Hybrid Cloud Database Solutions with SnapCenter

NetApp Solutions

NetApp
April 26, 2024

Table of Contents

Hybrid Cloud Database Solutions with SnapCenter

TR-4908: Hybrid Cloud Database Solutions with SnapCenter Overview

Alan Cao, Felix Melligan, NetApp

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

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

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

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

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

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

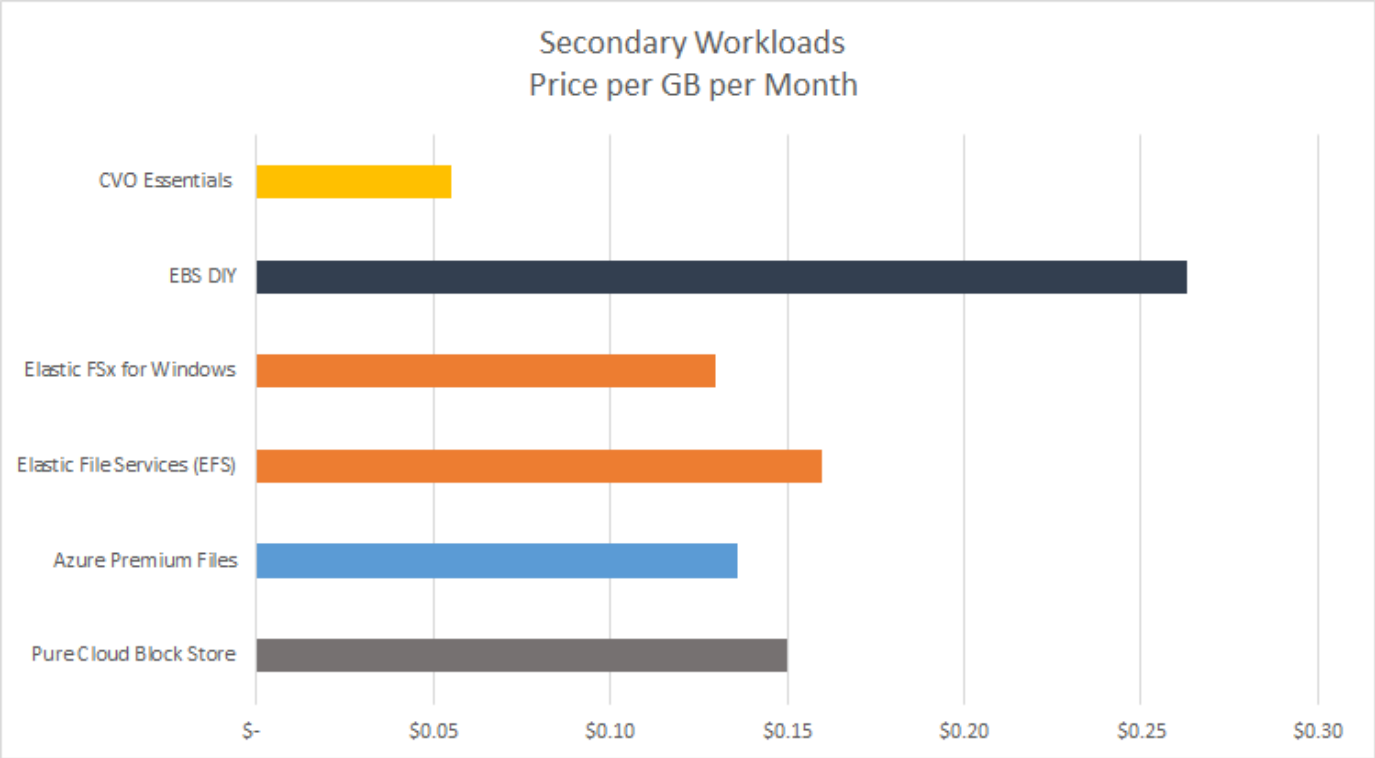

This solution demonstrates that, by using the SnapCenter GUI-based software tool and NetApp SnapMirror technology, hybrid cloud database operations can be easily setup, implemented, and operated.

The following videos demonstrate SnapCenter in action:

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

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

To test drive the solution and use cases for yourself, a NetApp Lab-on-Demand SL10680 can be requested at following xref:./databases/ TL_AWS_004 HCoD: AWS - NW,SnapCenter(OnPrem).

Solution Architecture

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

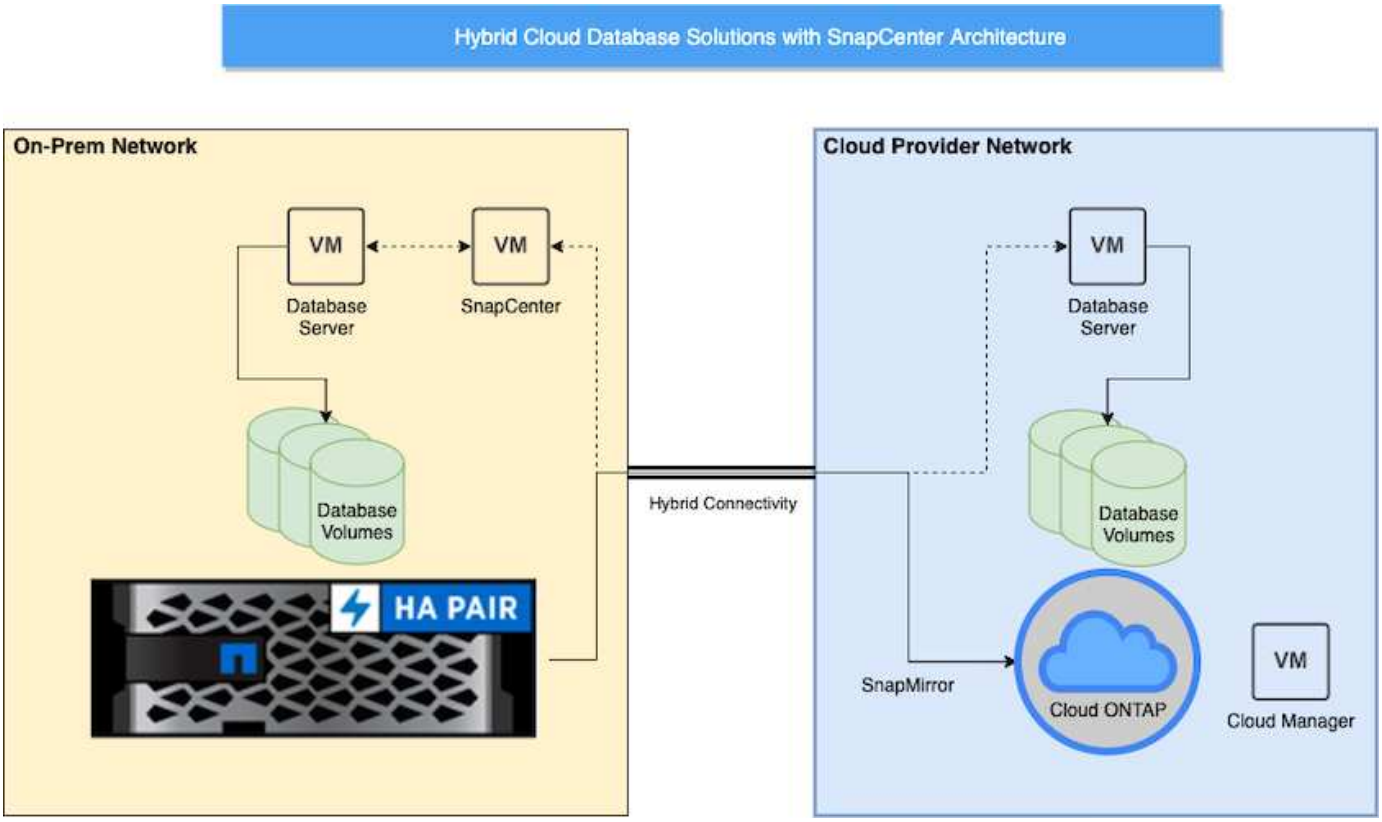

In normal business operations, synchronized database volumes in the cloud can be cloned and mounted to dev/test database instances for applications development or testing. In the event of a failure, the synchronized database volumes in the cloud can then be activated for disaster recovery.

SnapCenter Requirements

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

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

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

Requirements

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

Prerequisites configuration

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

On premises

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

Public cloud

- A NetApp Cloud Central login
- Network access from a web browser to several endpoints

- A network location for a connector
- Cloud provider permissions
- Networking for individual services

Important considerations:

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

The following links provide further details:

On Premises

Public Cloud

Prerequisites on-premises

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

SnapCenter installation and configuration

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

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

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

On-premises database server storage configuration

Storage performance plays an important role in the overall performance of databases and applications. A well-designed storage layout can not only improve DB performance but also make it easy to manage database backup and recovery. Several factors should be considered when defining your storage layout, including the

size of the database, the rate of expected data change for the database, and the frequency with which you perform backups.

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

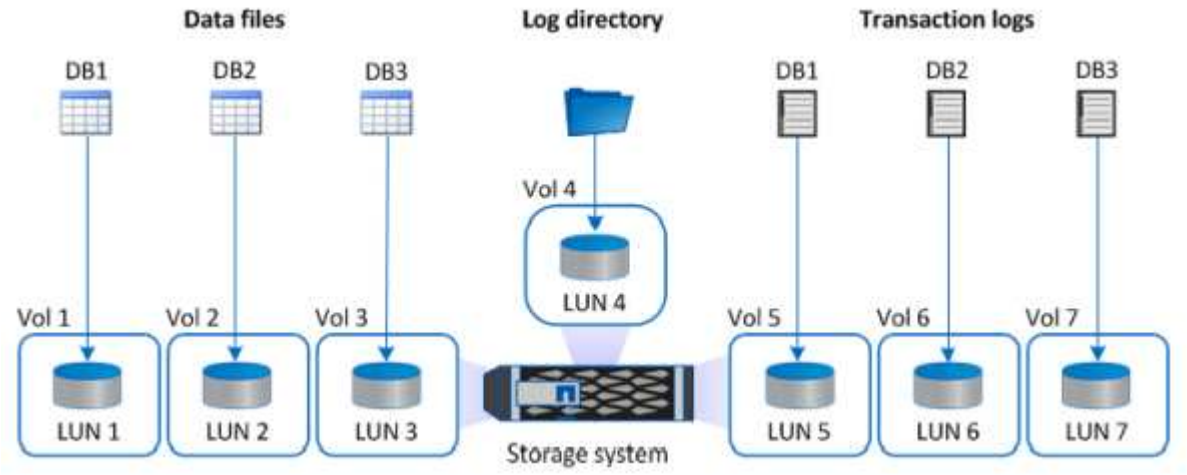

The following figure shows the NetApp recommended storage layout for small or medium SQL Server database on LUNs.

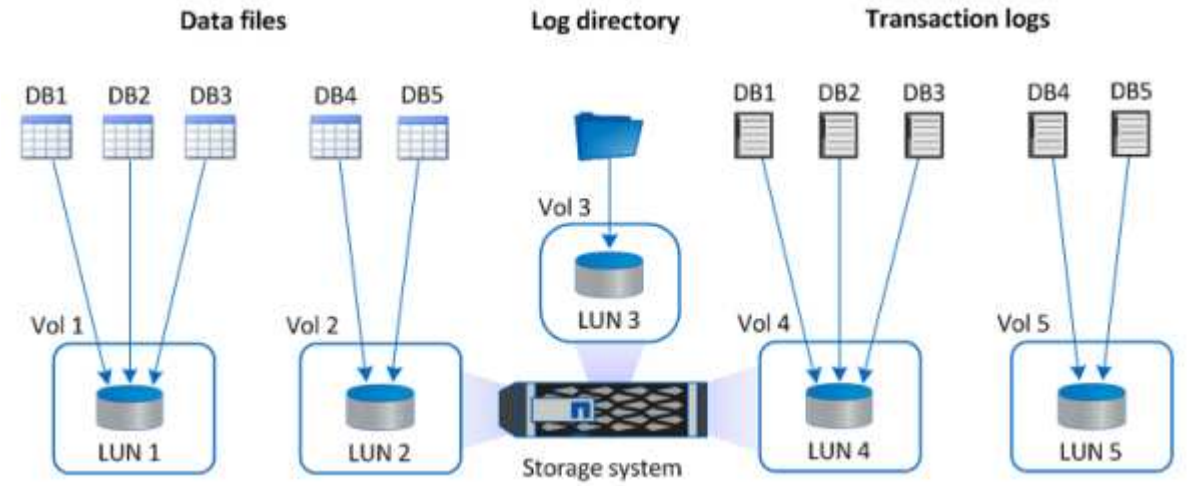

The Log directory is dedicated to SnapCenter to perform transaction log rollup for database recovery. For an extra large database, multiple LUNs can be allocated to a volume for better performance.

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

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

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

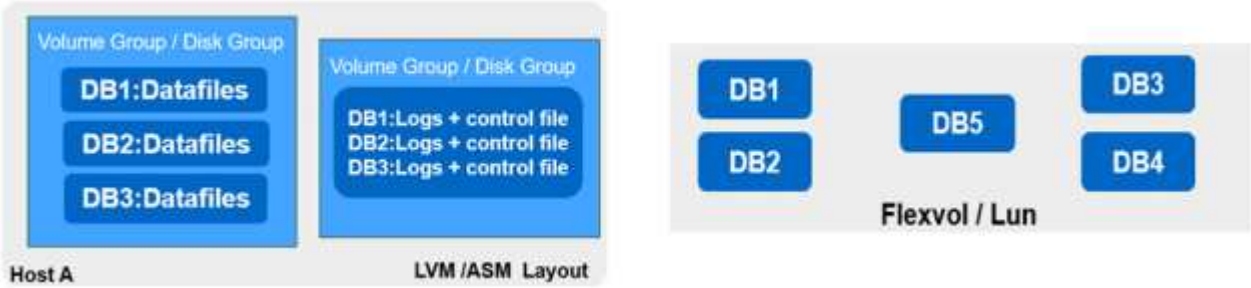

To facilitate the relocation of Oracle databases, the Oracle binary should be installed on a separate LUN that is included in the regular backup policy. This ensures that in the case of database relocation to a new server host, the Oracle stack can be started for recovery without any potential issues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

Licensing requirements

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

SnapCenter standard capacity-based licenses

Networking and security

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

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

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

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

- Security group rules for CVO - AWS
- Security group rules for CVO - Azure
- Firewall rules for CVO - GCP

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

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

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

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

Prerequisites for the public cloud

Before we install the Cloud Manager connector and Cloud Volumes ONTAP and configure SnapMirror, we must perform some preparation for our cloud environment. This page describes the work that needs to be done as well as the considerations when deploying Cloud Volumes ONTAP.

Cloud Manager and Cloud Volumes ONTAP deployment prerequisites checklist

- ☐ A NetApp Cloud Central login
- ☐ Network access from a web browser to several endpoints
- ☐ A network location for a Connector
- ☐ Cloud provider permissions
- ☐ Networking for individual services

For more information about what you need to get started, visit our cloud documentation.

Considerations

1. What is a Cloud Manager connector?

In most cases, a Cloud Central account admin must deploy a connector in your cloud or on-premises network. The connector enables Cloud Manager to manage resources and processes within your public cloud environment.

For more information about Connectors, visit our cloud documentation.

2. Cloud Volumes ONTAP sizing and architecture

When deploying Cloud Volumes ONTAP, you are given the choice of either a predefined package or the creation of your own configuration. Although many of these values can be changed later on nondisruptively, there are some key decisions that need to be made before deployment based on the workloads to be deployed in the cloud.

Each cloud provider has different options for deployment and almost every workload has its own unique properties. NetApp has a CVO sizing tool that can help size deployments correctly based on capacity and performance, but it has been built around some basic concepts which are worth considering:

- Capacity required
- Network capability of the cloud virtual machine
- Performance characteristics of cloud storage

The key is to plan for a configuration that not only satisfies the current capacity and performance requirements, but also looks at future growth. This is generally known as capacity headroom and performance headroom.

If you would like further information, read the documentation about planning correctly for AWS, Azure, and GCP.

3. Single node or high availability?

In all clouds, there is the option to deploy CVO in either a single node or in a clustered high availability pair with two nodes. Depending on the use case, you might wish to deploy a single node to save costs or an HA pair to provide further availability and redundancy.

For a DR use case or spinning up temporary storage for development and testing, single nodes are common since the impact of a sudden zonal or infrastructure outage is lower. However, for any production use case, when the data is in only a single location, or when the dataset must have more redundancy and availability, high availability is recommended.

For further information about the architecture of each cloud's version of high availability, visit the documentation for AWS, Azure and GCP.

Getting started overview

This section provides a summary of the tasks that must be completed to meet the prerequisite requirements as outlined in previous section. The following section provide a high level tasks list for both on-premises and public cloud operations. The detailed processes and procedures can be accessed by clicking on the relevant links.

On-premises

- Setup database admin user in SnapCenter
- SnapCenter plugin installation prerequisites
- SnapCenter host plugin installation
- DB resource discovery
- Setup storage cluster peering and DB volume replication
- Add CVO database storage SVM to SnapCenter
- Setup database backup policy in SnapCenter
- Implement backup policy to protect database
- Validate backup

AWS public cloud

- Pre-flight check
- Steps to deploy Cloud Manager and Cloud Volumes ONTAP in AWS
- Deploy EC2 compute instance for database workload

Click the following links for details:

On Premises, Public Cloud - AWS

Getting started on premises

The NetApp SnapCenter tool uses role based access control (RBAC) to manage user resources access and permission grants, and SnapCenter installation creates prepopulated roles. You can also create custom roles based on your needs or applications.

On Premises

1. Setup database admin user in SnapCenter

It makes sense to have a dedicated admin user ID for each database platform supported by SnapCenter for database backup, restoration, and/or disaster recovery. You can also use a single ID to manage all databases. In our test cases and demonstration, we created a dedicated admin user for both Oracle and SQL Server, respectively.

Certain SnapCenter resources can only be provisioned with the SnapCenterAdmin role. Resources can then be assigned to other user IDs for access.

In a pre-installed and configured on-premises SnapCenter environment, the following tasks might have already have been completed. If not, the following steps create a database admin user:

1. Add the admin user to Windows Active Directory.
2. Log into SnapCenter using an ID granted with the SnapCenterAdmin role.
3. Navigate to the Access tab under Settings and Users, and click Add to add a new user. The new user ID is linked to the admin user created in Windows Active Directory in step 1. . Assign the proper role to the user

as needed. Assign resources to the admin user as applicable.

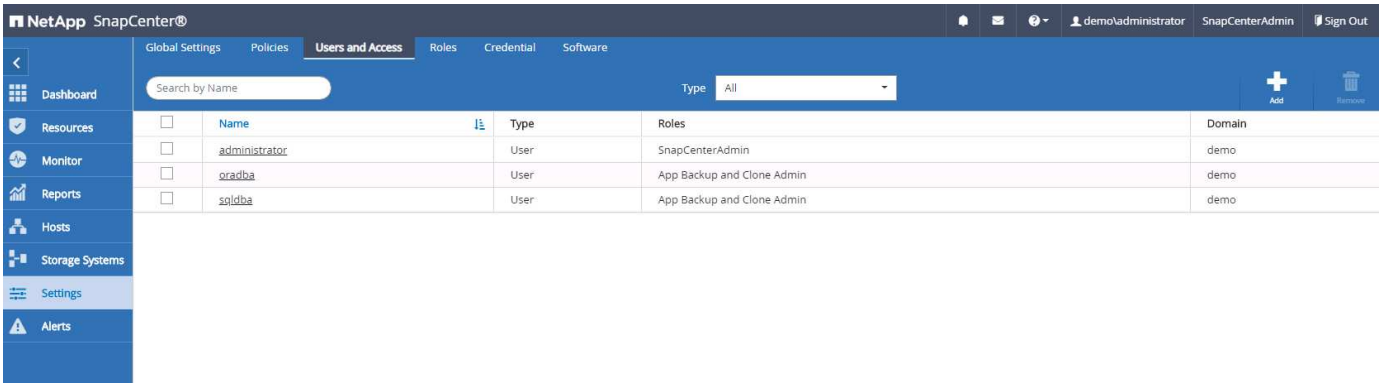

2. SnapCenter plugin installation prerequisites

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

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

Windows host for SQL server:

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

Unix host for Oracle:

