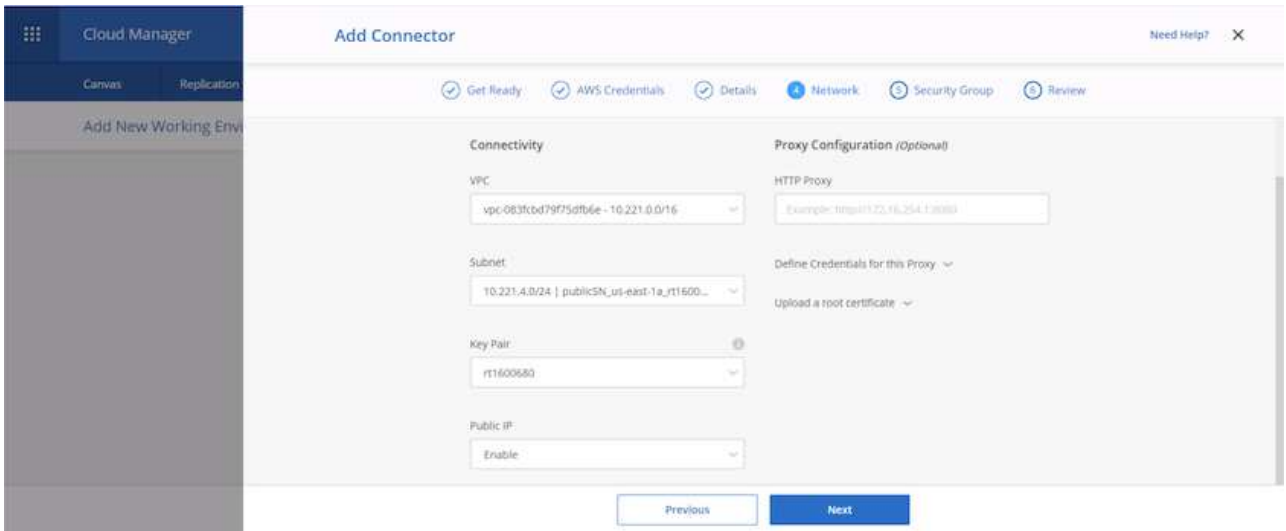
6. Enter your secret key and access key. Make sure that your user has the correct permissions outlined on the [NetApp policies page](#).



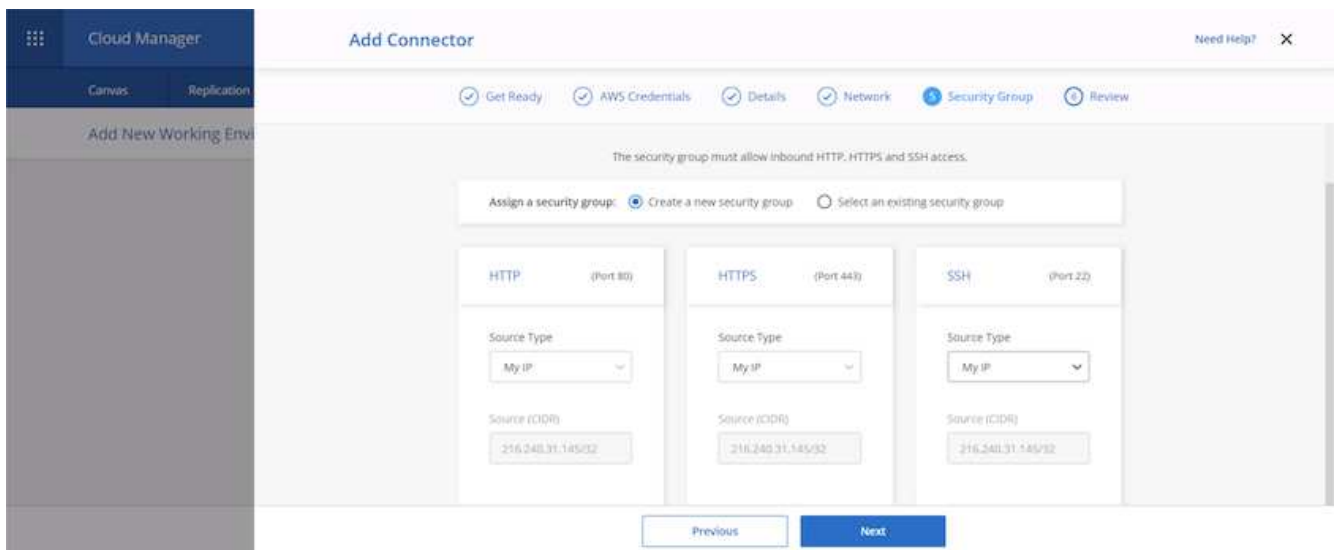
7. Give the connector a name and either use a predefined role as described on the [NetApp policies page](#) or ask Cloud Manager to create the role for you.



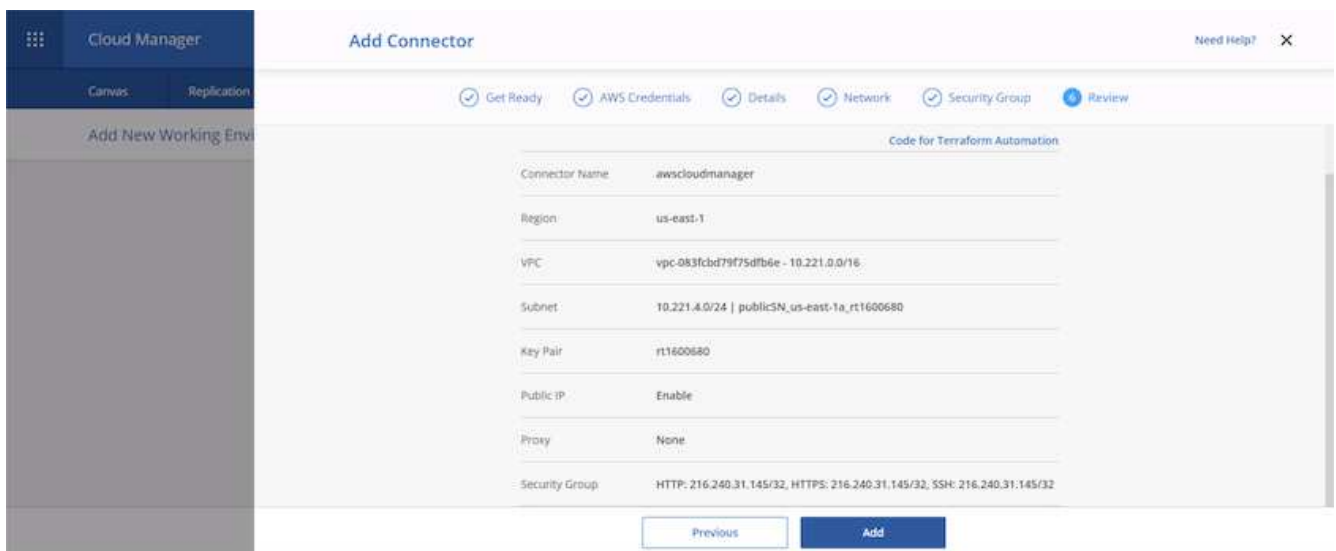
8. Give the networking information needed to deploy the connector. Verify that outbound internet access is enabled by:
 - a. Giving the connector a public IP address
 - b. Giving the connector a proxy to work through
 - c. Giving the connector a route to the public internet through an Internet Gateway