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

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

3. SnapCenter host plugin installation

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

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

Configure SnapCenter VMware global settings

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

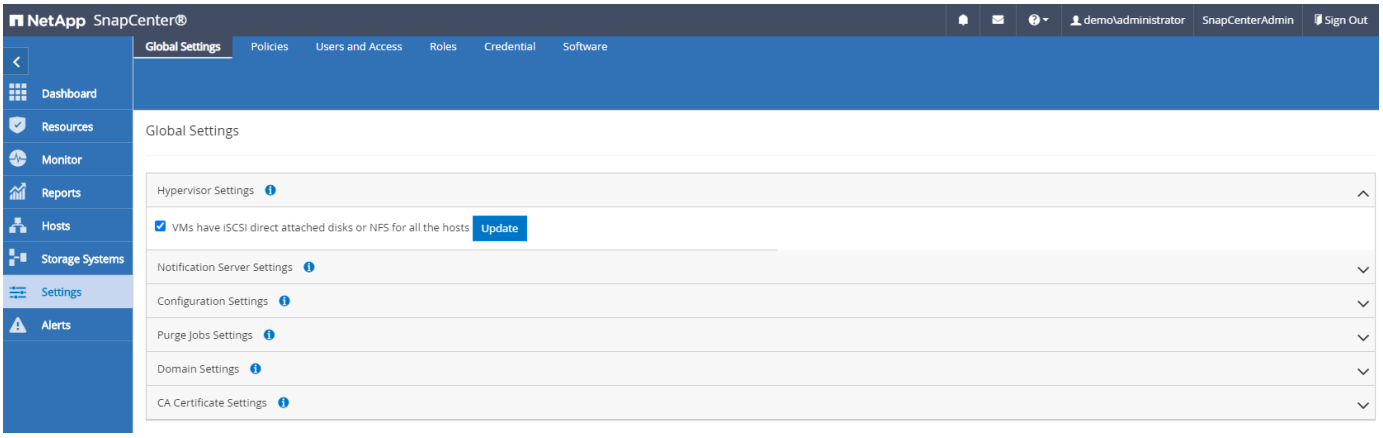

Add Windows host and installation of plugin on the host

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

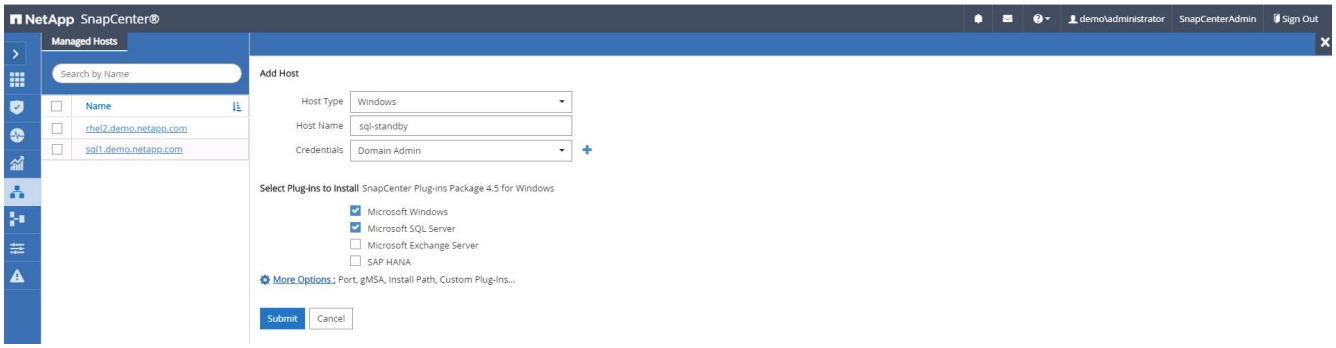

4. After the plugin is installed on a Windows host, its Overall Status is shown as "Configure log directory."

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

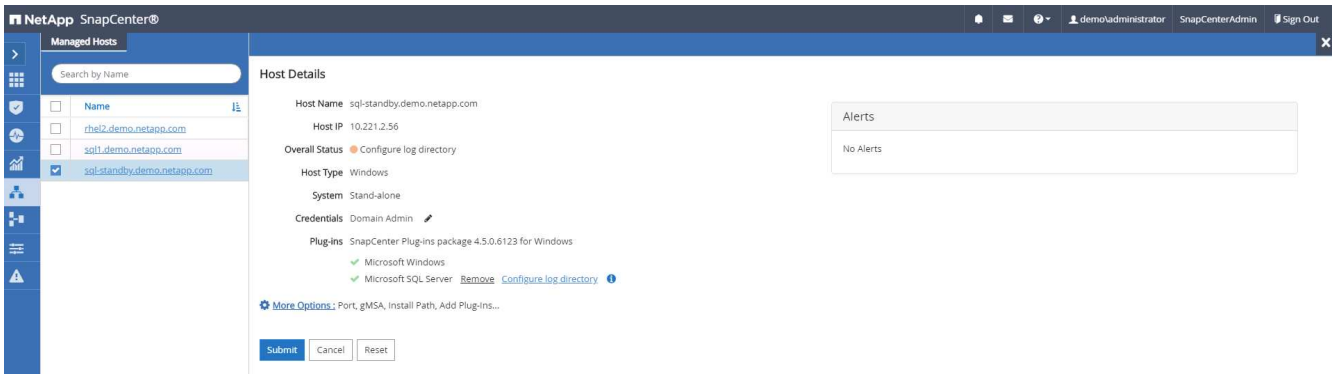

6. Click "Configure log directory" to open "Configure Plug-in for SQL Server."

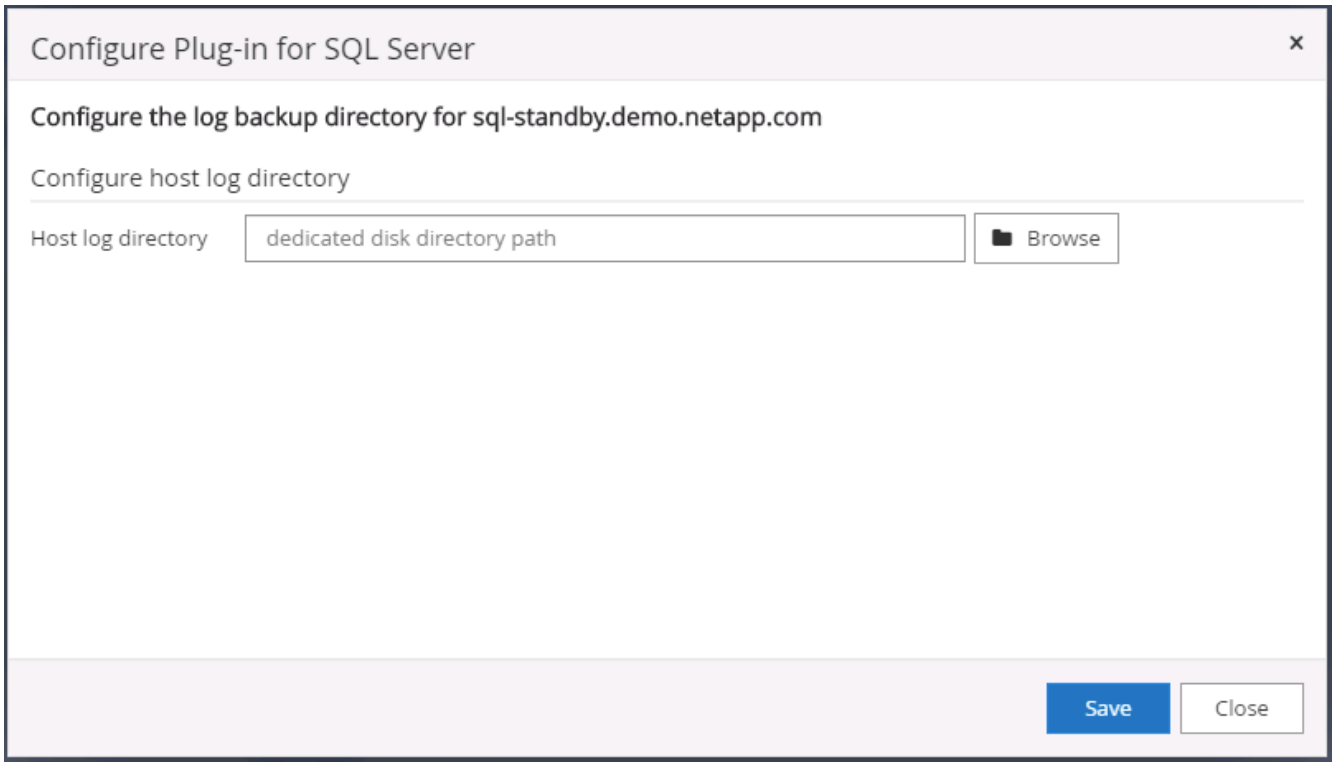

7. Click Browse to discover NetApp storage so that a log directory can be set; SnapCenter uses this log directory to roll up the SQL server transaction log files. Then click Save.

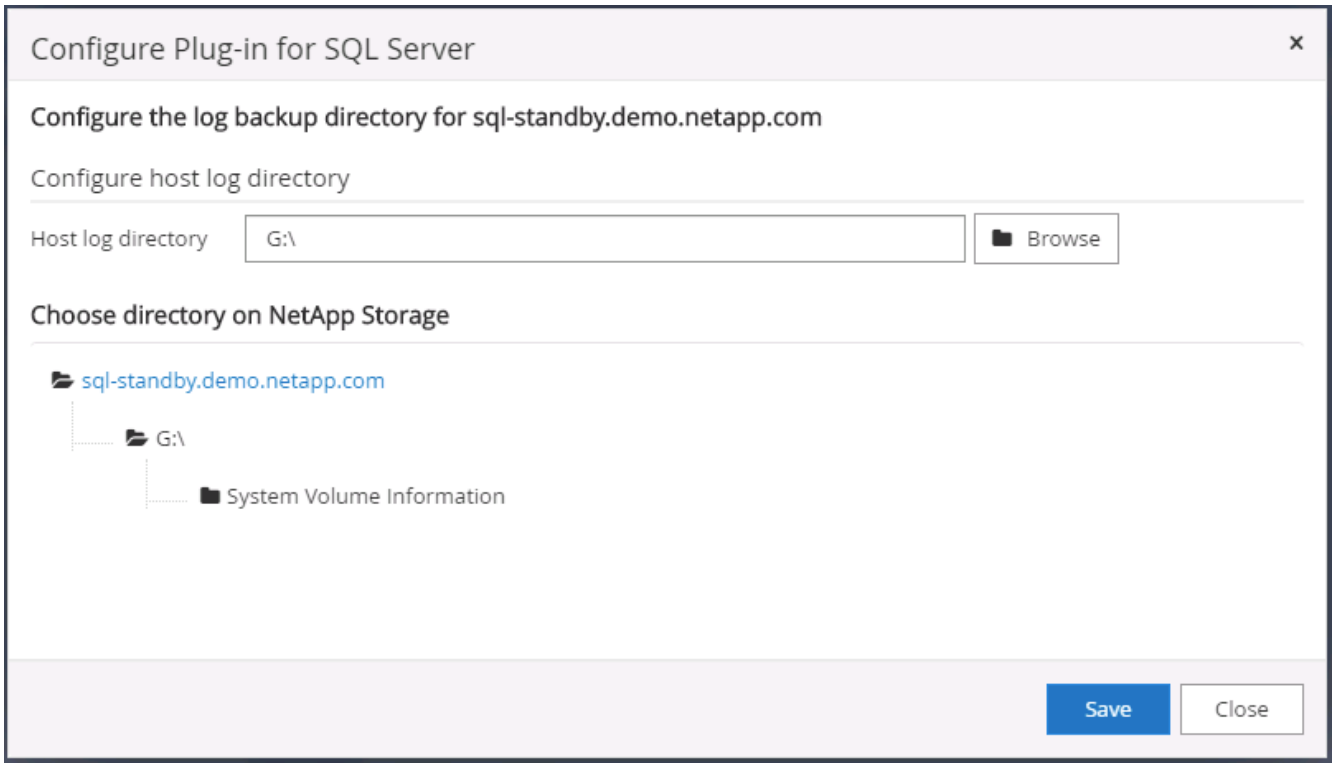

For NetApp storage provisioned to a DB host to be discovered, the storage (on-prem or CVO) must be added to SnapCenter, as illustrated in step 6 for CVO as an example.

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

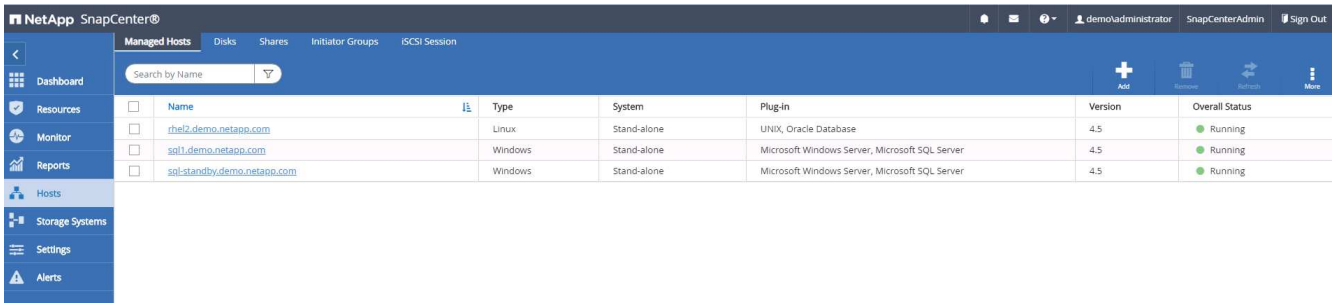

9. To assign the host to the database management user ID, navigate to the Access tab under Settings and Users, click the database management user ID (in our case the sqldba that the host needs to be assigned to), and click Save to complete host resource assignment.

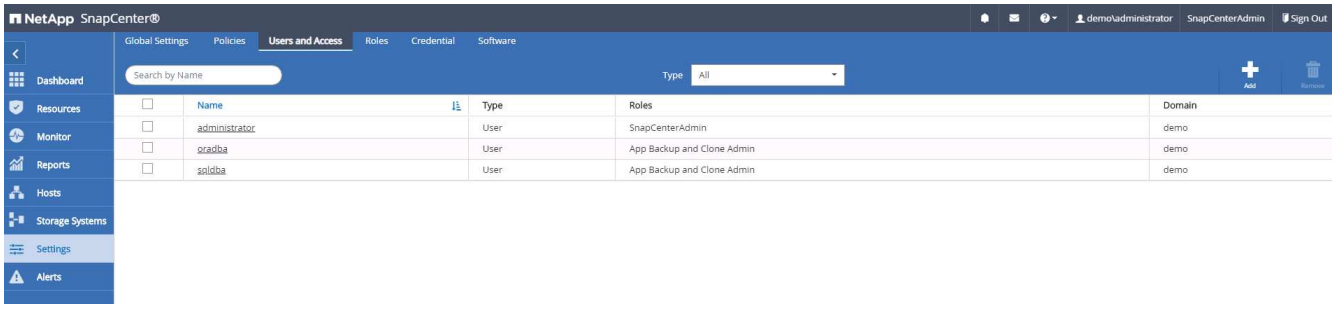

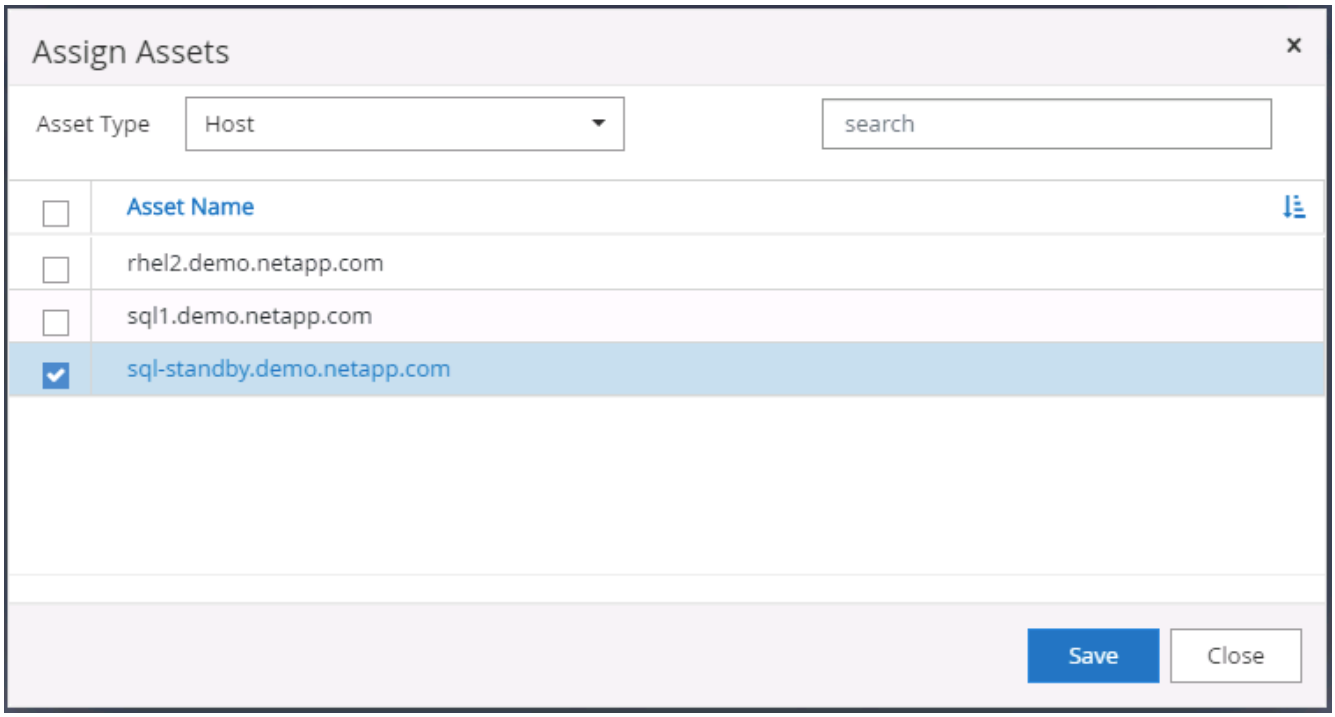

Add Unix host and installation of plugin on the host

1. Log into SnapCenter with a user ID with SnapCenterAdmin privileges.
2. Click the Hosts tab from left-hand menu, and click Add to open the Add Host workflow.
3. Choose Linux as the Host Type. The Host Name can be either the host name or an IP address. However, the host name must be resolved to correct host IP address from SnapCenter host. Choose host credentials created in step 2. The host credentials require sudo privileges. Check Oracle Database as the plug-in to be installed, which installs both Oracle and Linux host plugins.

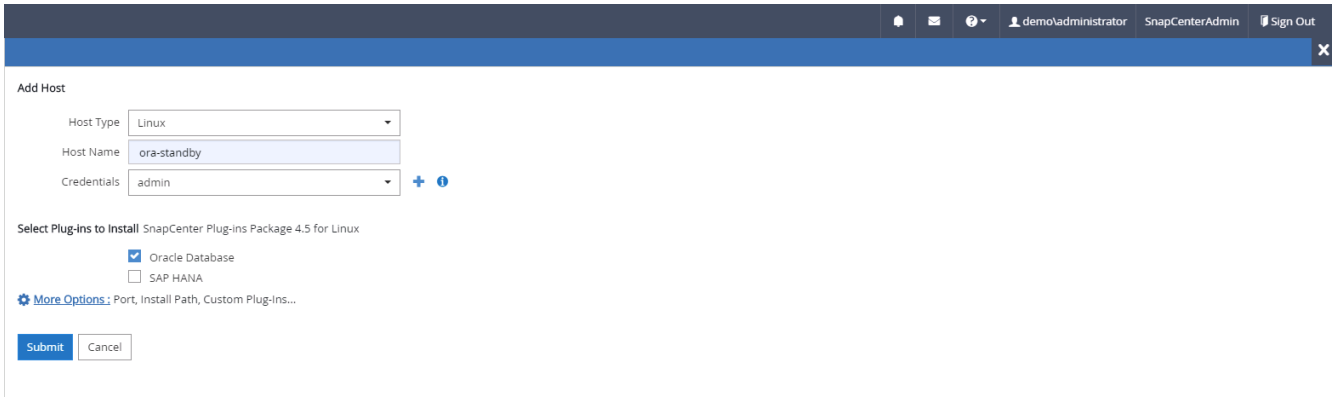

4. Click More Options and select "Skip preinstall checks." You are prompted to confirm the skipping of the preinstall check. Click Yes and then Save.

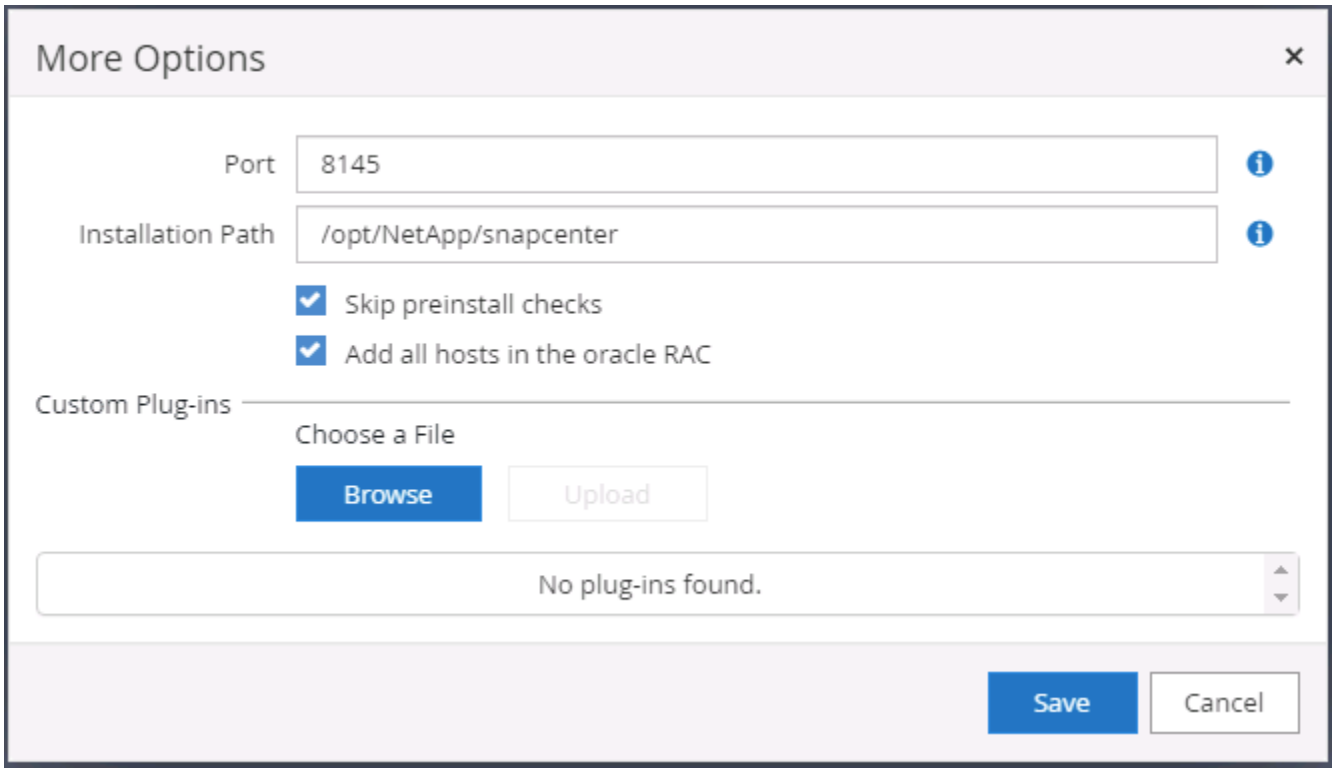

5. Click Submit to start the plugin installation. You are prompted to Confirm Fingerprint as shown below.

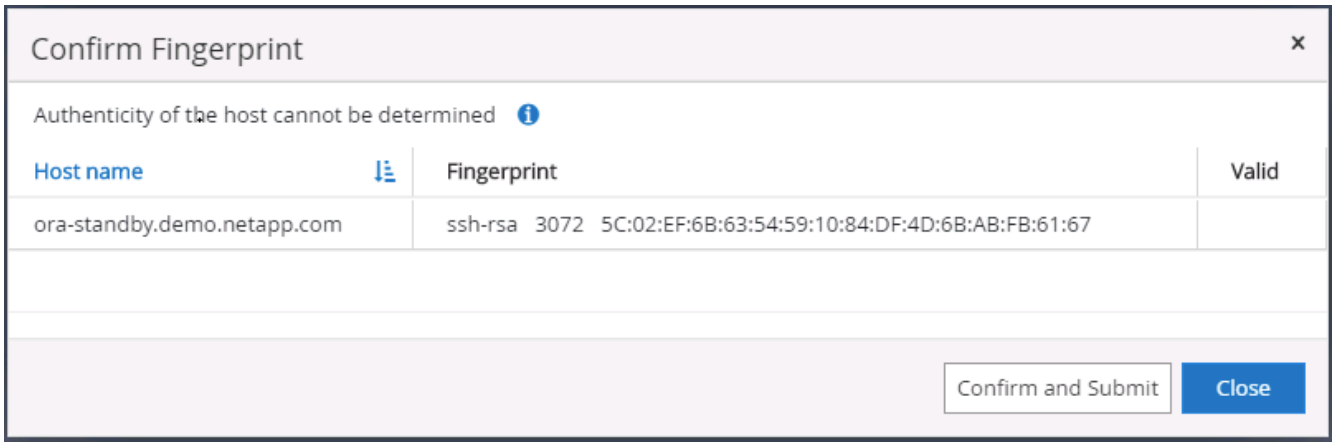

6. SnapCenter performs host validation and registration, and then the plugin is installed on the Linux host. The status is changed from Installing Plugin to Running.