9. Provide communication with the connector via SSH, HTTP, and HTTPS by either providing a security group or creating a new security group. I have enabled access to the connector from my IP address only.



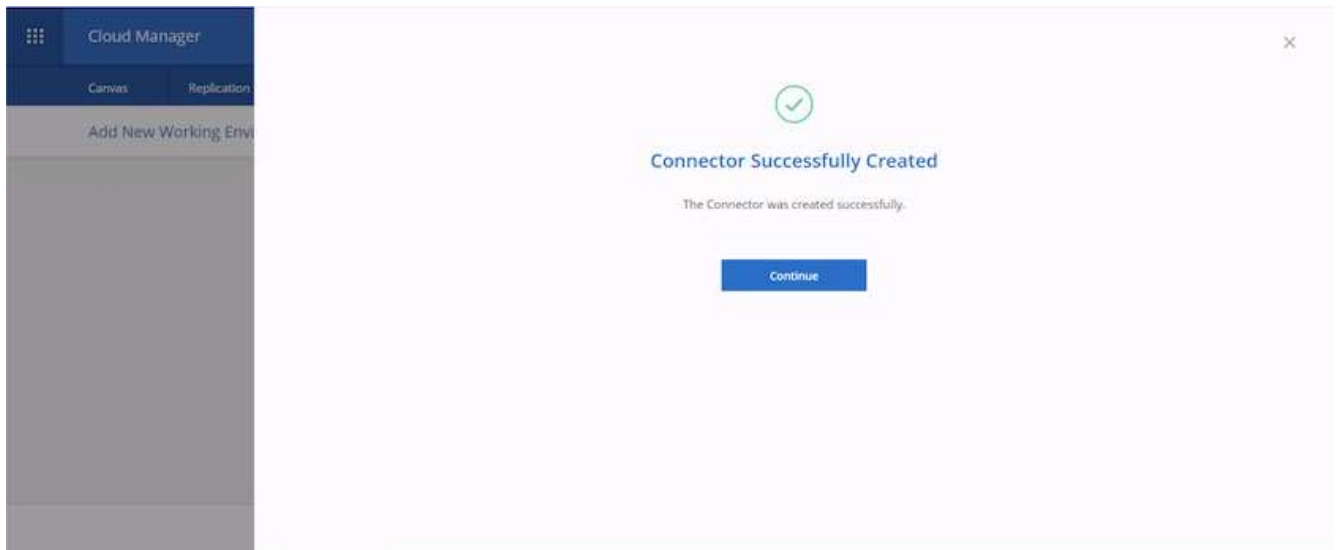
10. Review the information on the summary page and click Add to deploy the connector.



11. The connector now deploys using a cloud formation stack. You can monitor its progress from Cloud Manager or through AWS.

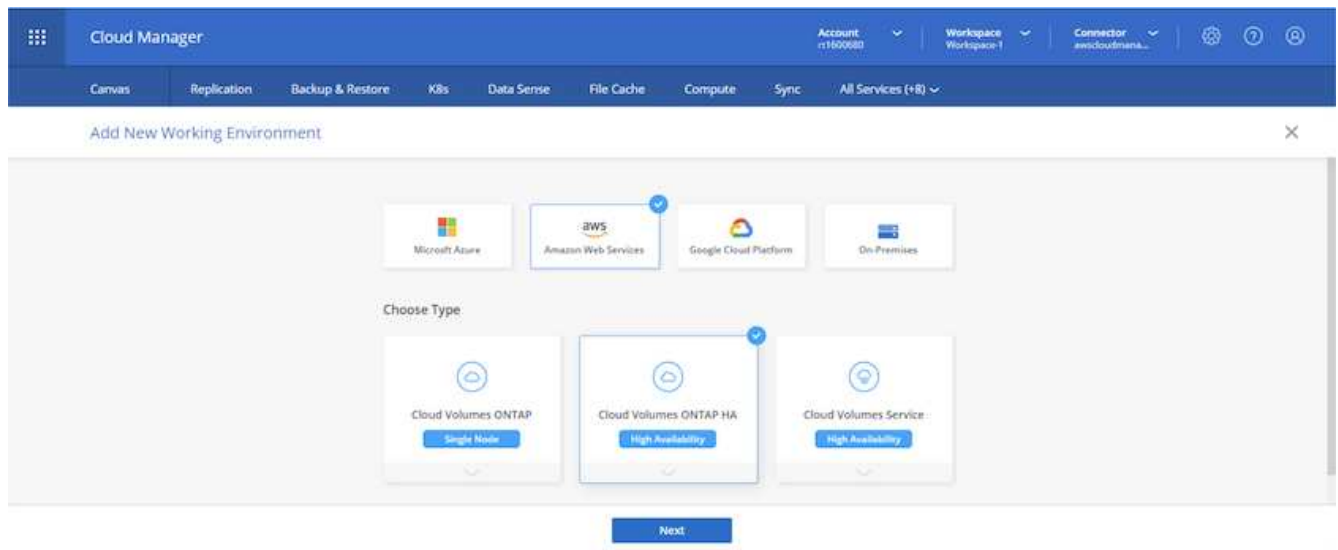


12. When the deployment is complete, a success page appears.

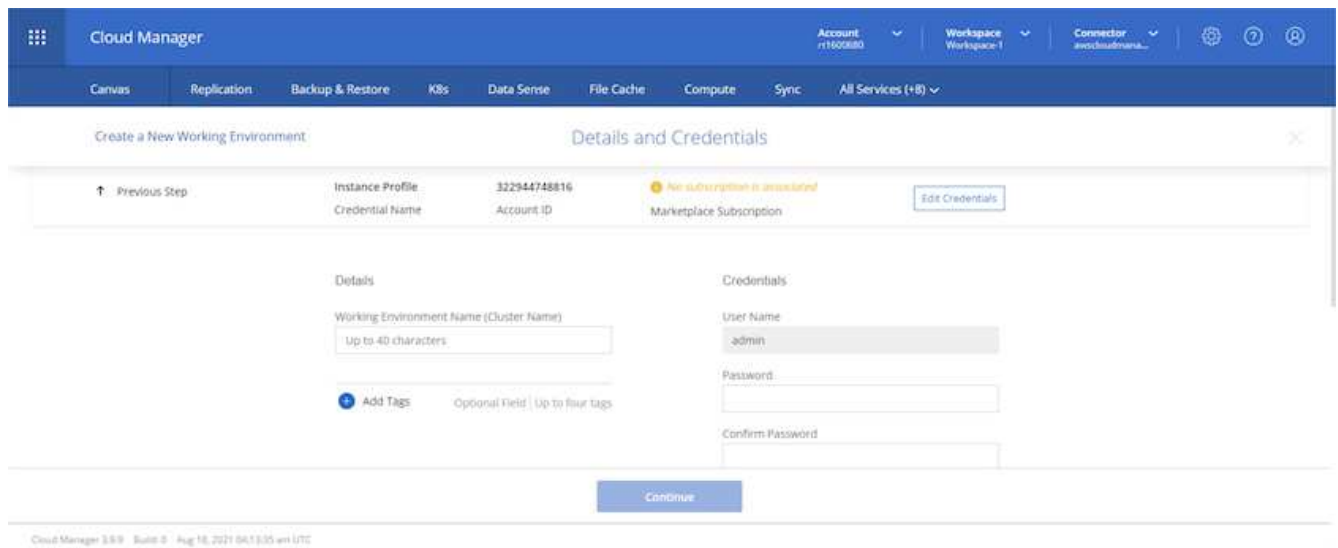


Deploy Cloud Volumes ONTAP


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



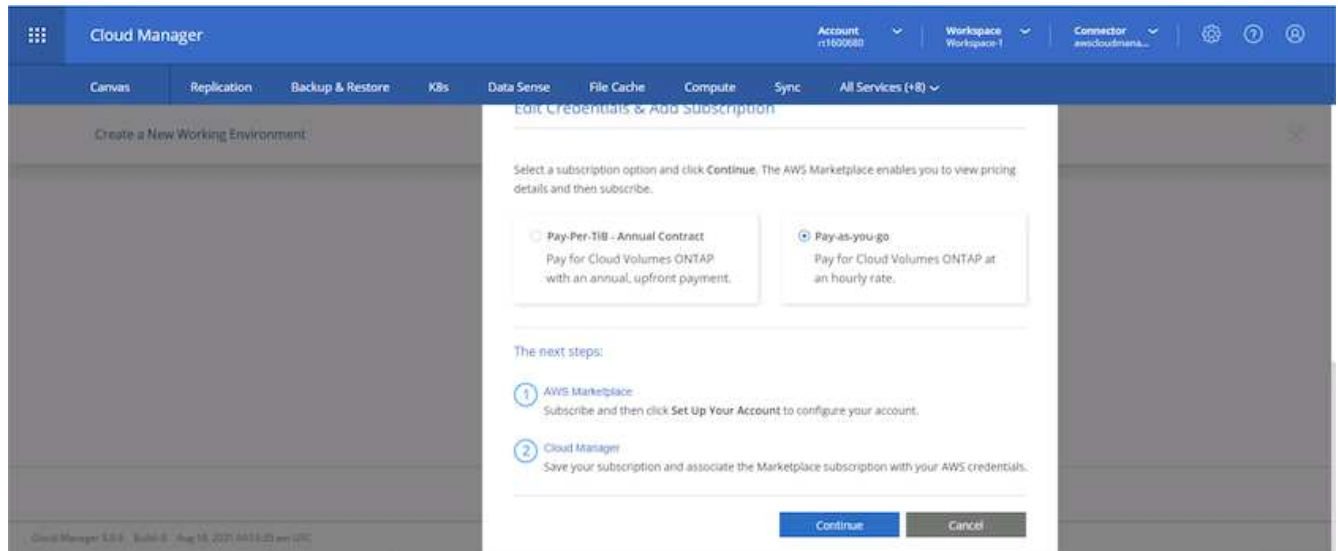
2. If no subscription has been assigned and you wish to purchase with PAYGO, choose Edit Credentials.