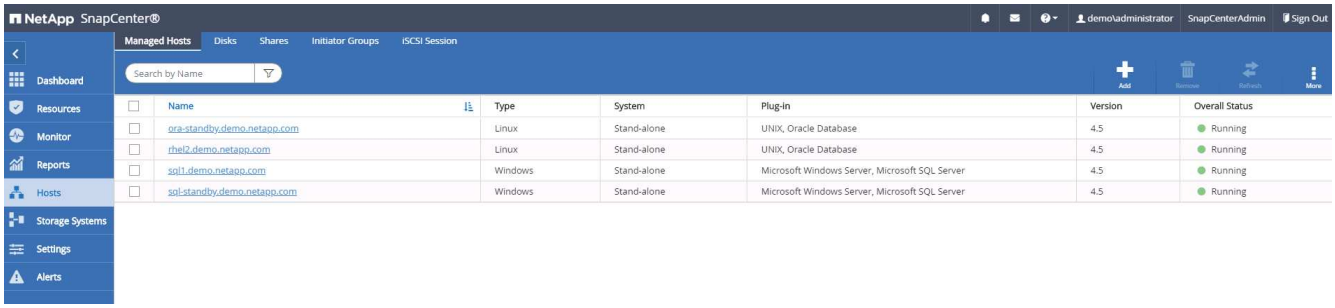

7. Assign the newly added host to the proper database management user ID (in our case, oradba).

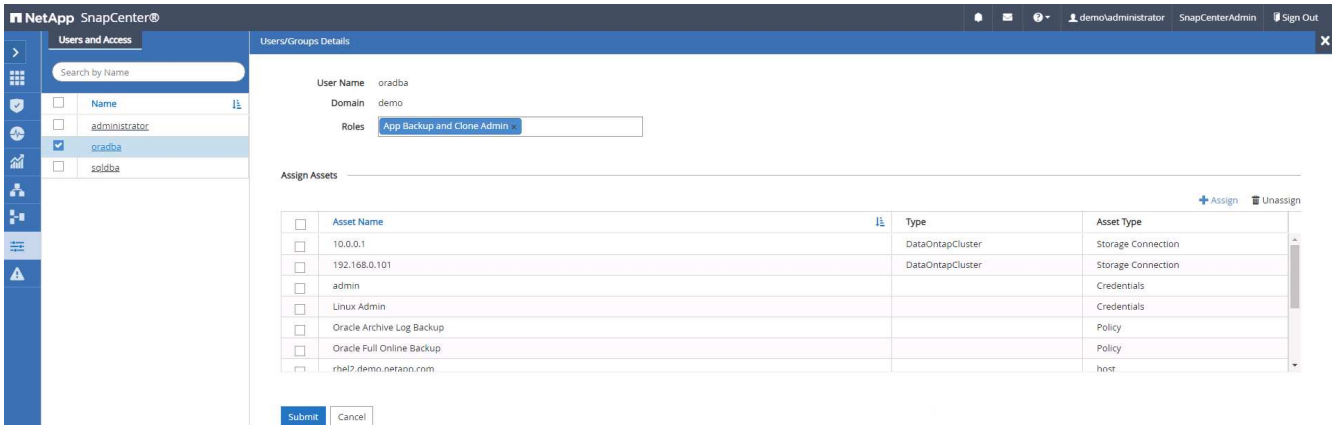

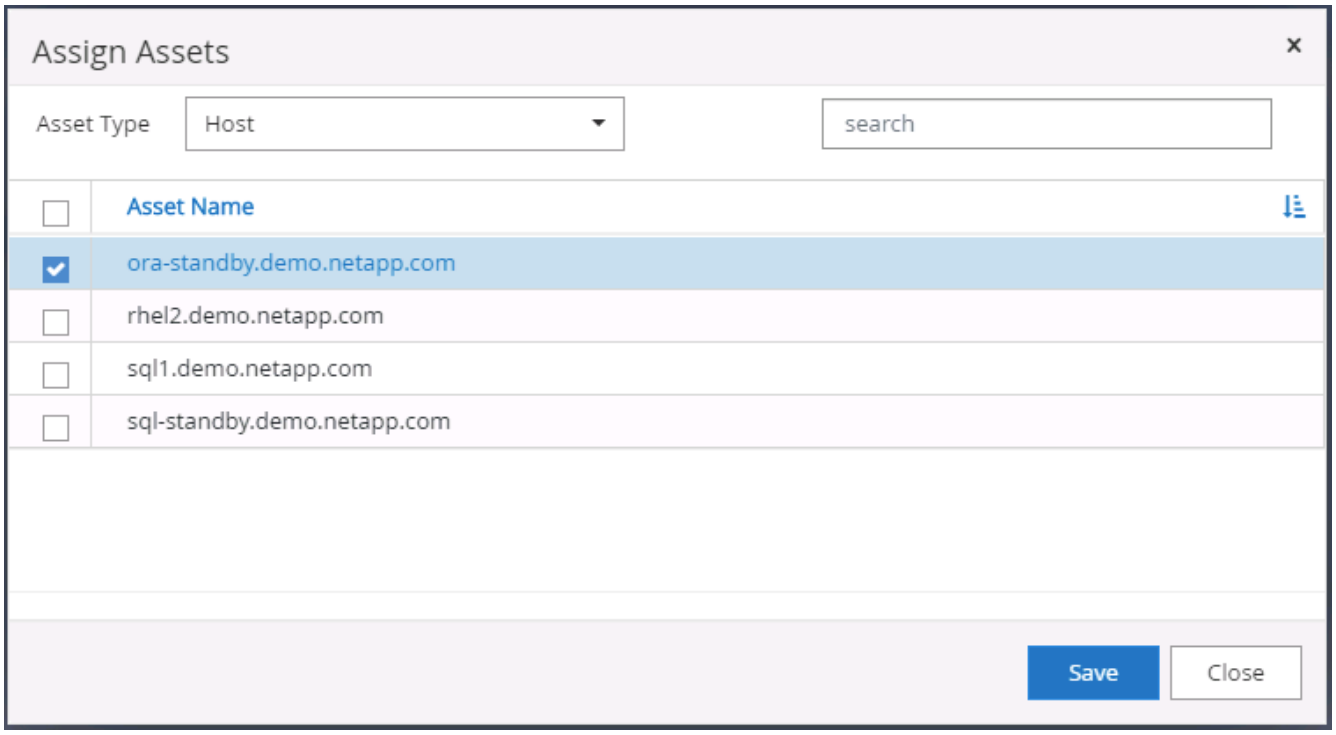

4. Database resource discovery

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.

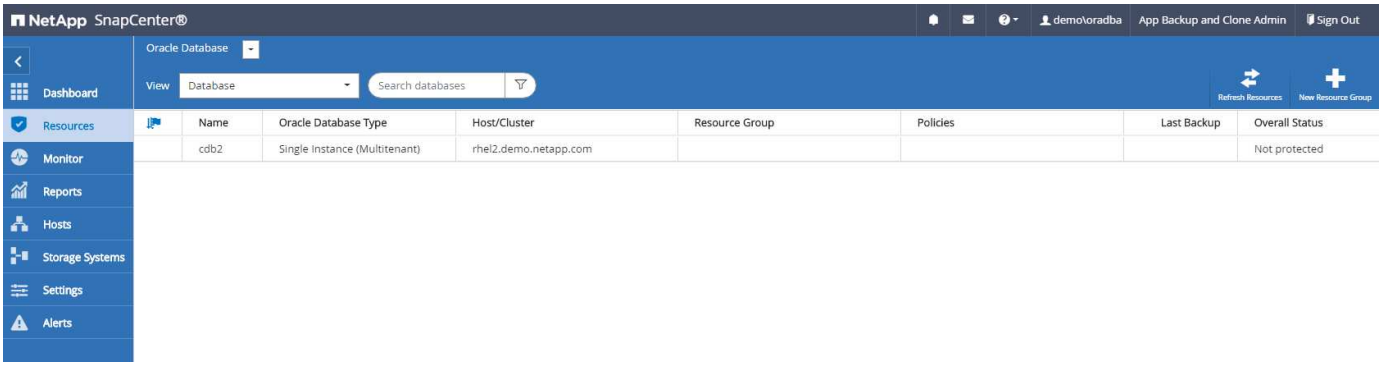

When the database is initially discovered, the Overall Status is shown as "Not protected." The previous screenshot shows an Oracle database not protected yet by a backup policy.

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.

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.

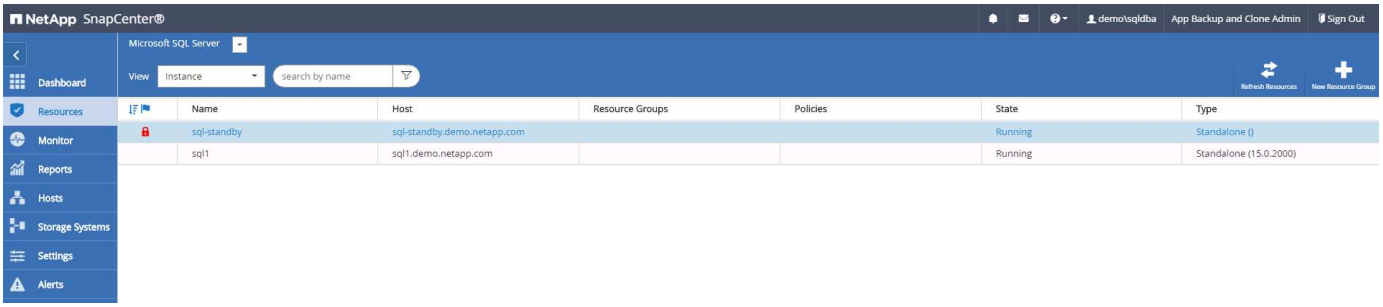

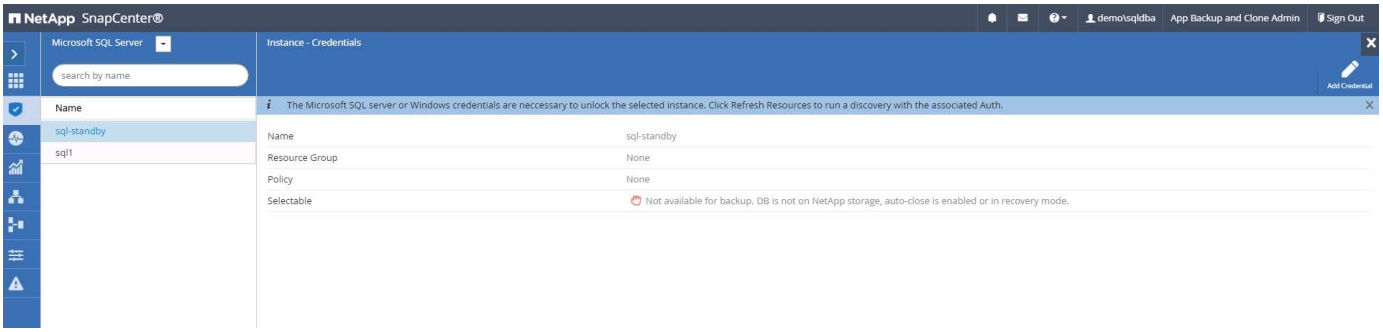

After the appropriate credentials are configured either at the Windows level or the database level, the red lock disappears and SQL Server Type information is gathered and reviewed.

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 Manger. A default CVO deployment has inter-cluster LIFs configured automatically.

   On-premises cluster:

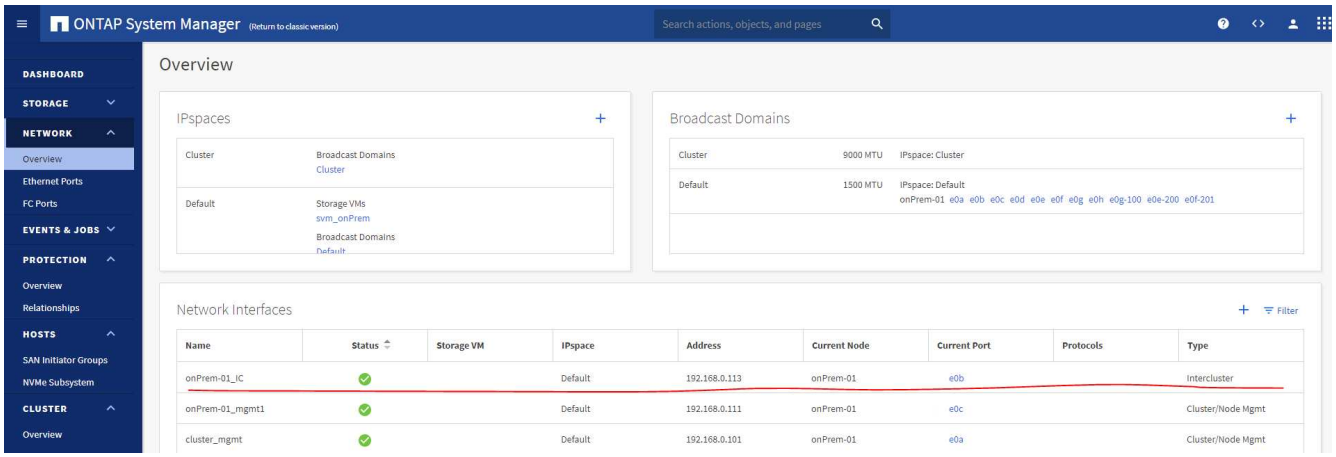

   Target CVO cluster:

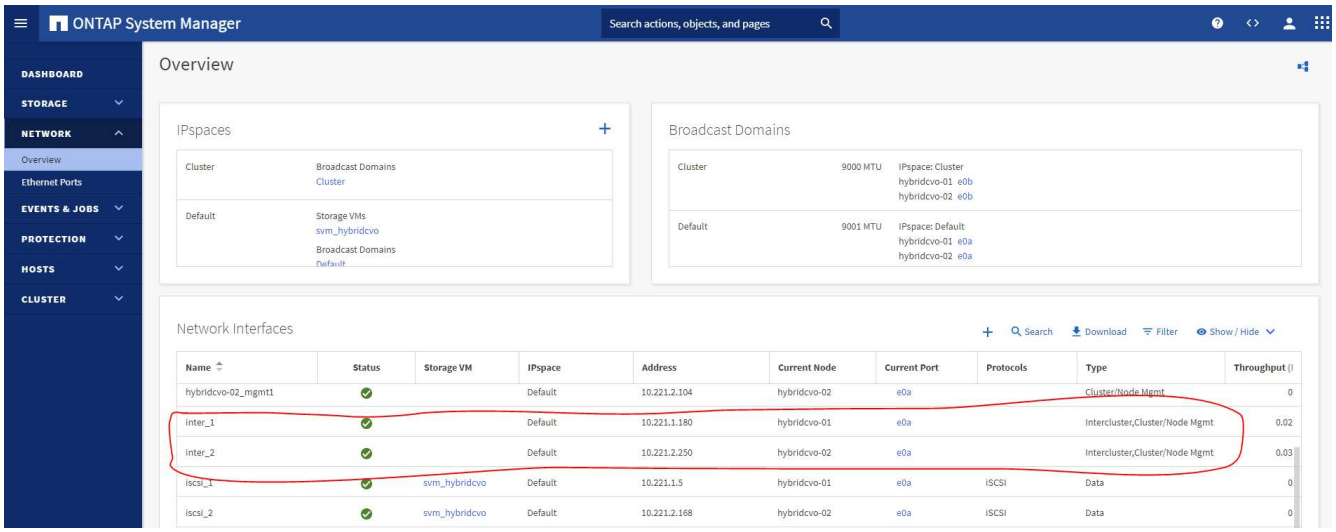

2. With the intercluster LIFs configured, cluster peering and volume replication can be set up by using drag-and-drop in NetApp Cloud Manager. See "Getting Started - AWS Public Cloud" for details.

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

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

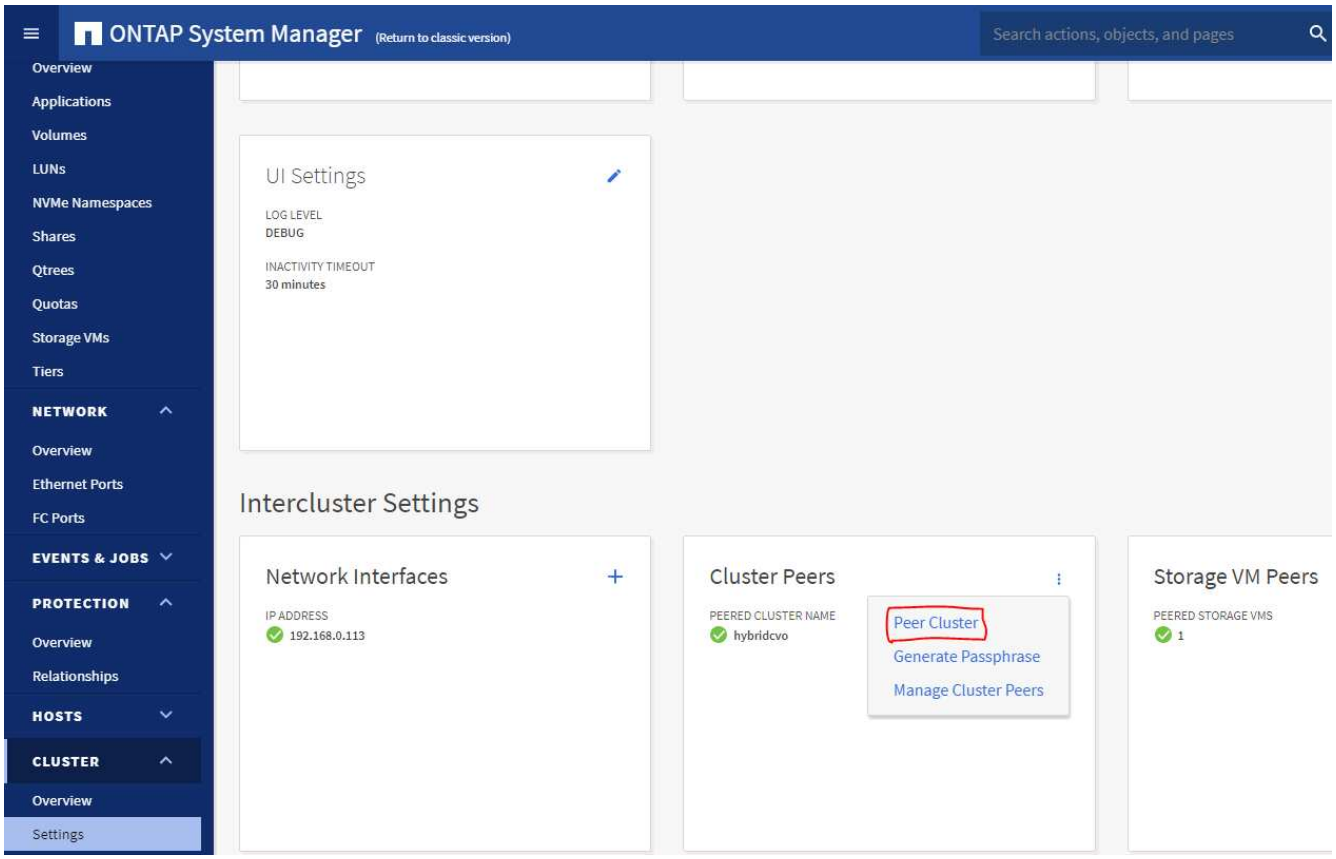

4. Go to the Volumes tab. Select the database volume to be replicated and click Protect.

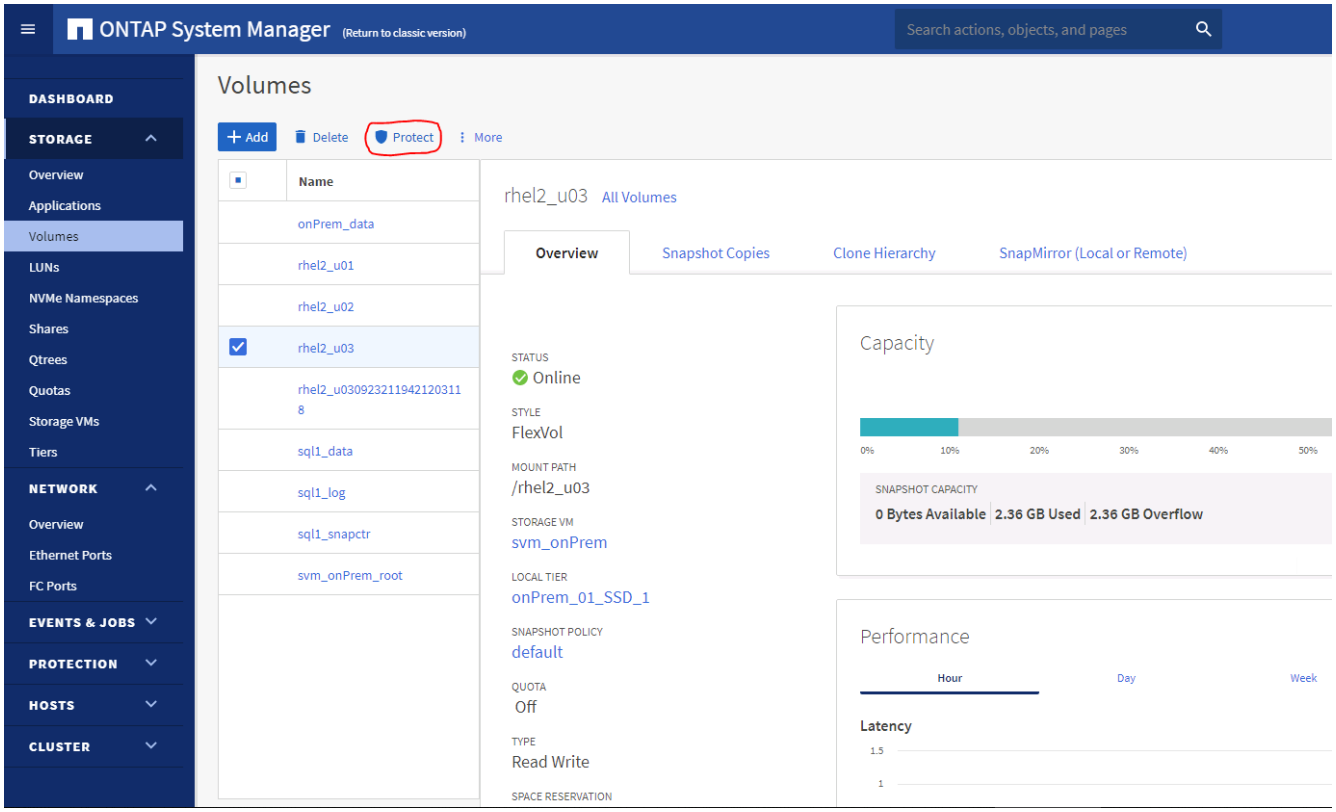

5. Set the protection policy to Asynchronous. Select the destination cluster and storage SVM.

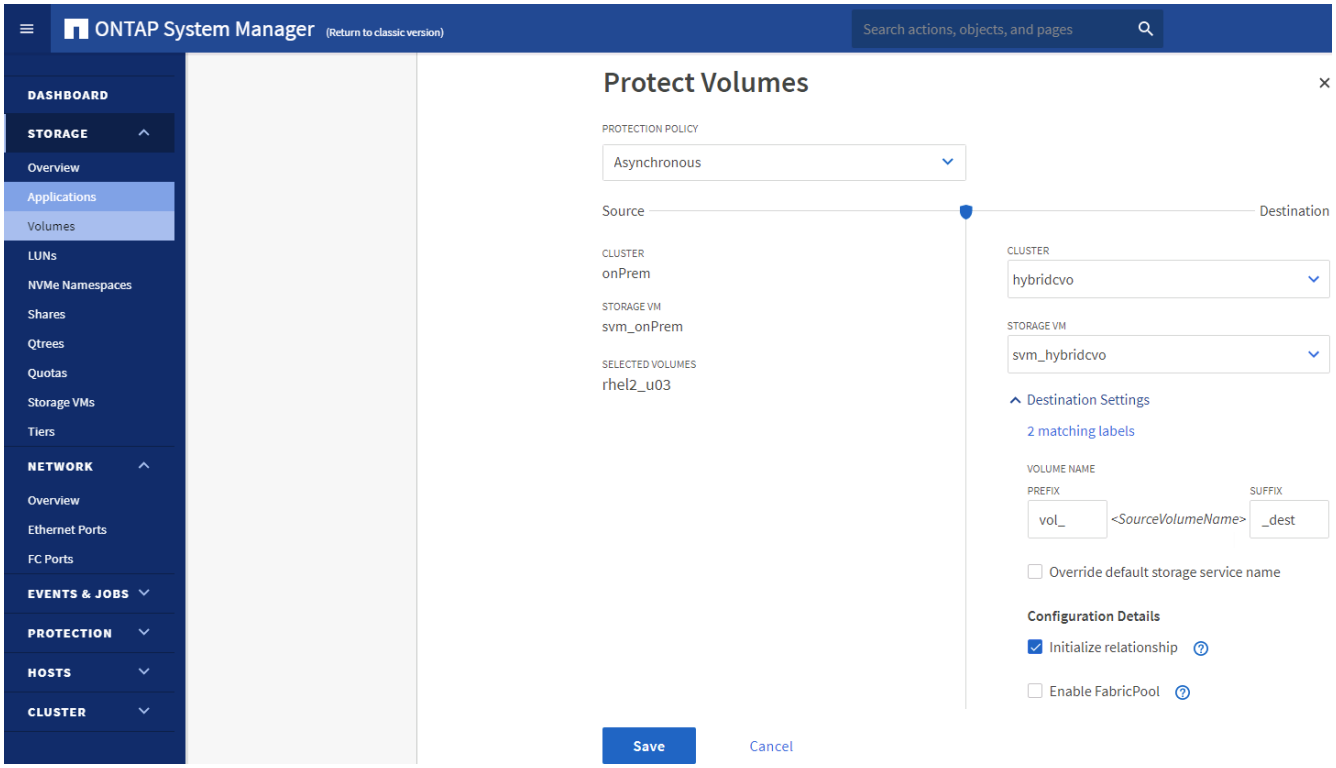

6. Validate that the volume is synced between the source and target and that the replication relationship is healthy.

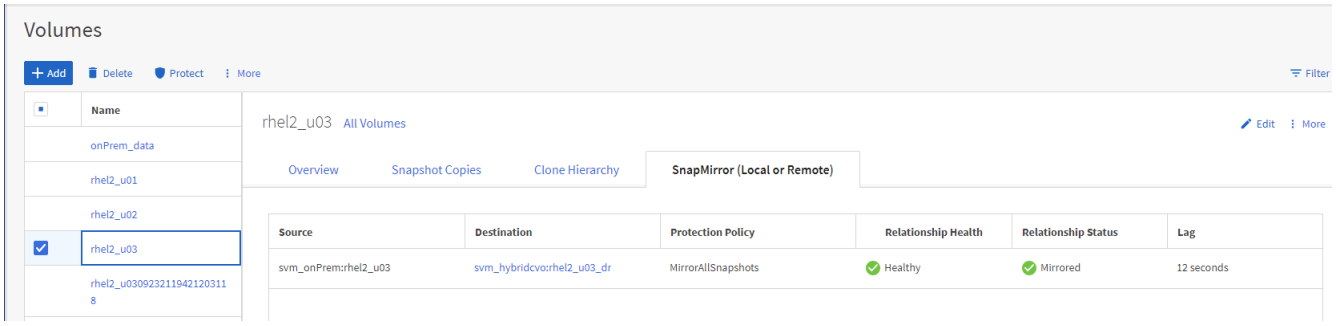

6. Add CVO database storage SVM to SnapCenter

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

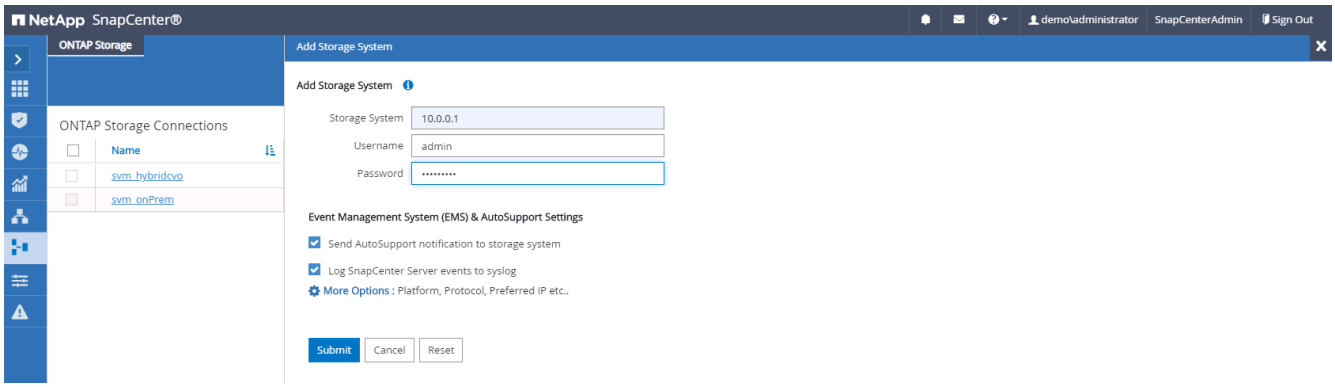

3. Click More Options to open additional storage configuration options. In the Platform field, select Cloud Volumes ONTAP, check Secondary, and then click Save.

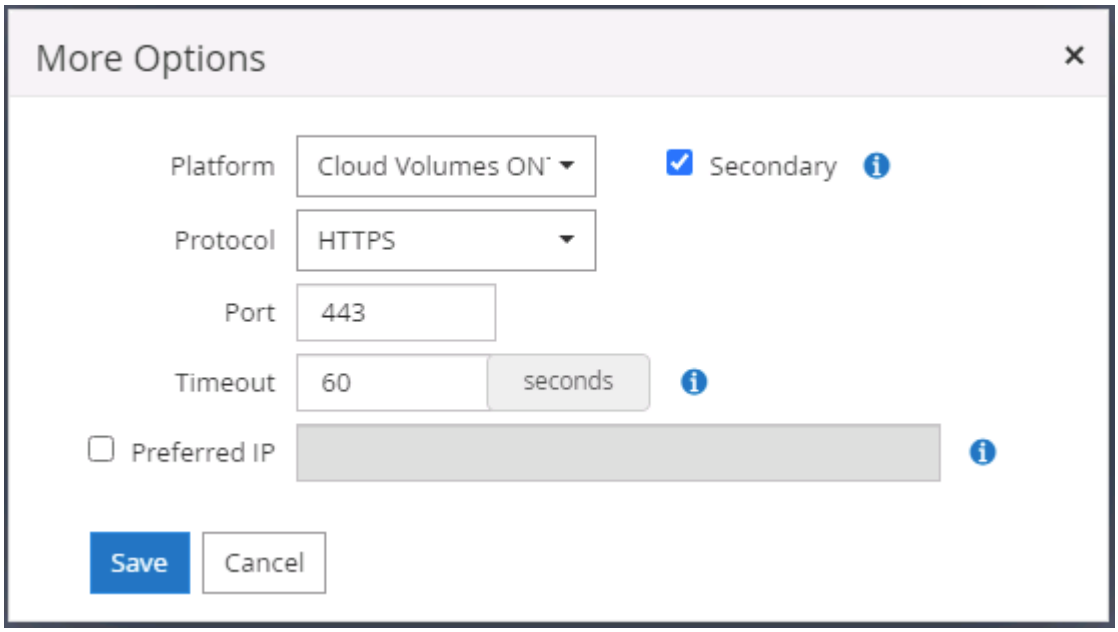

4. Assign the storage systems to SnapCenter database management user IDs as shown in 3. SnapCenter host plugin installation.

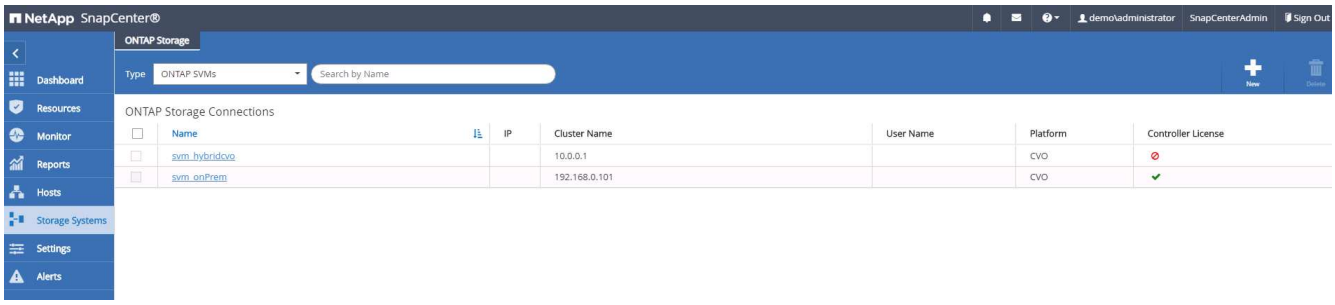

7. Setup database backup policy in SnapCenter

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

Create a full database backup policy for Oracle

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

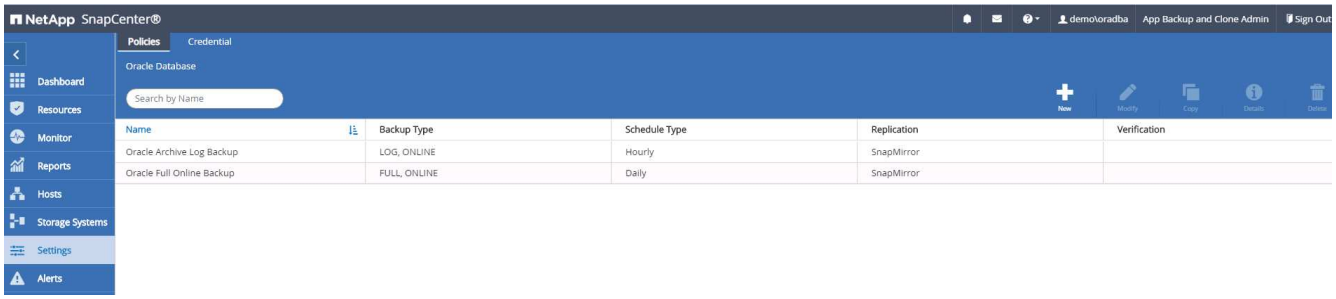

2. Click New to launch a new backup policy creation workflow or choose an existing policy for modification.

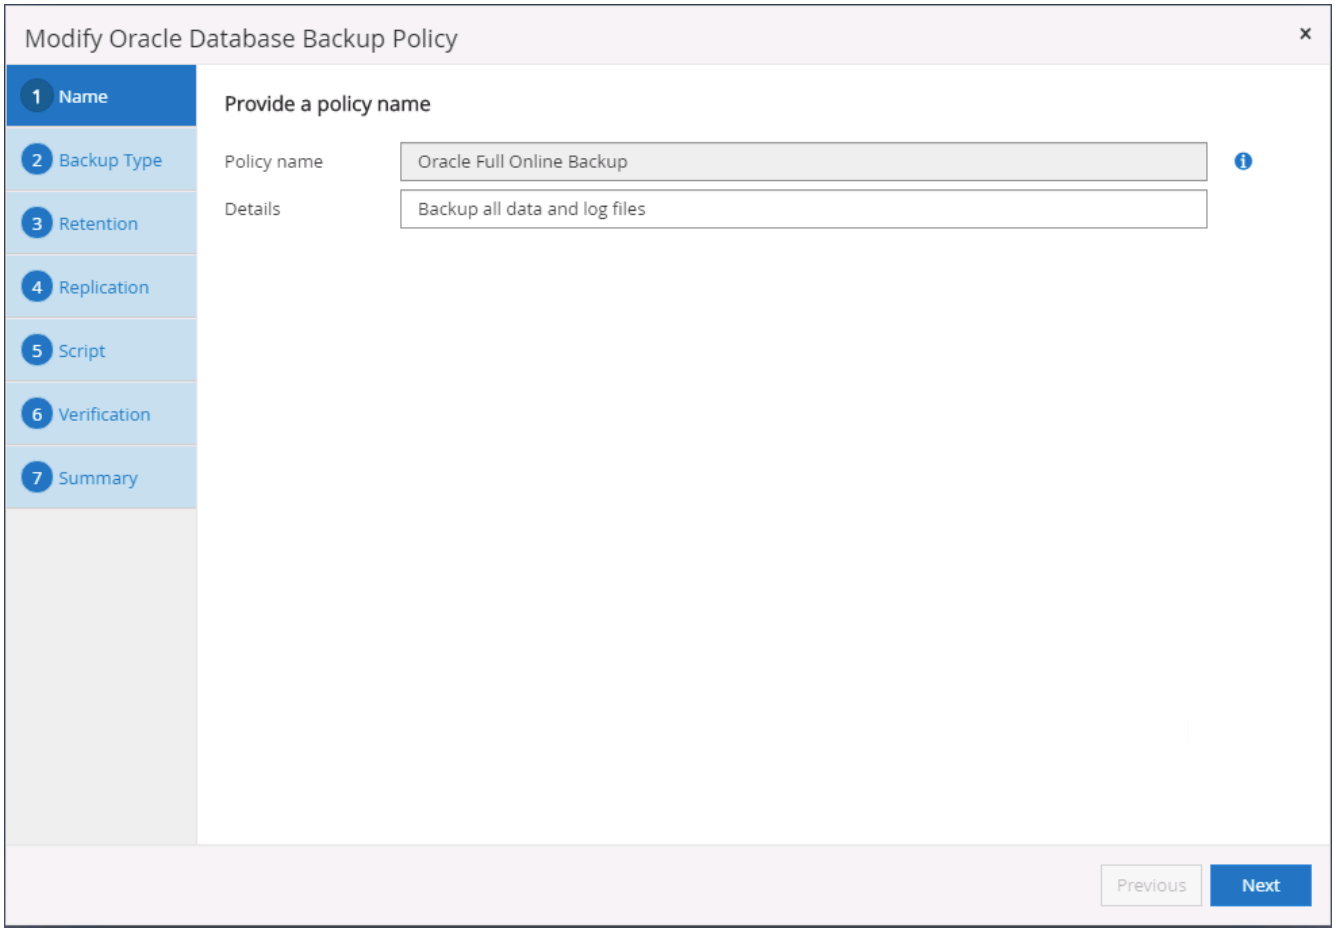

3. Select the backup type and schedule frequency.

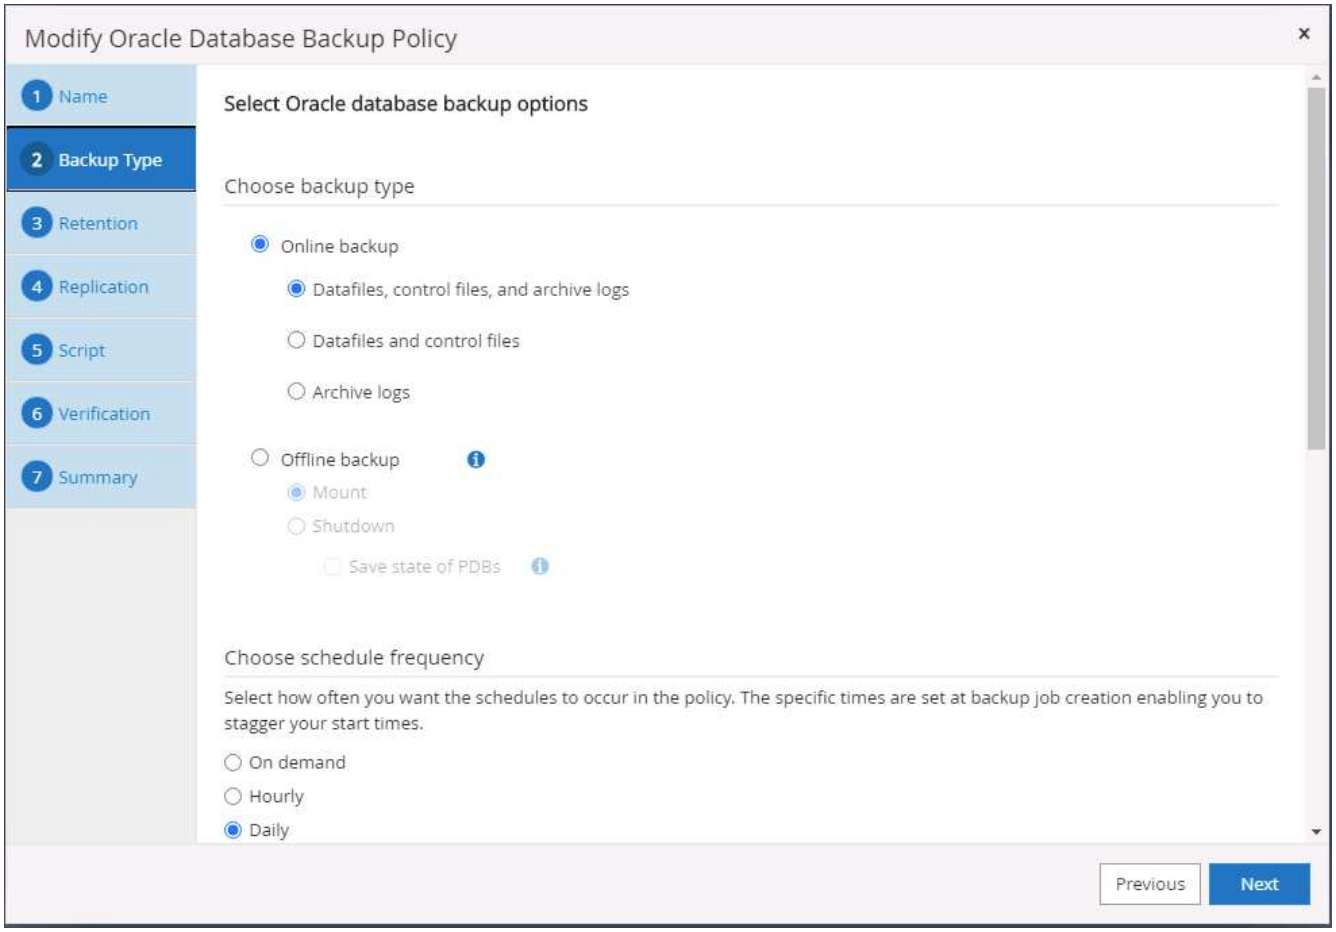

4. Set the backup retention setting. This defines how many full database backup copies to keep.

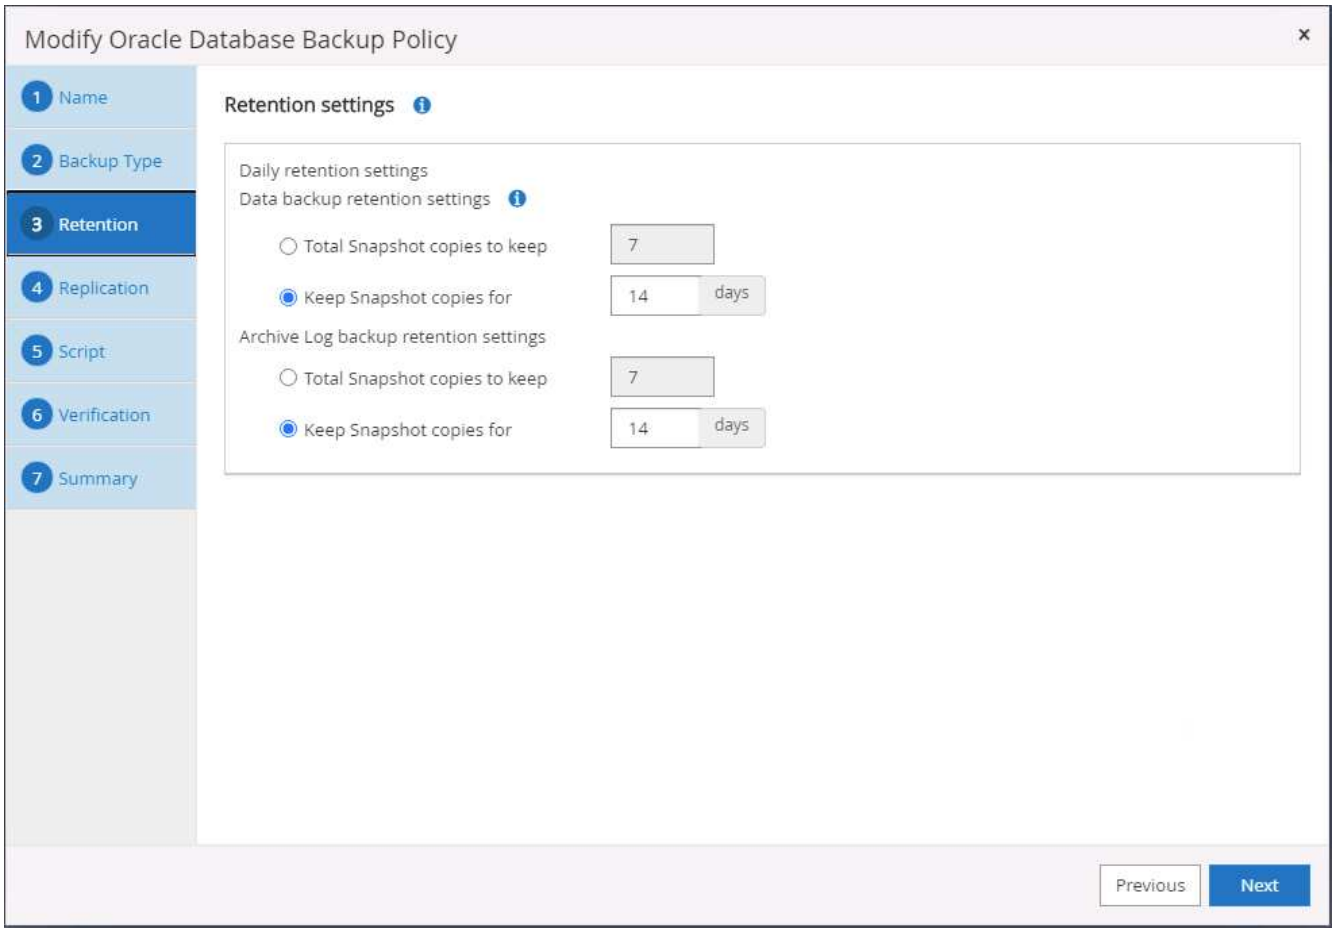

5. Select the secondary replication options to push local primary snapshots backups to be replicated to a secondary location in cloud.