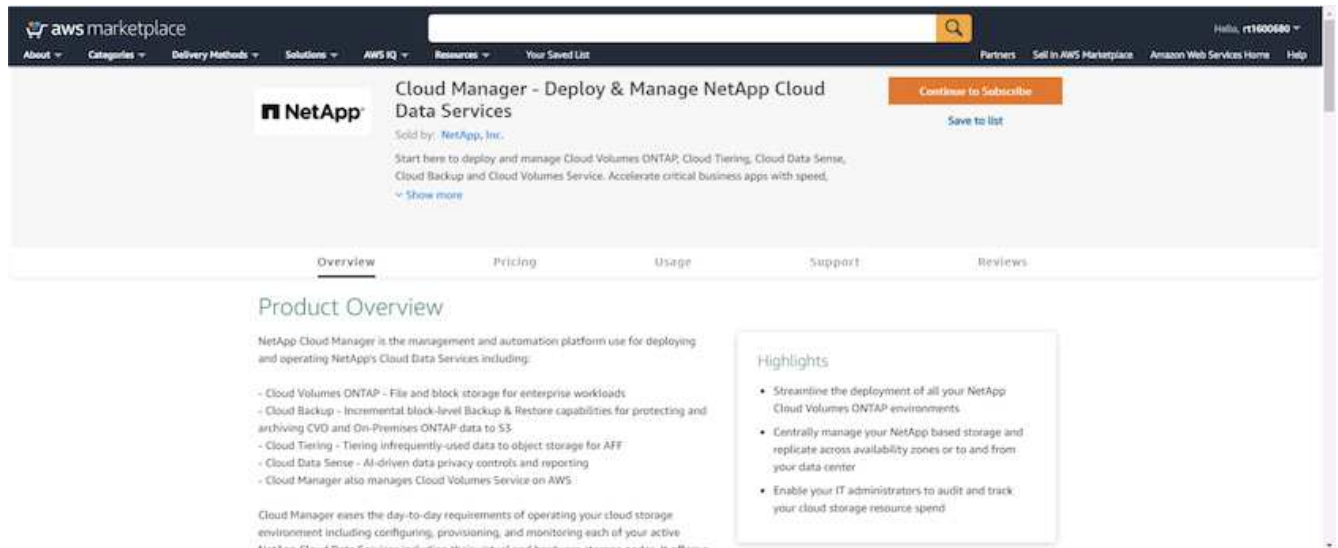
3. Choose Add Subscription.




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



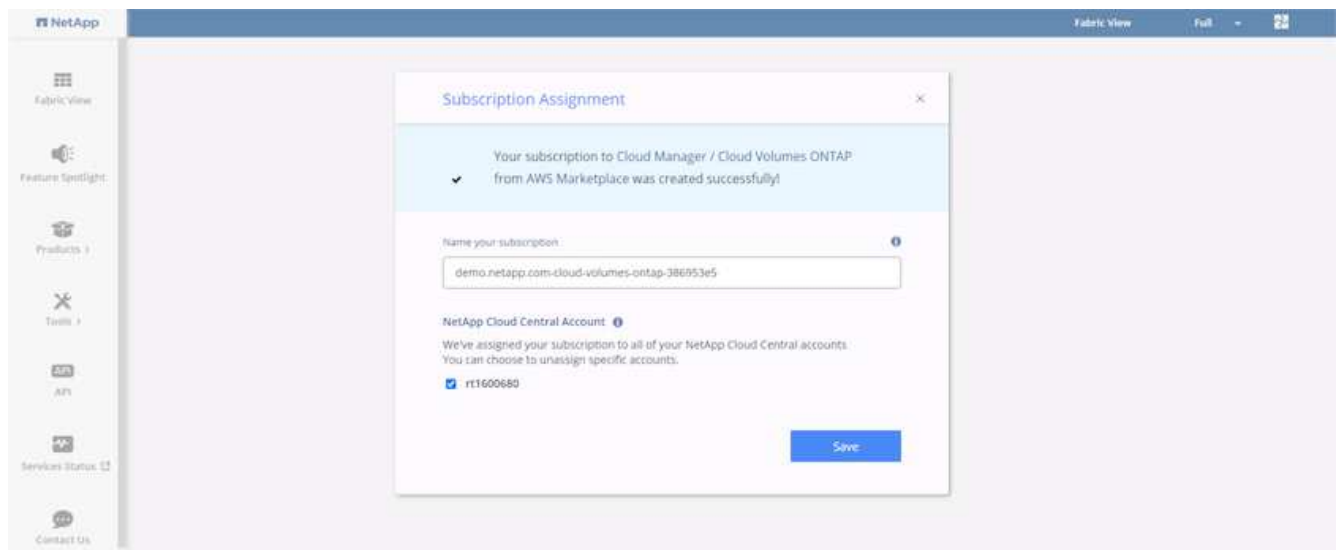
5. You are redirected to AWS; choose Continue to Subscribe.



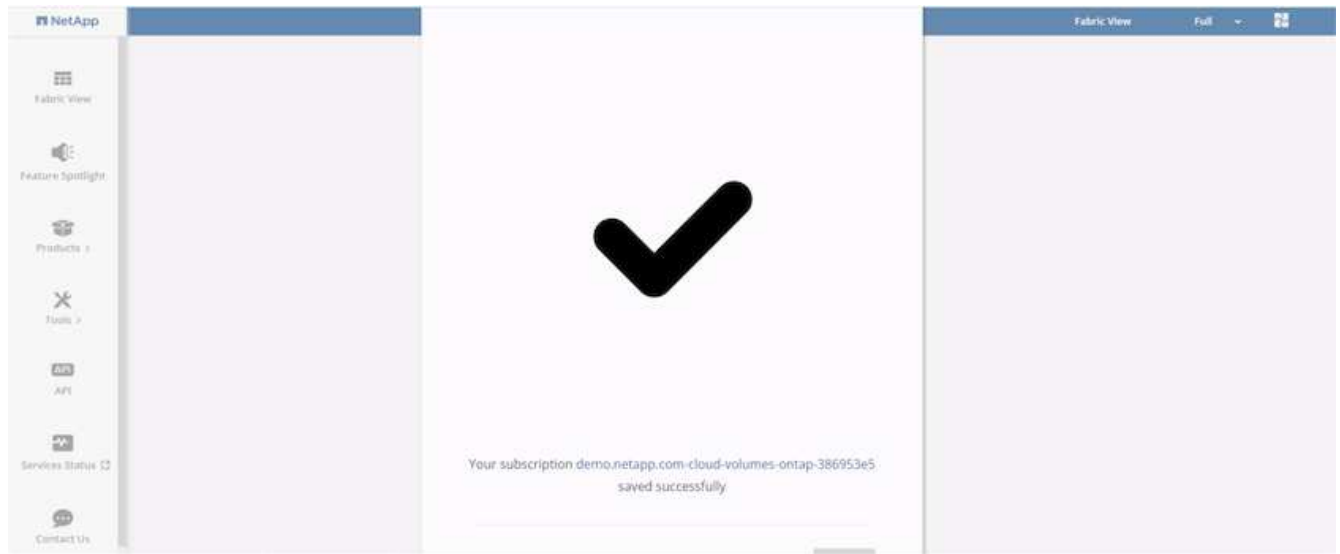
6. Subscribe and you are redirected back to NetApp Cloud Central. If you have already subscribed and don't get redirected, choose the "Click here" link.




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



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




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




10. Enter the working environment details such as:

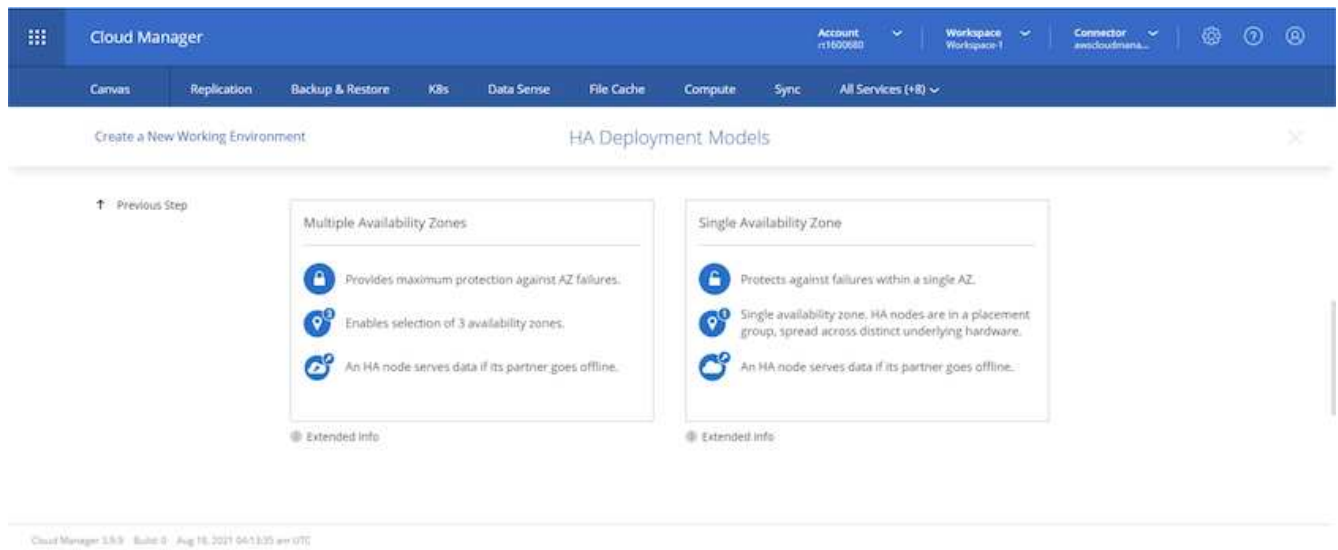
- a. Cluster name
- b. Cluster password
- c. AWS tags (Optional)




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




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



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



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





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

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

The screenshot shows the 'Floating IPs' configuration step in the Cloud Manager console. The page title is 'Floating IPs' and it is part of a 'Create a New Working Environment' wizard. The instructions state: '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.'

Four input fields are provided for floating 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
- Floating IP address for SVM management (Optional): Enter Floating IP Address

A 'Continue' button is located at the bottom of the configuration area.

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' and it is part of a 'Create a New Working Environment' wizard. The instructions state: 'Select the route tables that should include routes to the floating IP addresses. This enables client access to the Cloud Volumes ONTAP HA pair. If you leave a route table unselected, clients that are associated with the route table cannot access the HA pair.'

Additional information is provided:


Additional information ⓘ

<input checked="" type="checkbox"/>	Name	Main	ID	Associate with Subnet	Tags
<input checked="" type="checkbox"/>	private_rt_rt1600680	No	rtb-08b4cb88f5c826a5	3 Subnets	1 Tags
<input checked="" type="checkbox"/>	public_rt_rt1600680	Yes	rtb-0e46720d0da10c593	1 Subnets	1 Tags

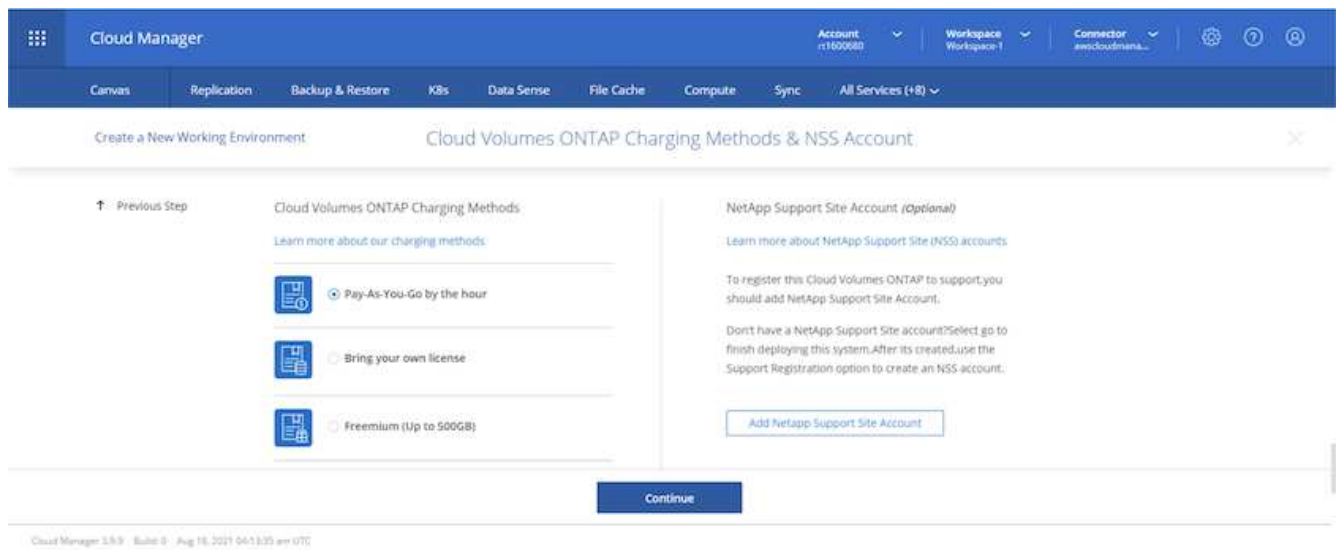
2 Route Tables | The main route table is the default for the VPC

A 'Continue' button is located at the bottom of the configuration area.