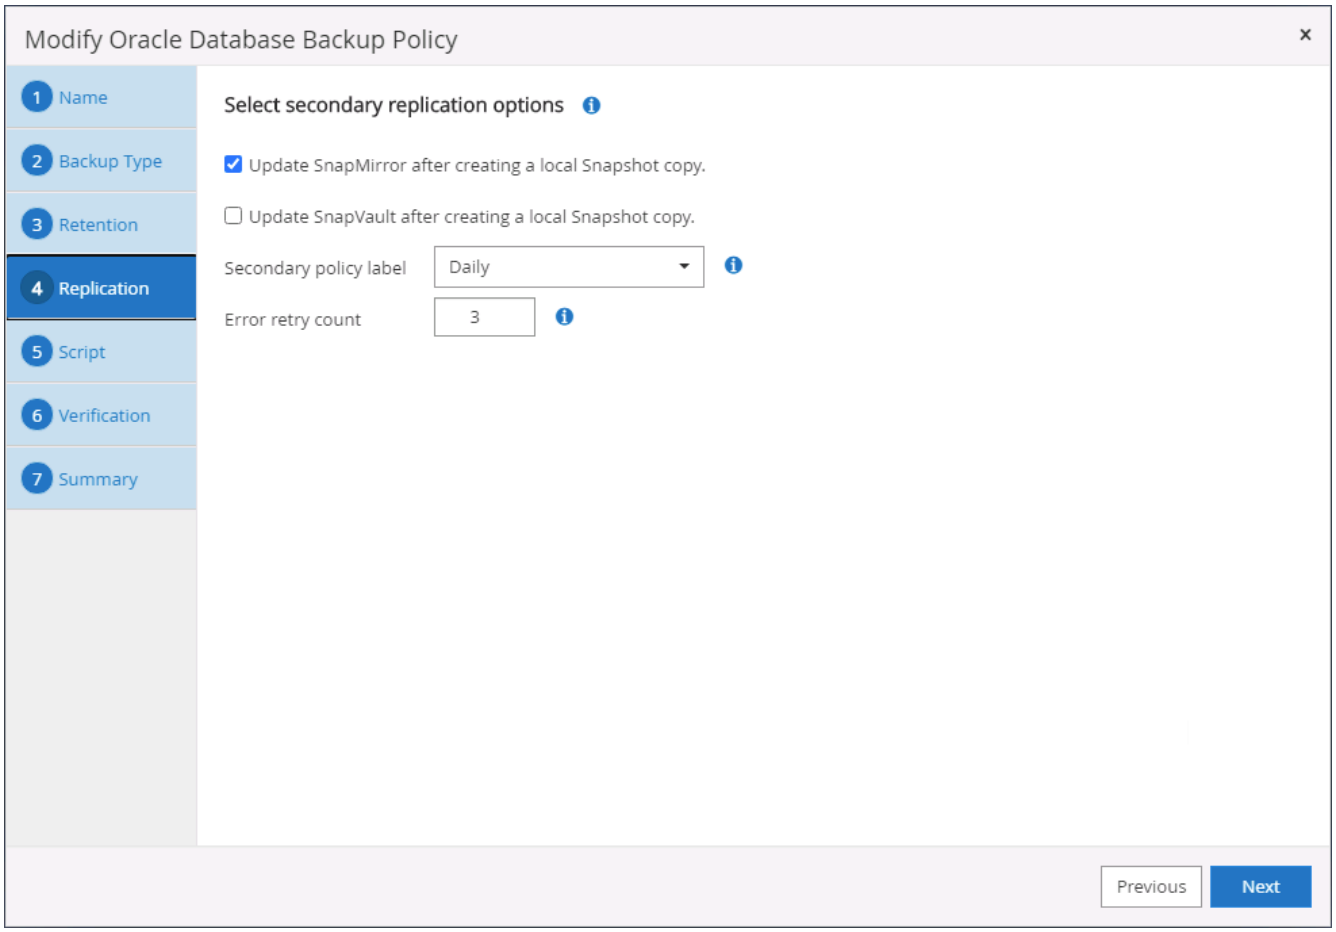

6. Specify any optional script to run before and after a backup run.

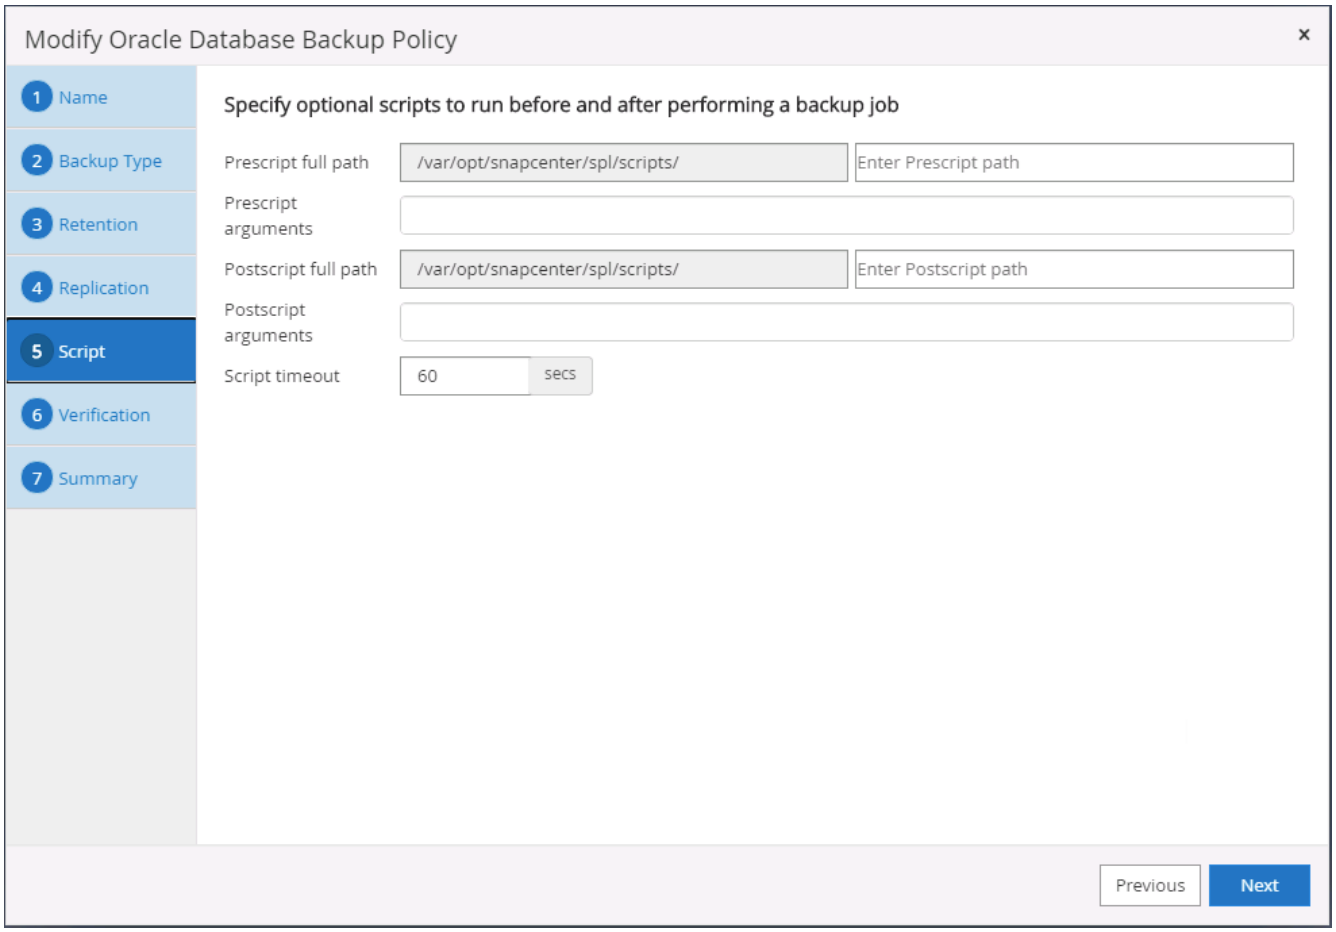

7. Run backup verification if desired.

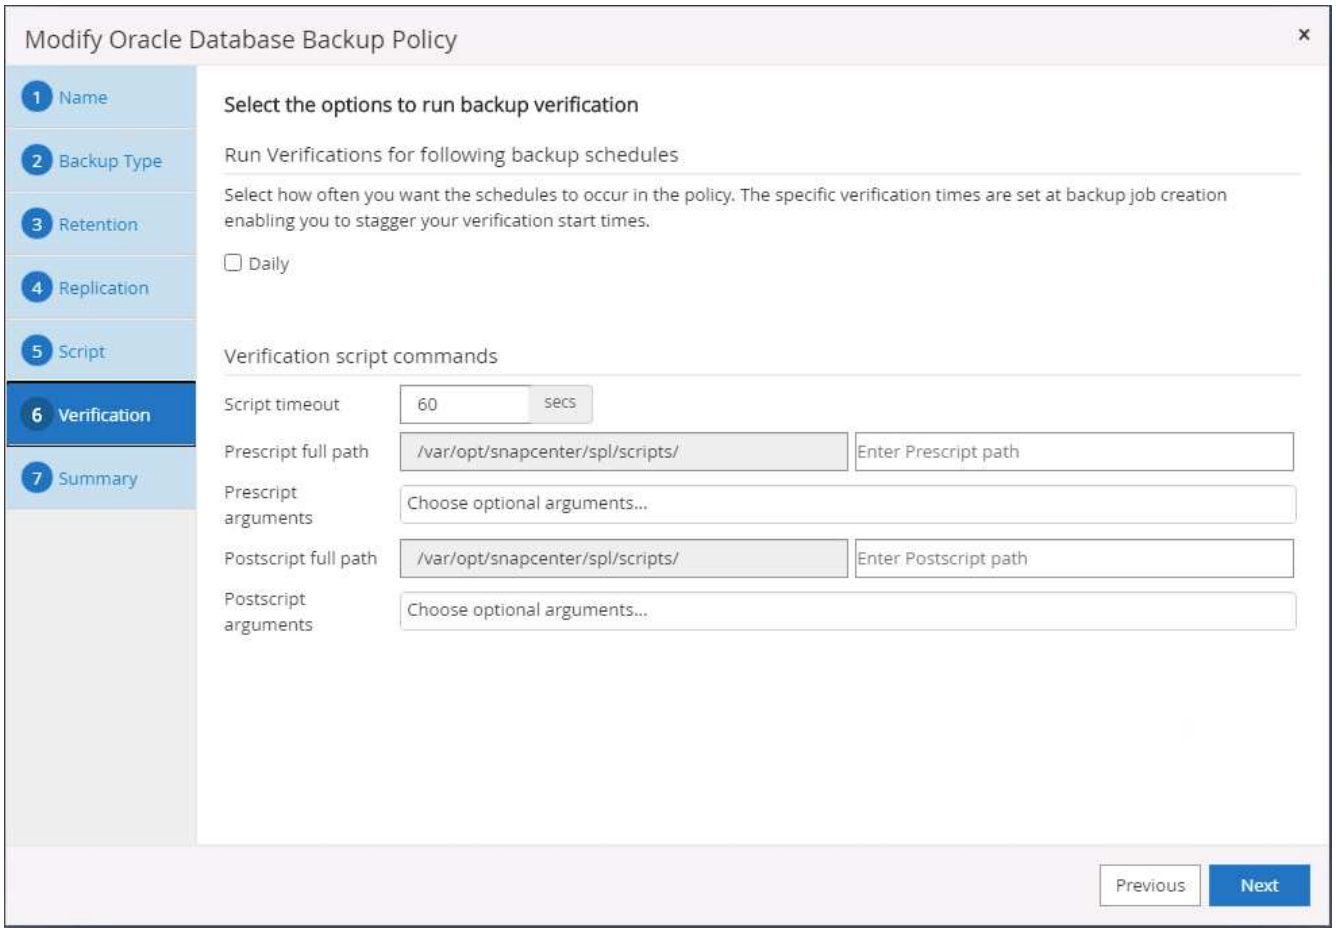
Modify Oracle Database Backup Policy
1 Name
2 Backup Type
3 Retention
4 Replication
5 Script
6 Verification
7 Summary
Select the options to run backup verification
Run Verifications for following backup schedules
Select how often you want the schedules to occur in the policy. The specific verification 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.

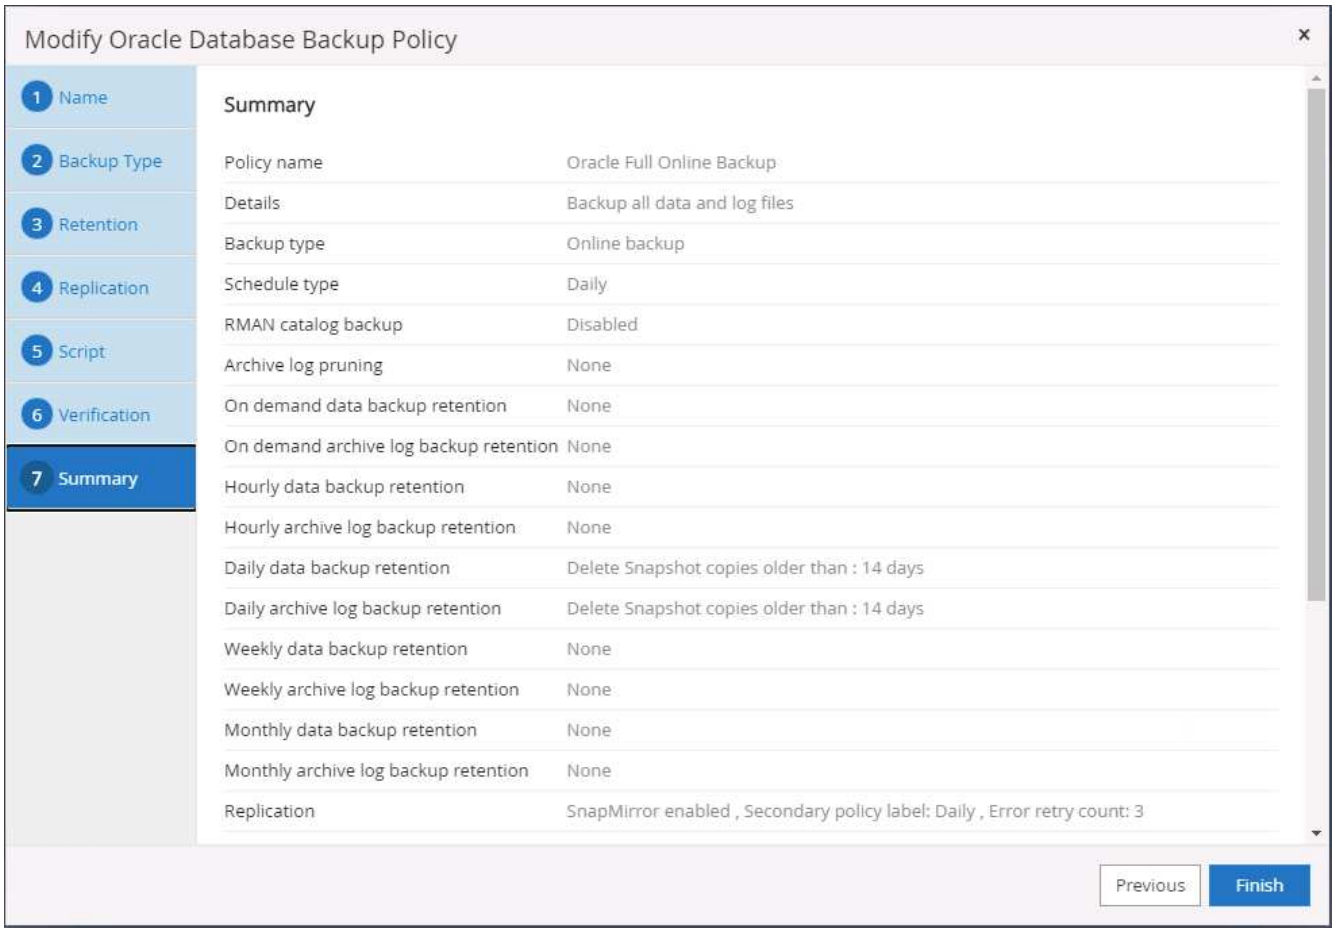

Create a database log backup policy for Oracle

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

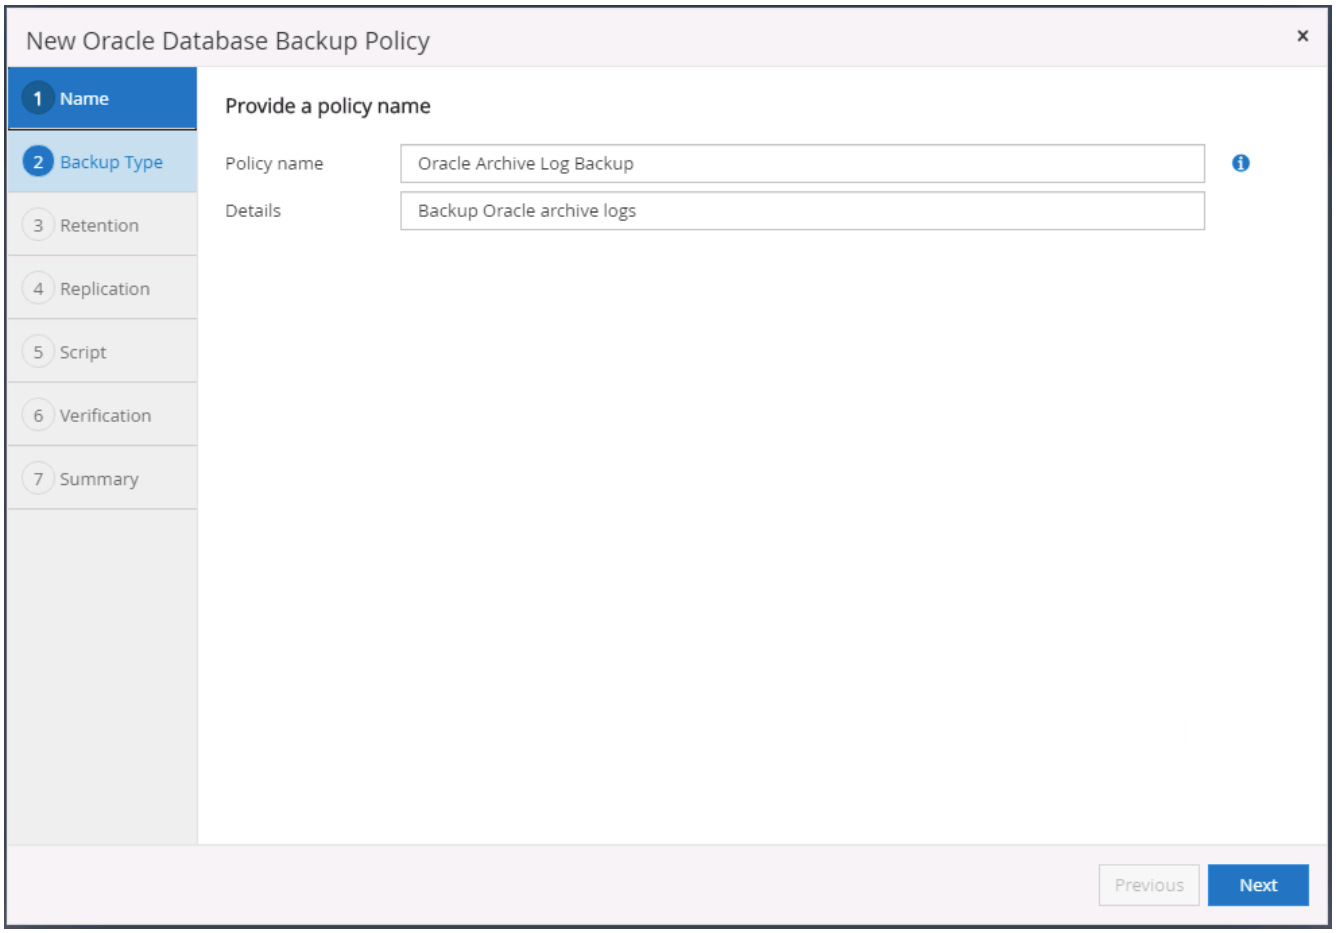

3. Select the backup type and schedule frequency.

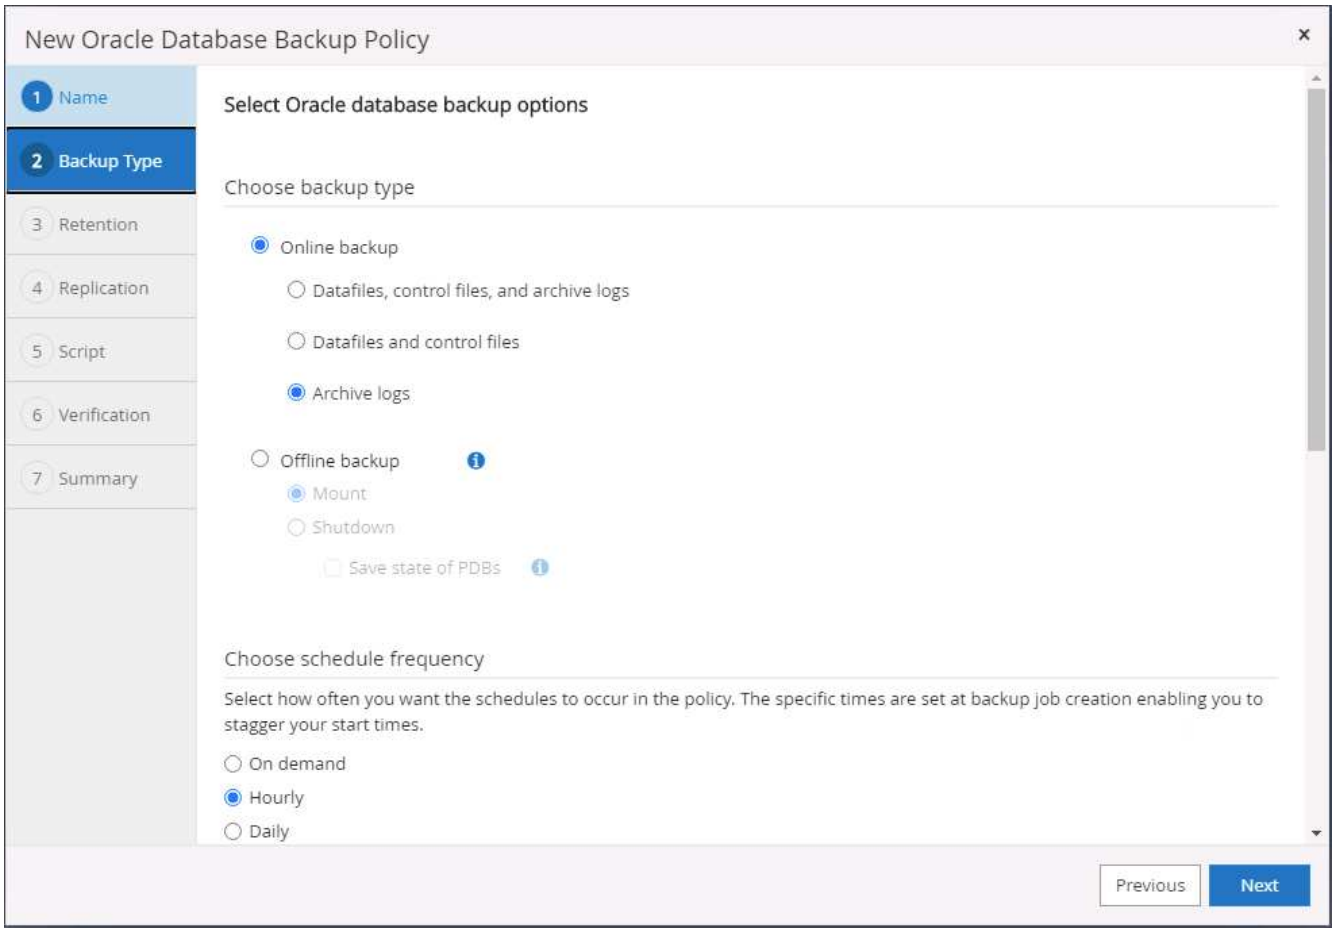

4. Set the log retention period.

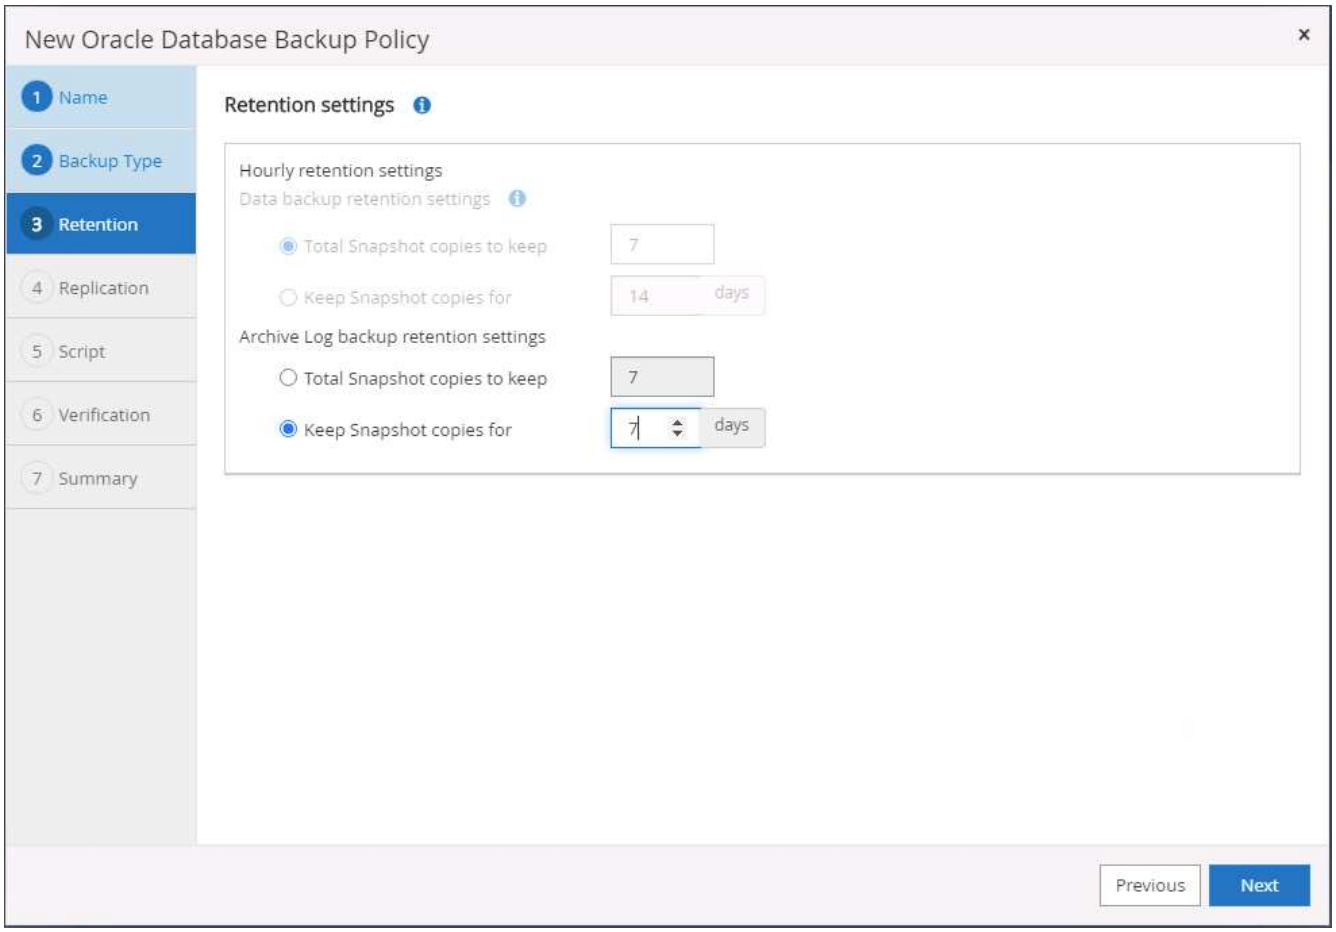

5. Enable replication to a secondary location in the public cloud.

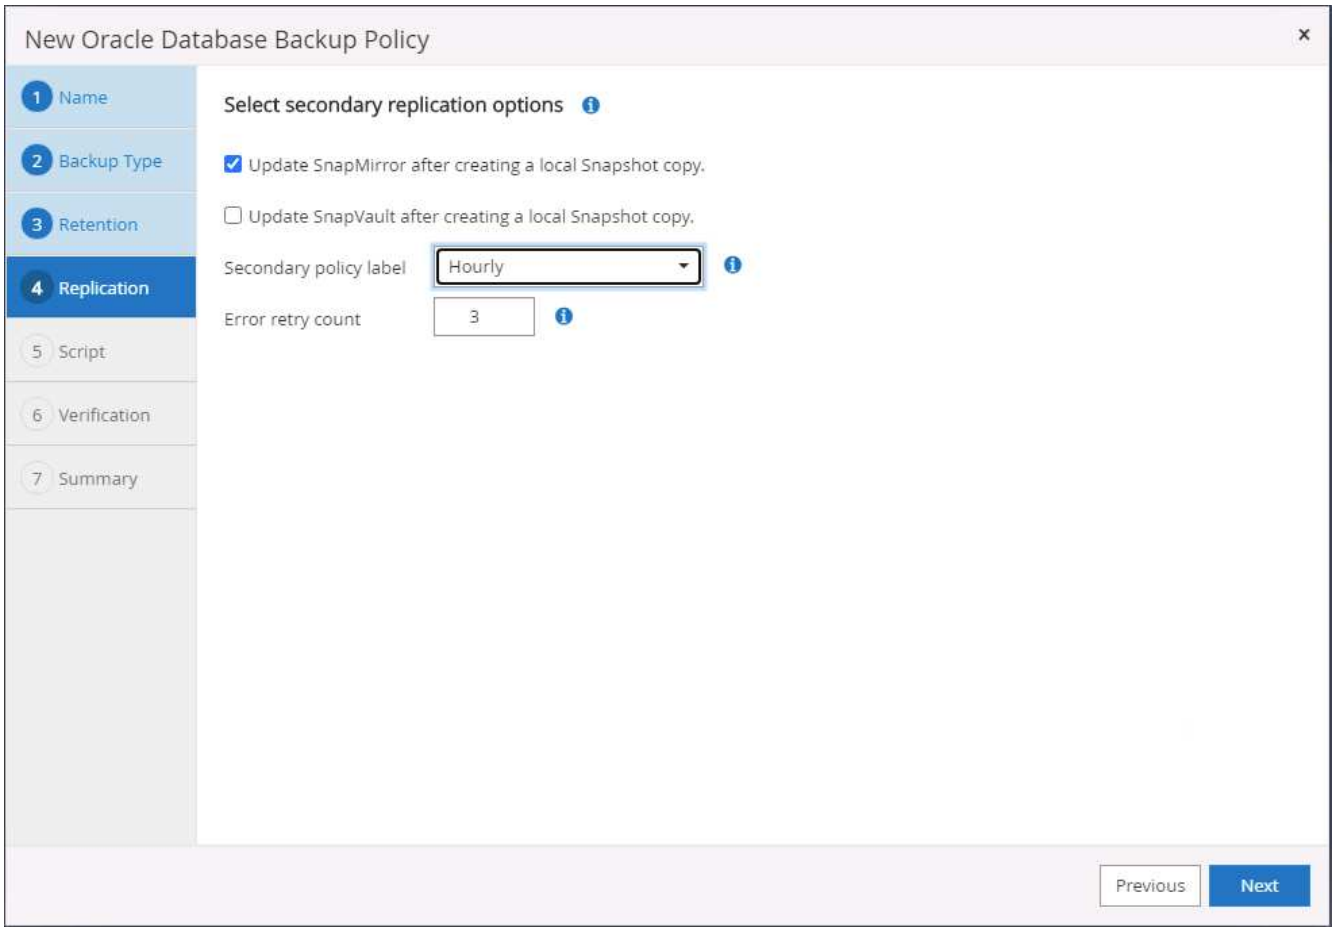

6. Specify any optional scripts to run before and after log backup.

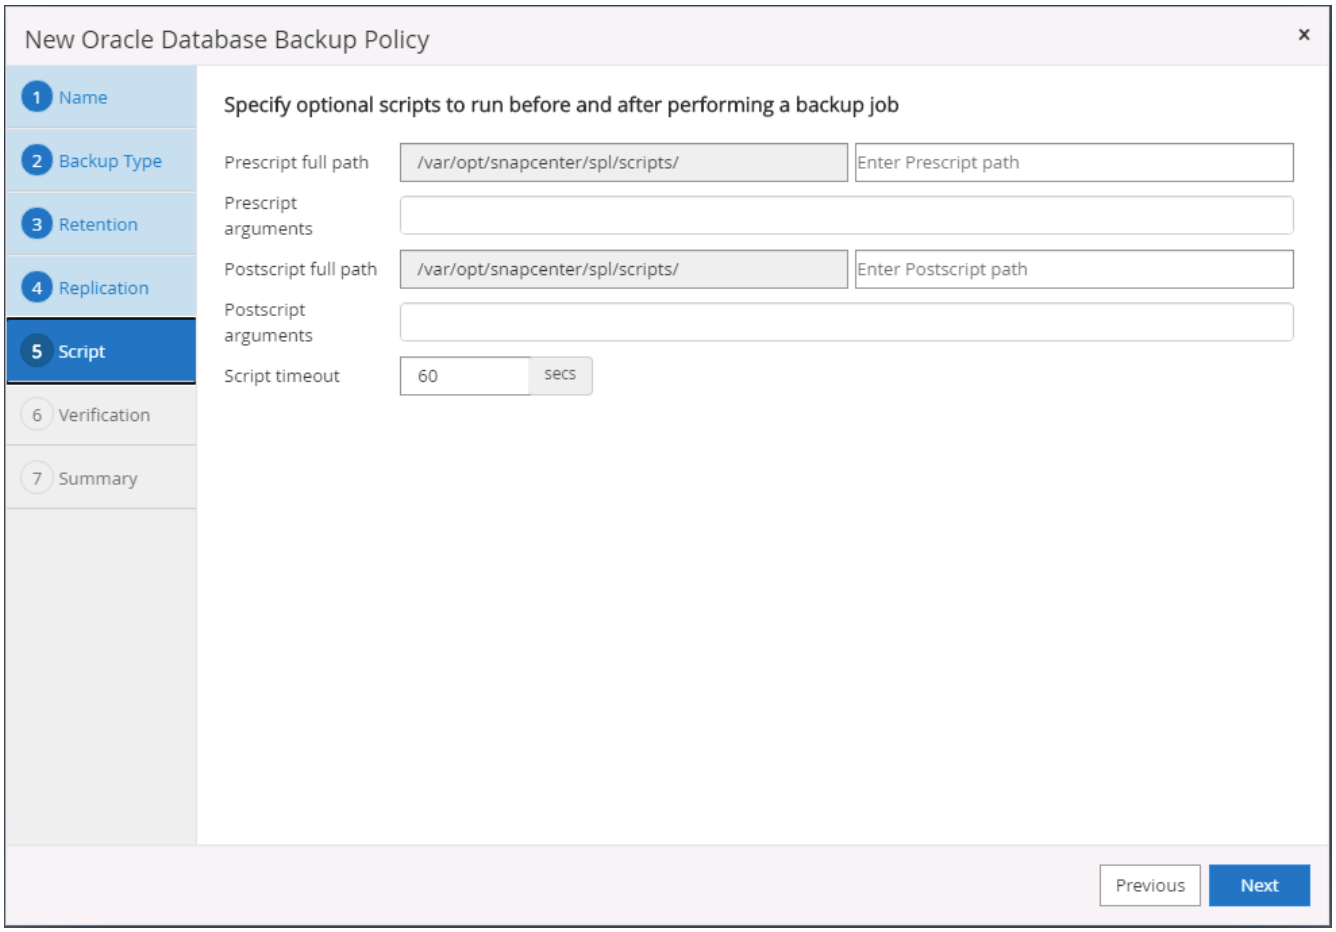

7. Specify any backup verification scripts.

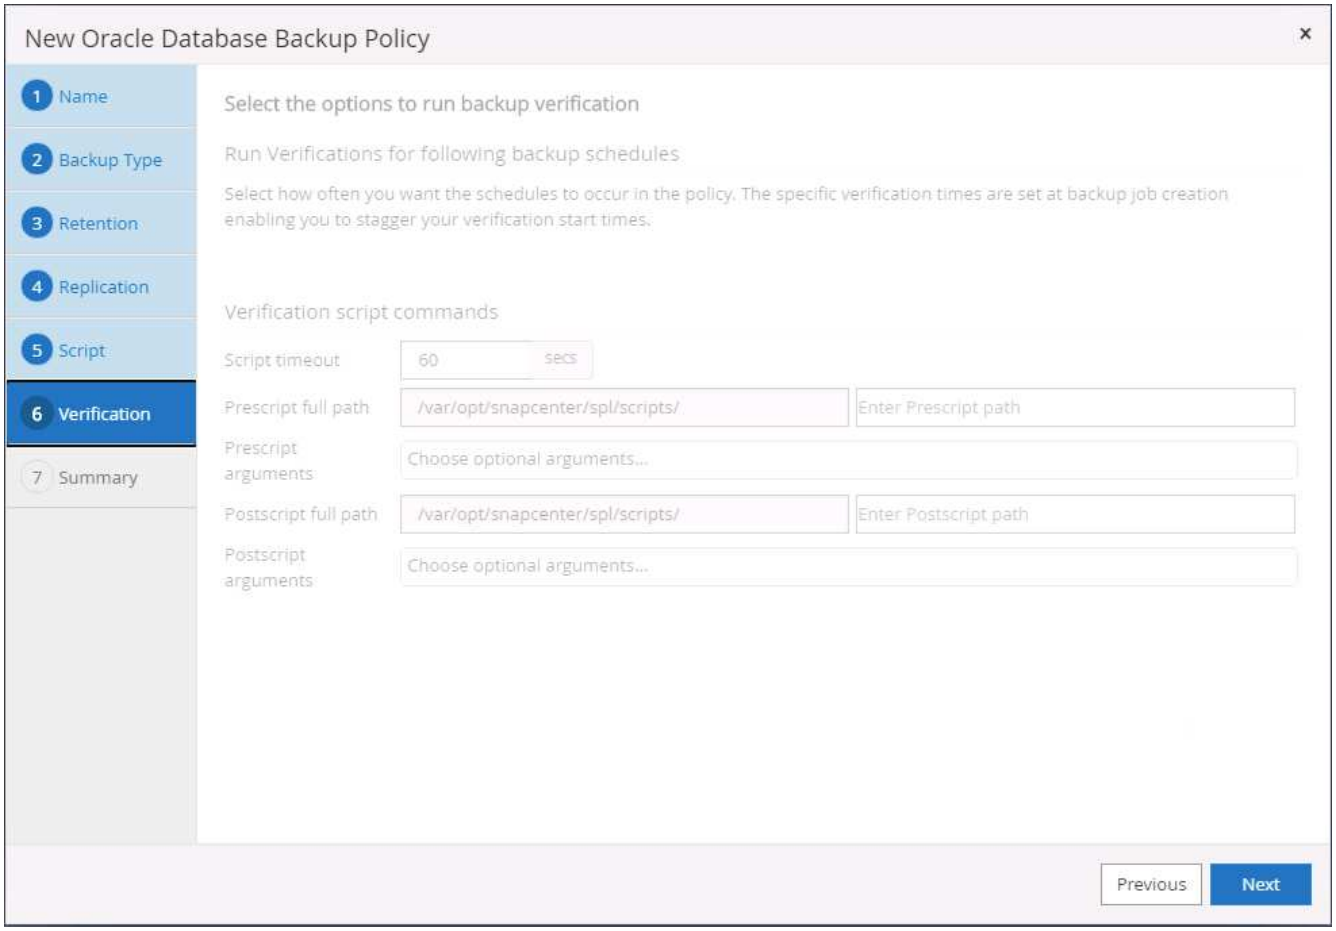
New Oracle Database Backup Policy
1 Name
2 Backup Type
3 Retention
4 Replication
5 Script
6 Verification
7 Summary
Select the options to run backup verification
Run Verifications for following backup schedules
Select how often you want the schedules to occur in the policy. The specific verification times are set at backup job creation enabling you to stagger your verification start times.
Verification script commands
Script timeout
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.