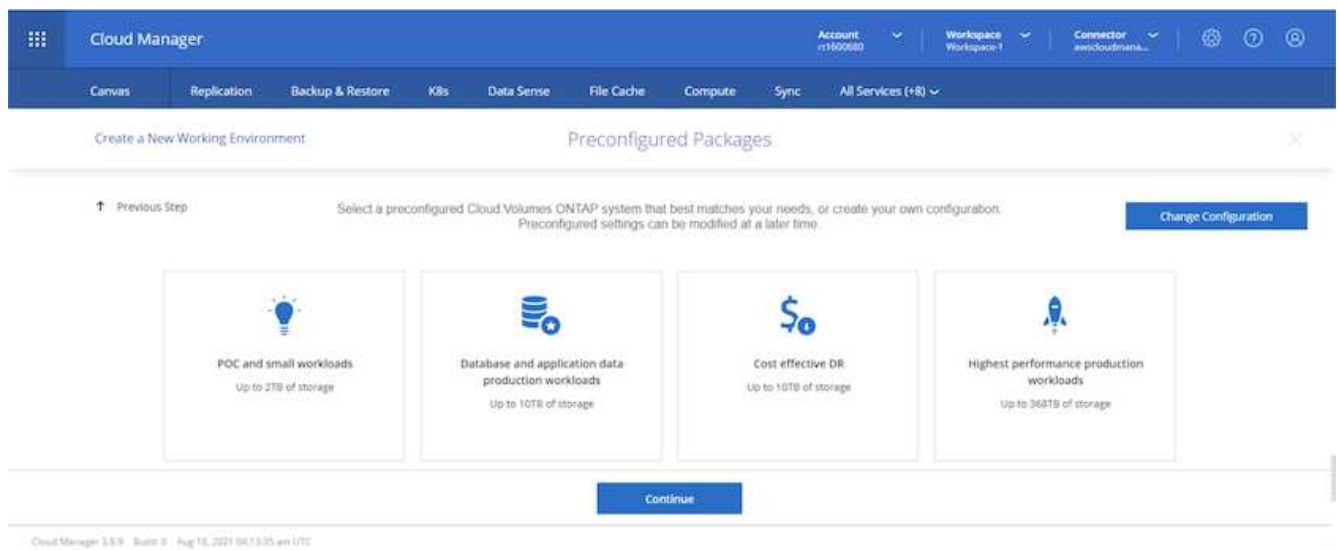
3. Choose whether to enable AWS managed encryption or AWS KMS to encrypt the ONTAP root, boot, and data disks.



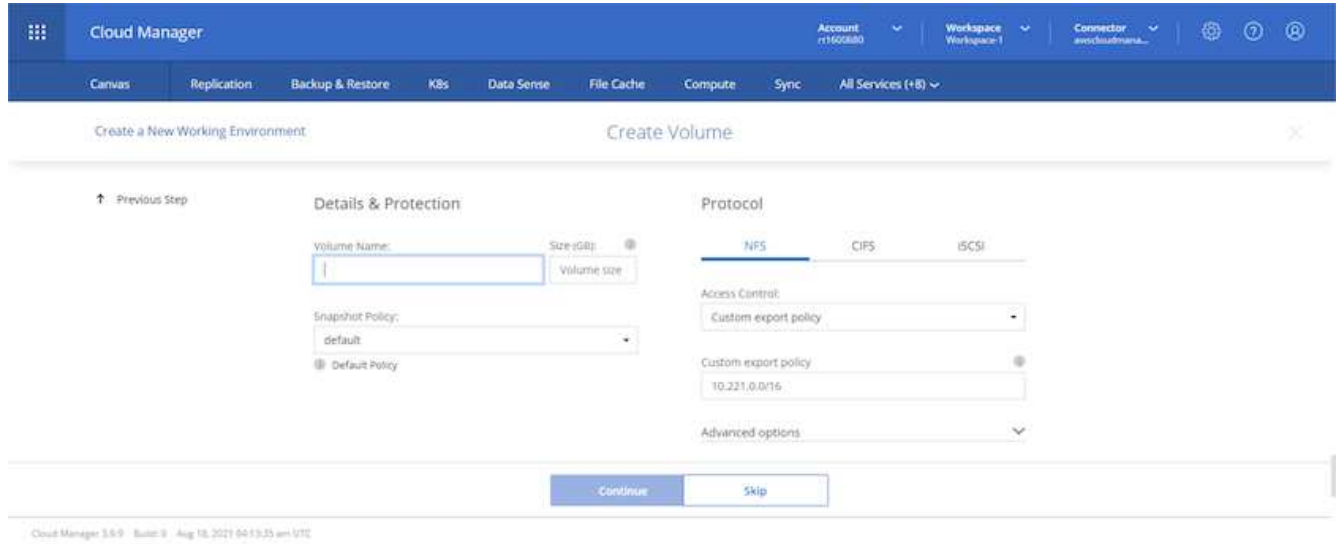
4. Choose your licensing model. If you don't know which to choose, contact your NetApp representative.