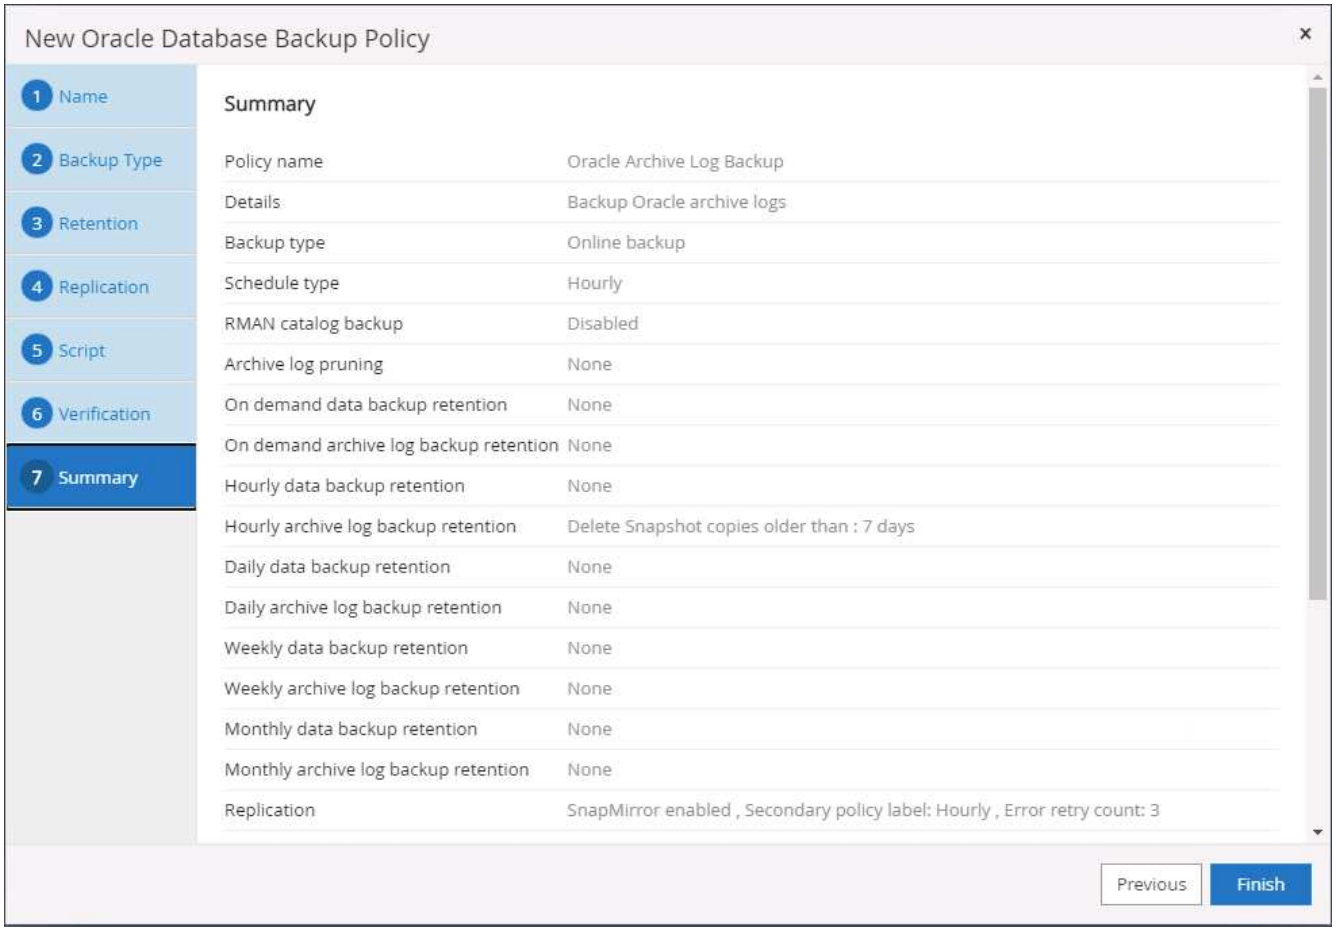

Create a full database backup policy for SQL

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

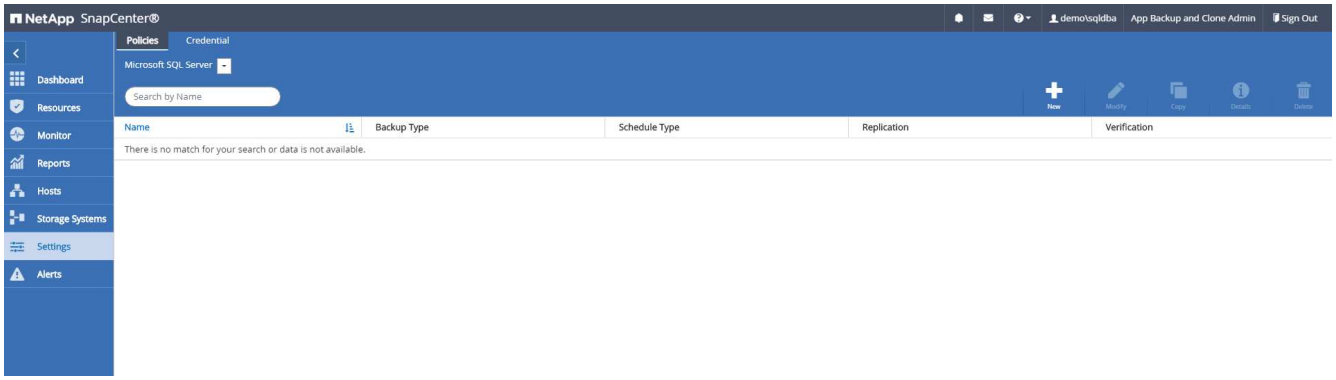

2. Click New to launch a new backup policy creation workflow, or choose an existing policy for modification.

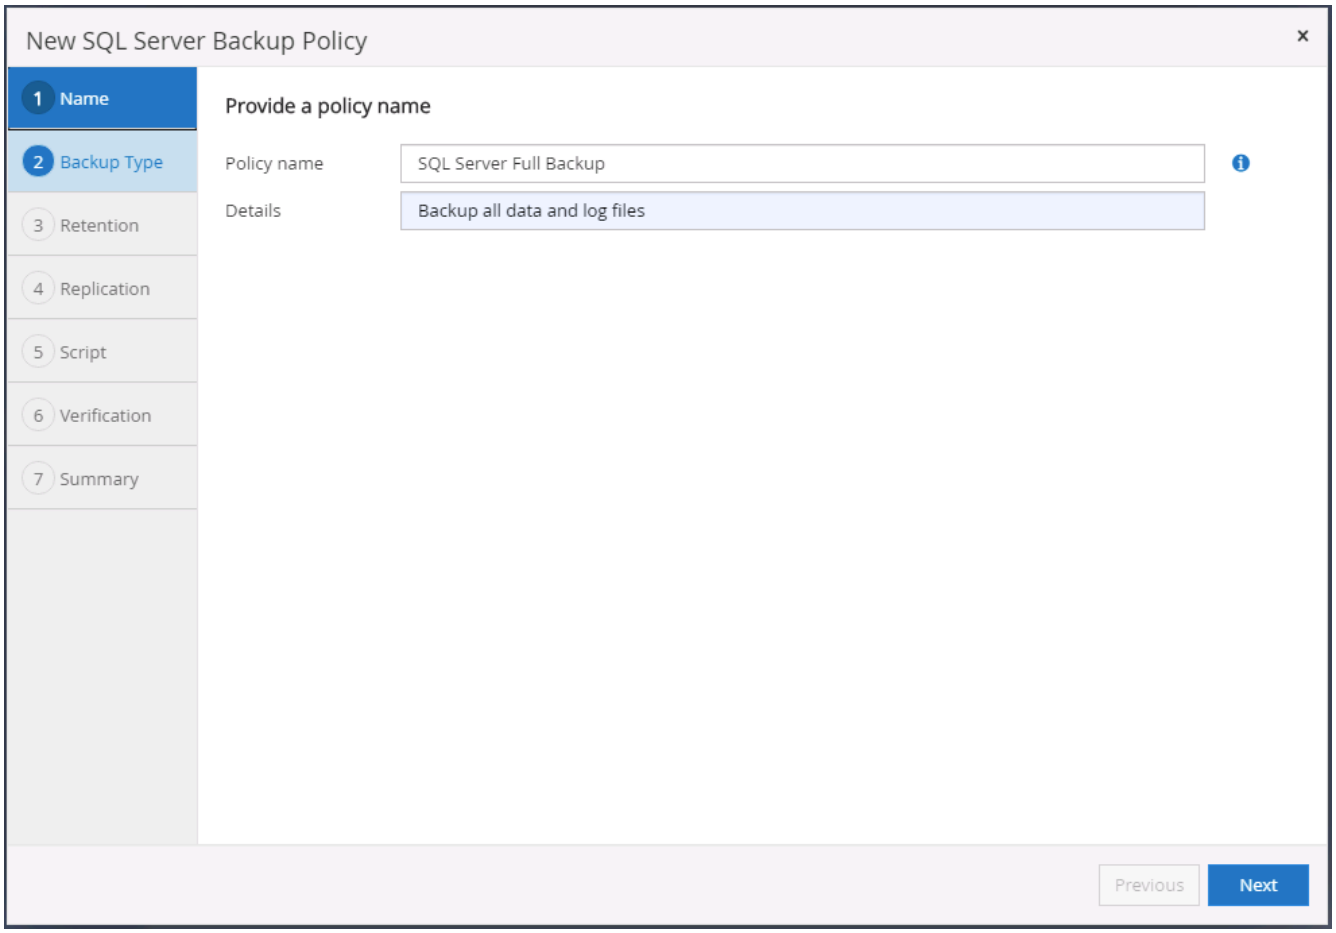

3. Define the backup option and schedule frequency. For SQL Server configured with an availability group, a preferred backup replica can be set.

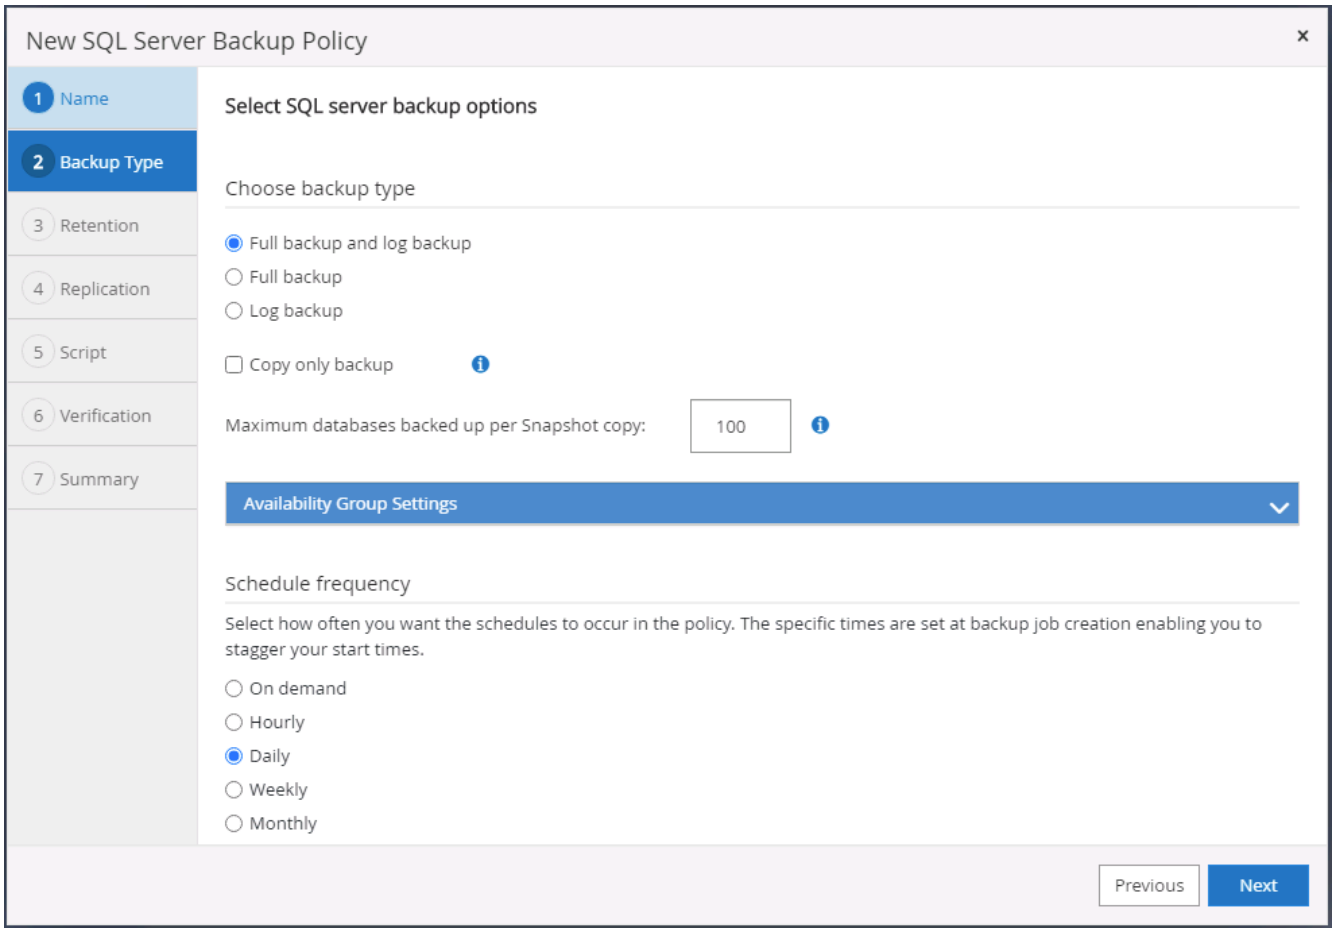

4. Set the backup retention period.

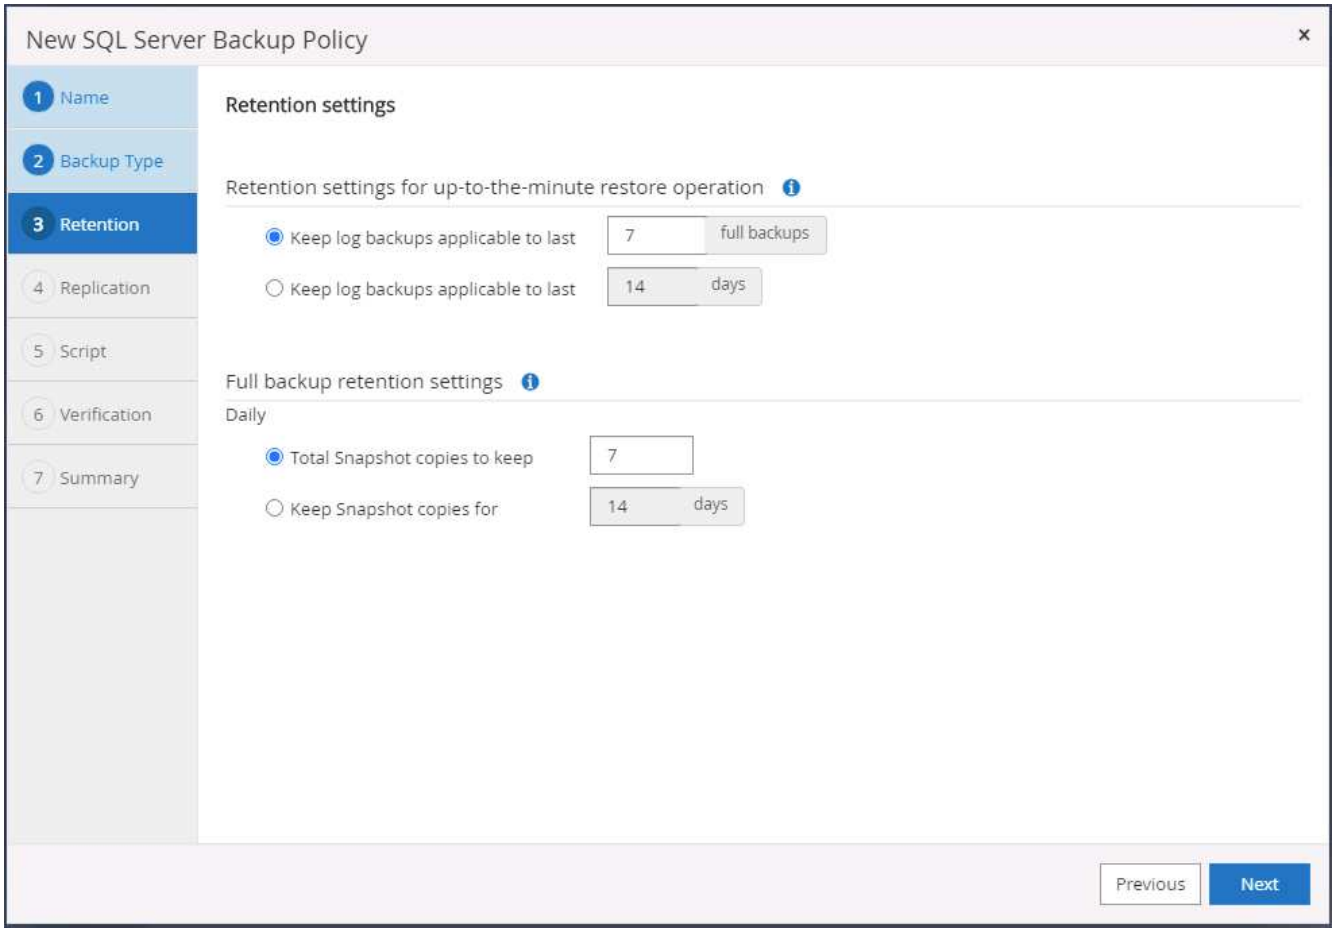

5. Enable backup copy replication to a secondary location in cloud.

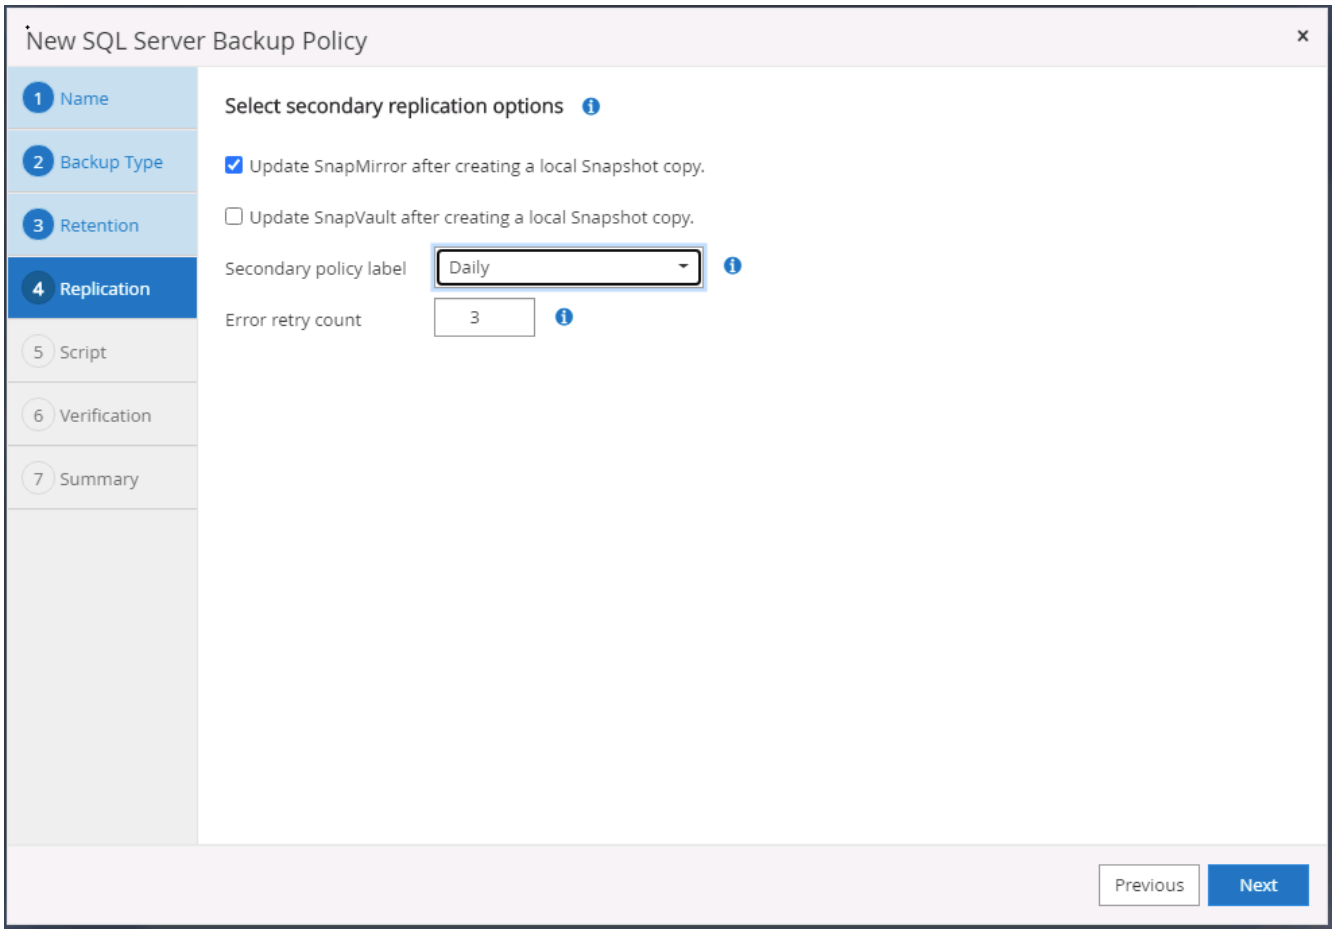

6. Specify any optional scripts to run before or after a backup job.

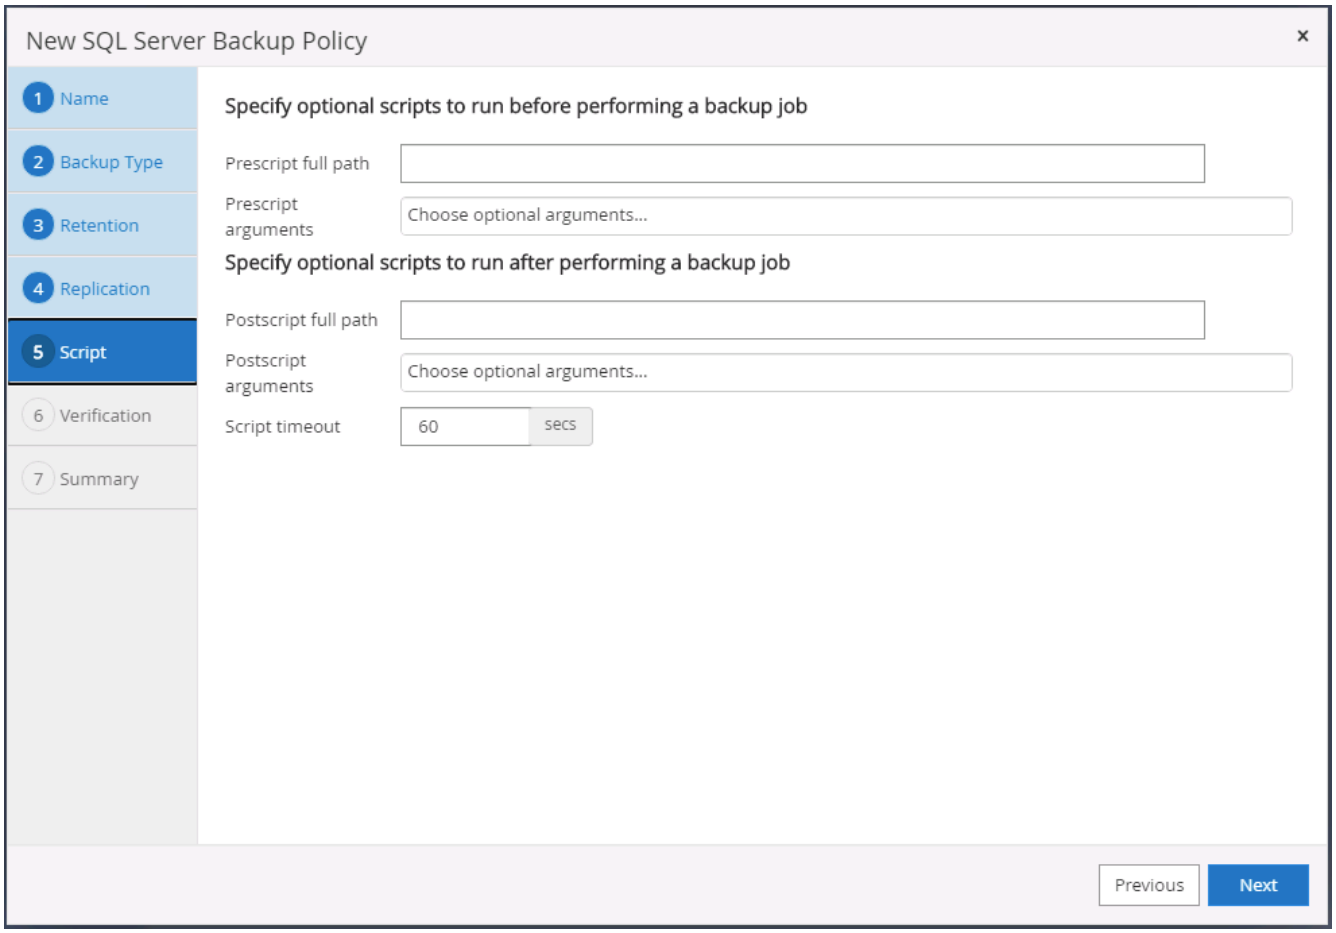

7. Specify the options to run backup verification.

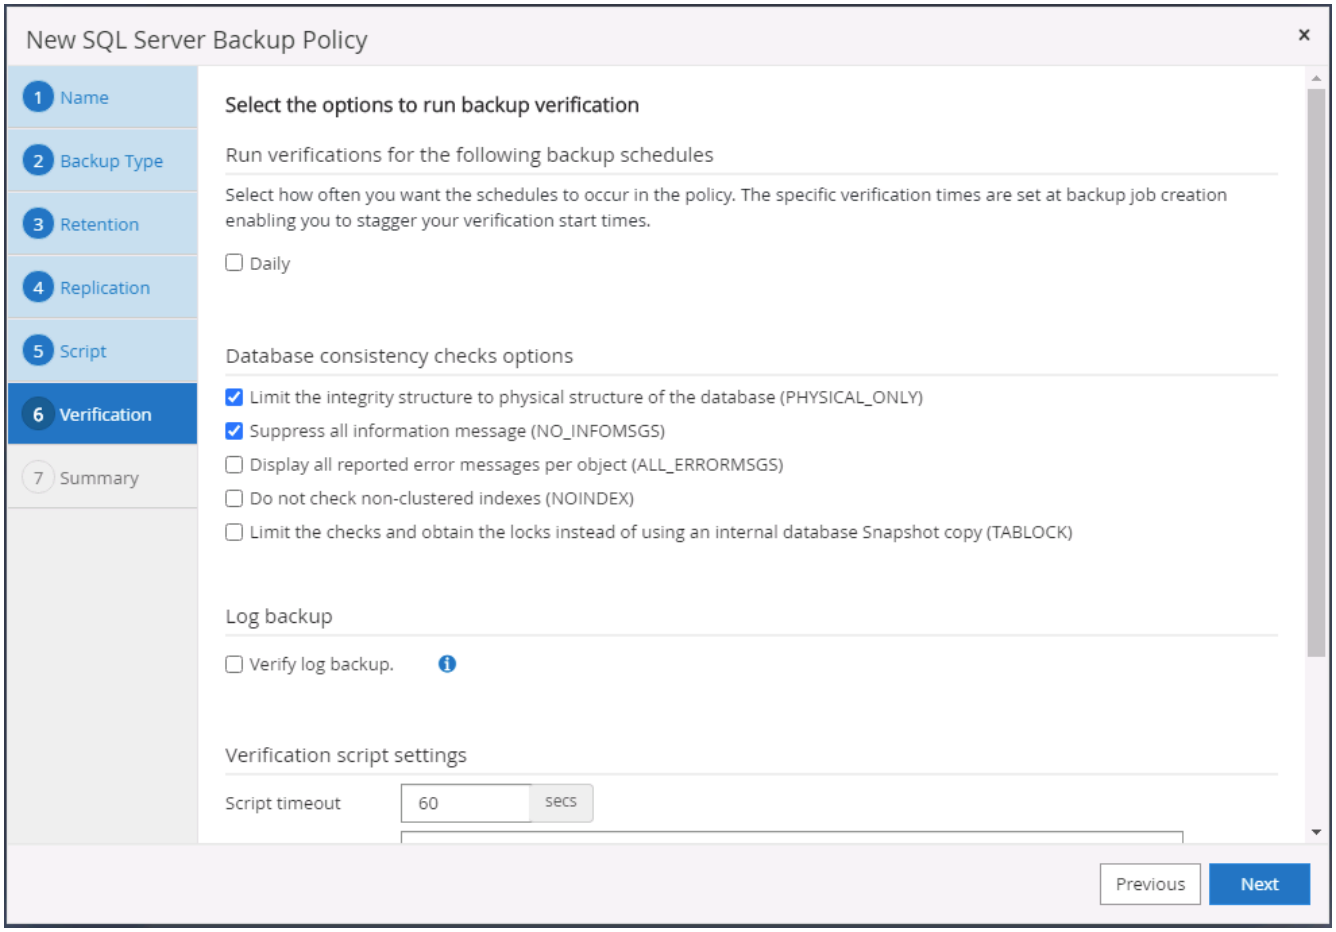
New SQL Server Backup Policy
1 Name
2 Backup Type
3 Retention
4 Replication
5 Script
6 Verification
7 Summary
Select the options to run backup verification
Run verifications for the following backup schedules
Select how often you want the schedules to occur in the policy. The specific verification times 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.
Verification script settings
Script timeout
60
secs
Previous
Next