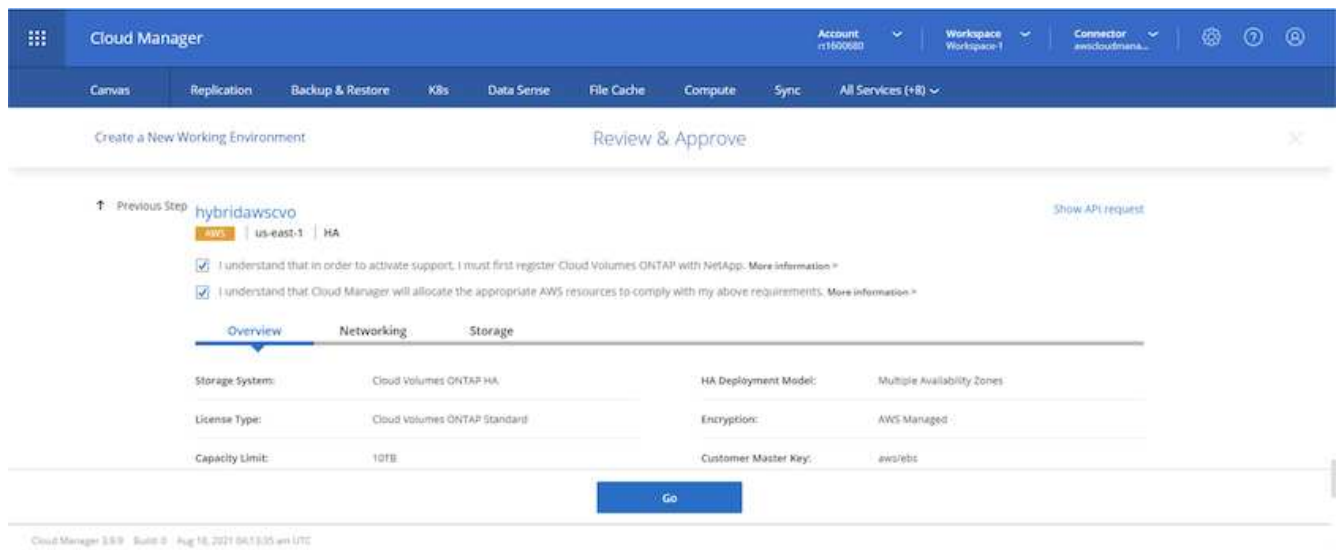
5. Select which configuration best suits your use case. This is related to the sizing considerations covered in the prerequisites page.



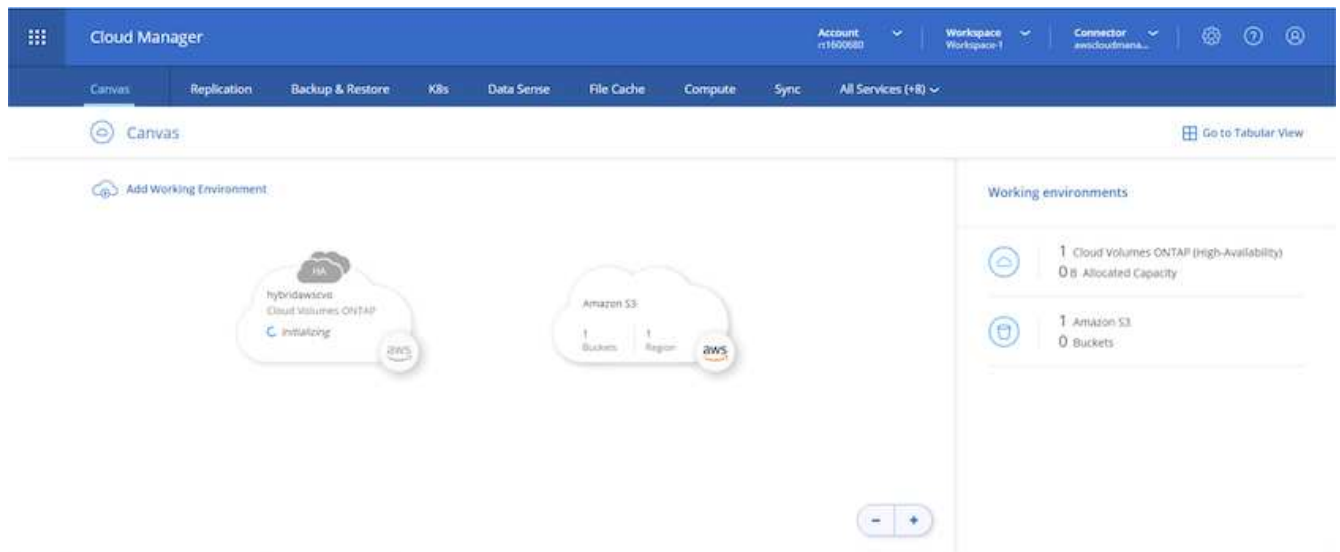
6. Optionally, create a volume. This is not required, because the next steps use SnapMirror, which creates the volumes for us.



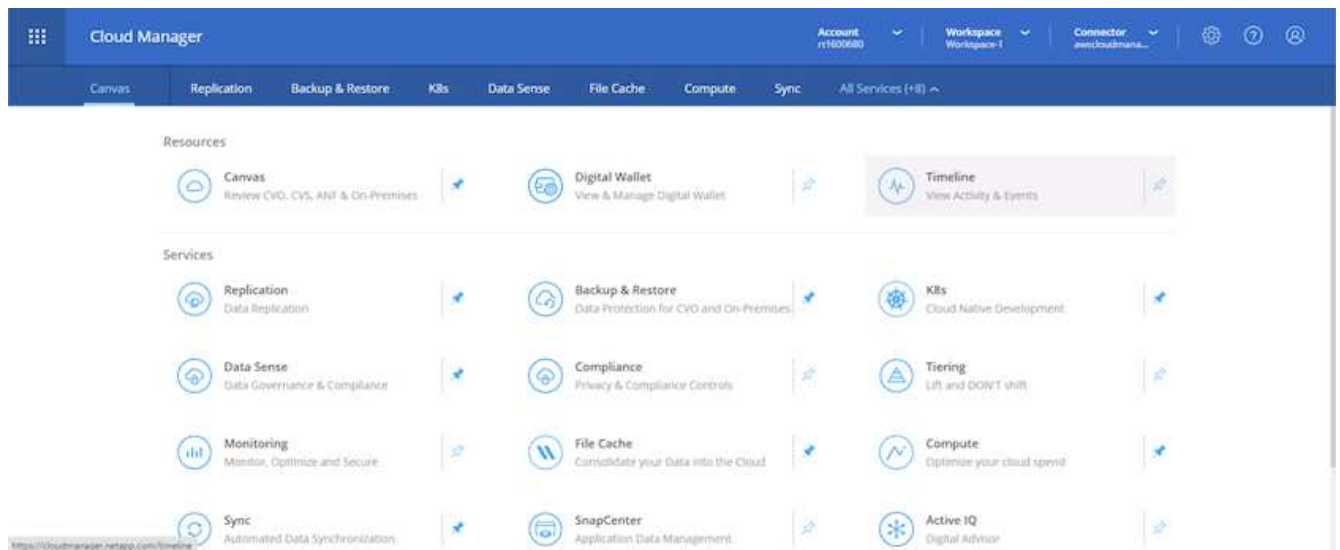
7. Review the selections made and tick the boxes to verify that you understand that Cloud Manager deploys resources into your AWS environment. When ready, click Go.




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.



11. After deployment is complete, the CVO cluster appears on the Canvas, which the current capacity. The ONTAP cluster in its current state is fully configured to allow a true, out-of-the-box experience.




Configure SnapMirror from on-premises to cloud

Now that you have a source ONTAP system and a destination ONTAP system deployed, you can replicate volumes containing database data into the cloud.

For a guide on compatible ONTAP versions for SnapMirror, see the [SnapMirror Compatibility Matrix](#).

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




Select Enable.


SERVICES

	Replication ■ Off	<input type="button" value="Enable"/>	
---	-----------------------------	---------------------------------------	---


Or Options.



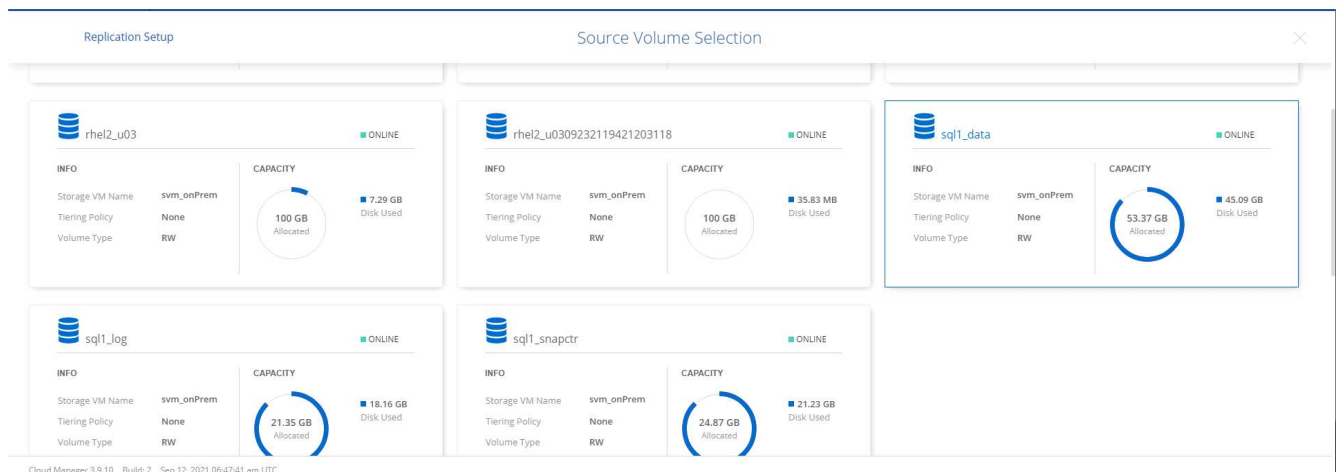
Replicate.




2. If you did not drag and drop, choose the destination cluster to replicate to.



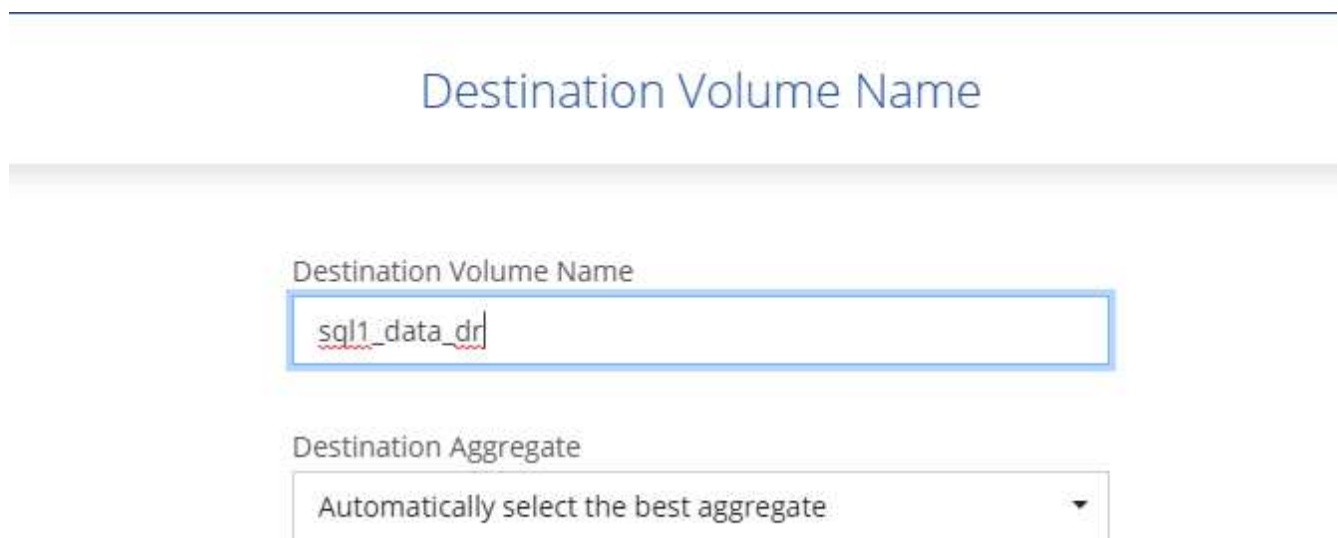
3. Choose the volume that you'd like to replicate. We replicated the data and all log volumes.



4. Choose the destination disk type and tiering policy. For disaster recovery, we recommend an SSD as the disk type and to maintain data tiering. Data tiering tiers the mirrored data into low-cost object storage and saves you money on local disks. When you break the relationship or clone the volume, the data uses the fast, local storage.



5. Select the destination volume name: we chose [source_volume_name]_dr.



6. Select the maximum transfer rate for the replication. This enables you to save bandwidth if you have a low bandwidth connection to the cloud such as a VPN.

Max Transfer Rate


You should limit the transfer rate. An unlimited rate might negatively impact the performance of other applications and it might impact your Internet performance.

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

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


Replication Policy

Default Policies Additional Policies

 Mirror

Typically used for disaster recovery


[More info](#)

 Mirror and Backup (1 month retention)


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

[More info](#)


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



- 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.



- Continue this process for data volumes and log volumes.
- 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.



- 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](#).

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.