8. Summary.

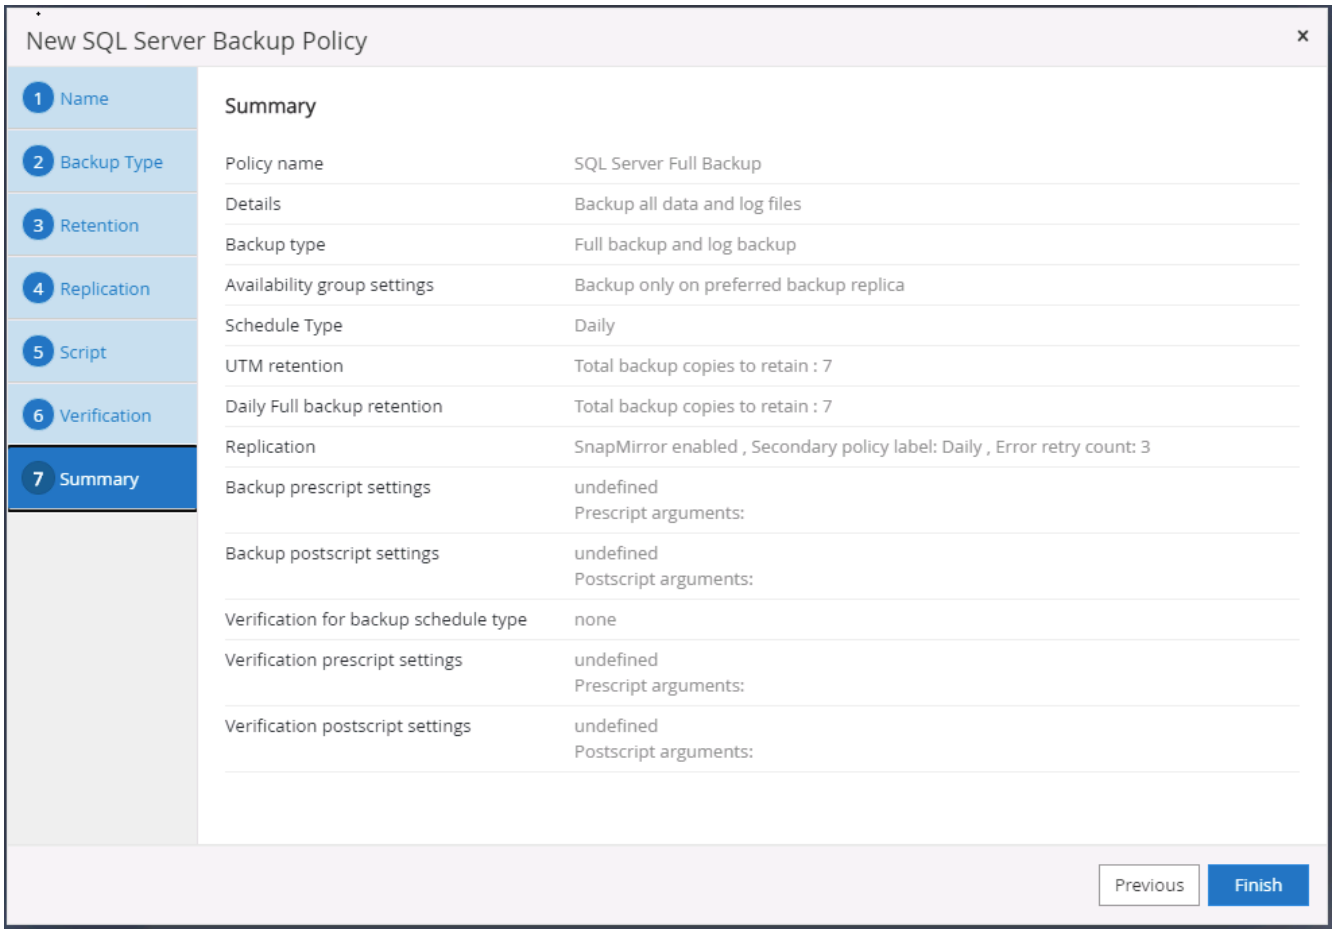

Create a database log backup policy for SQL.

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

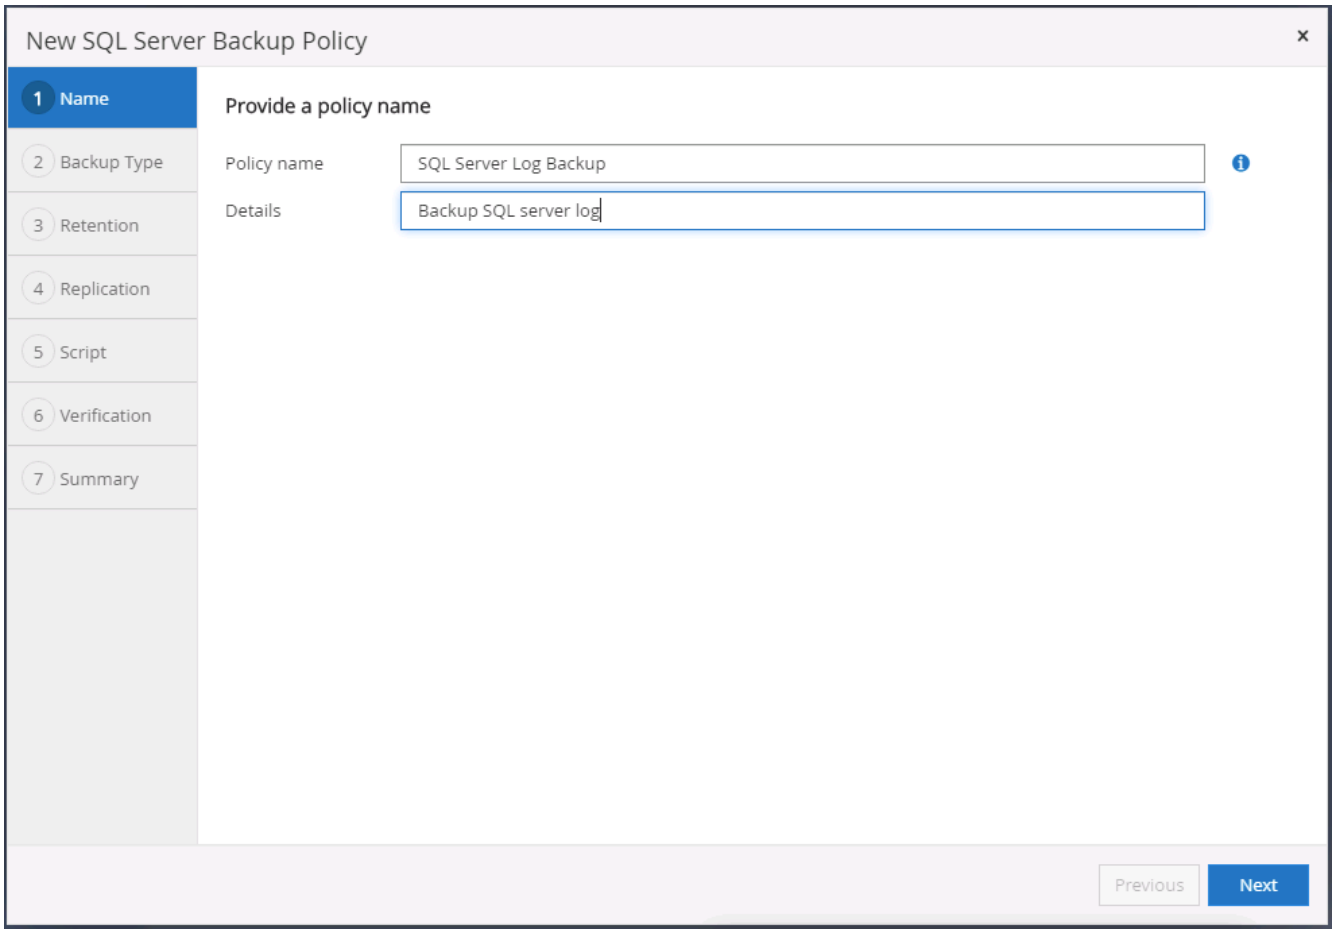

2. Define the log backup option and schedule frequency. For SQL Server configured with a availability group, a preferred backup replica can be set.

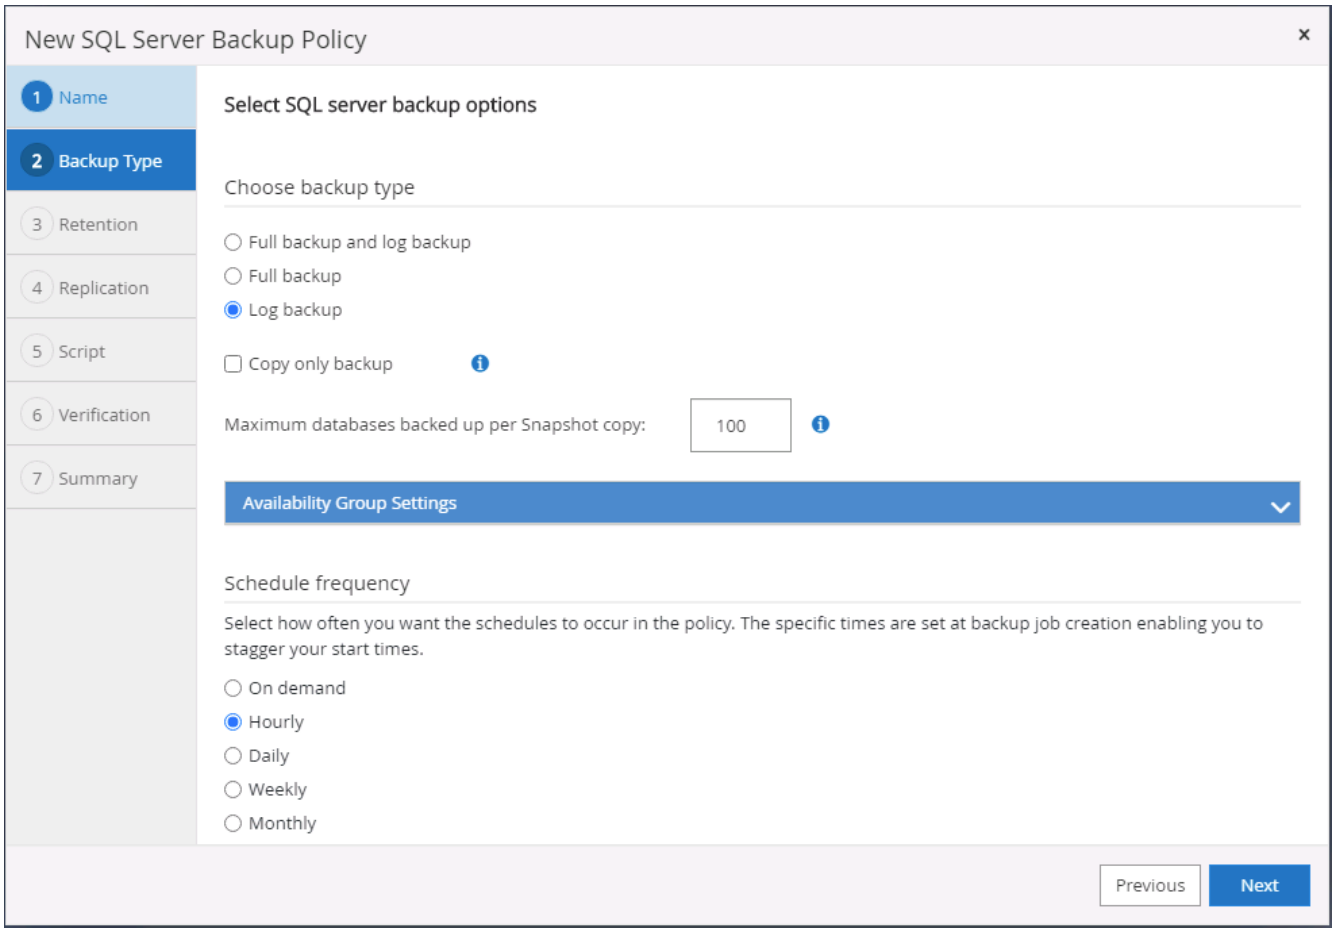

3. SQL server data backup policy defines the log backup retention; accept the defaults here.

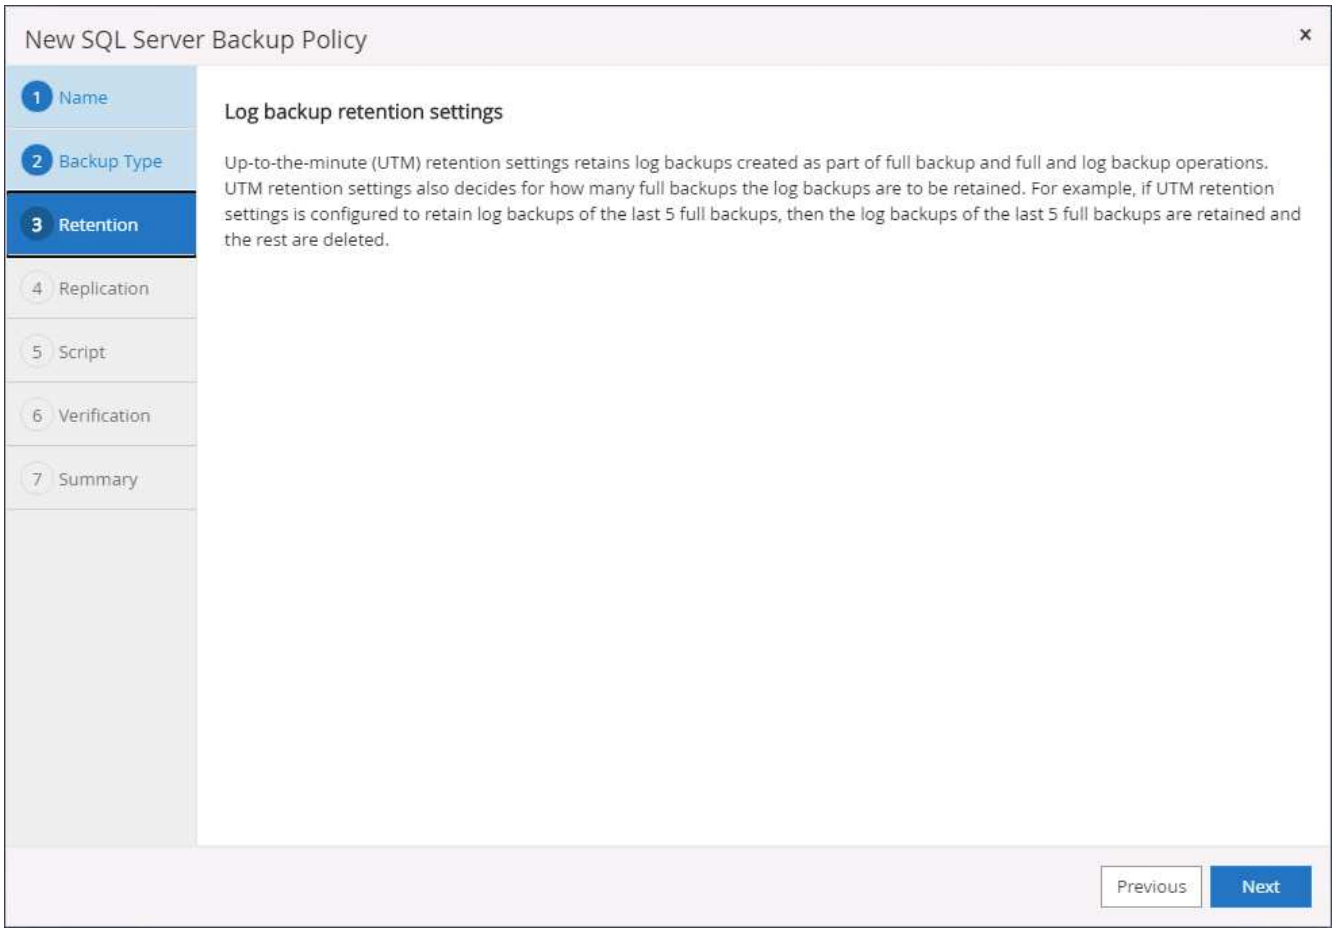

4. Enable log backup replication to secondary in the cloud.

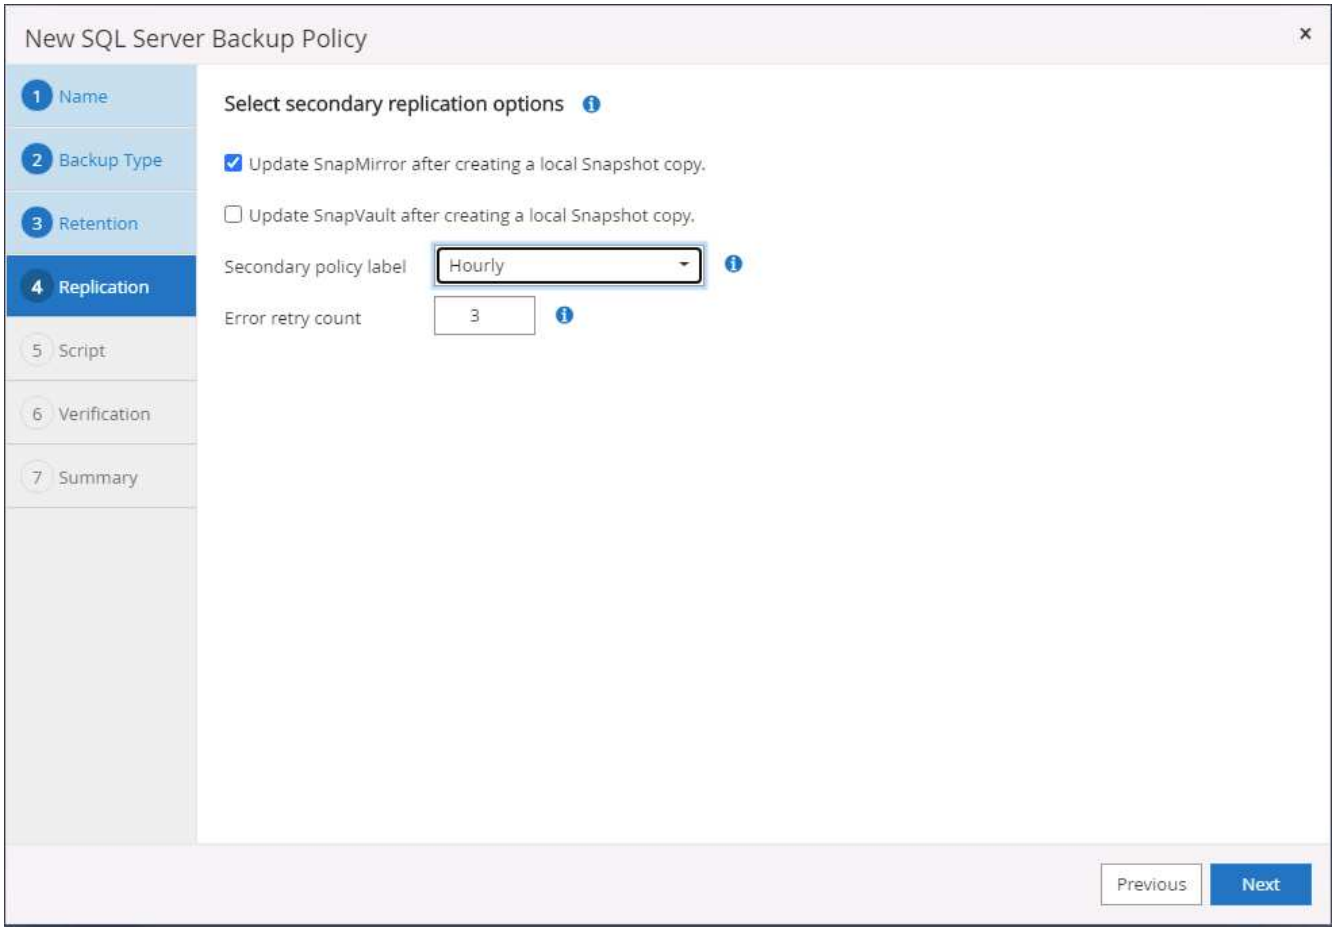

5. Specify any optional scripts to run before or after a backup job.

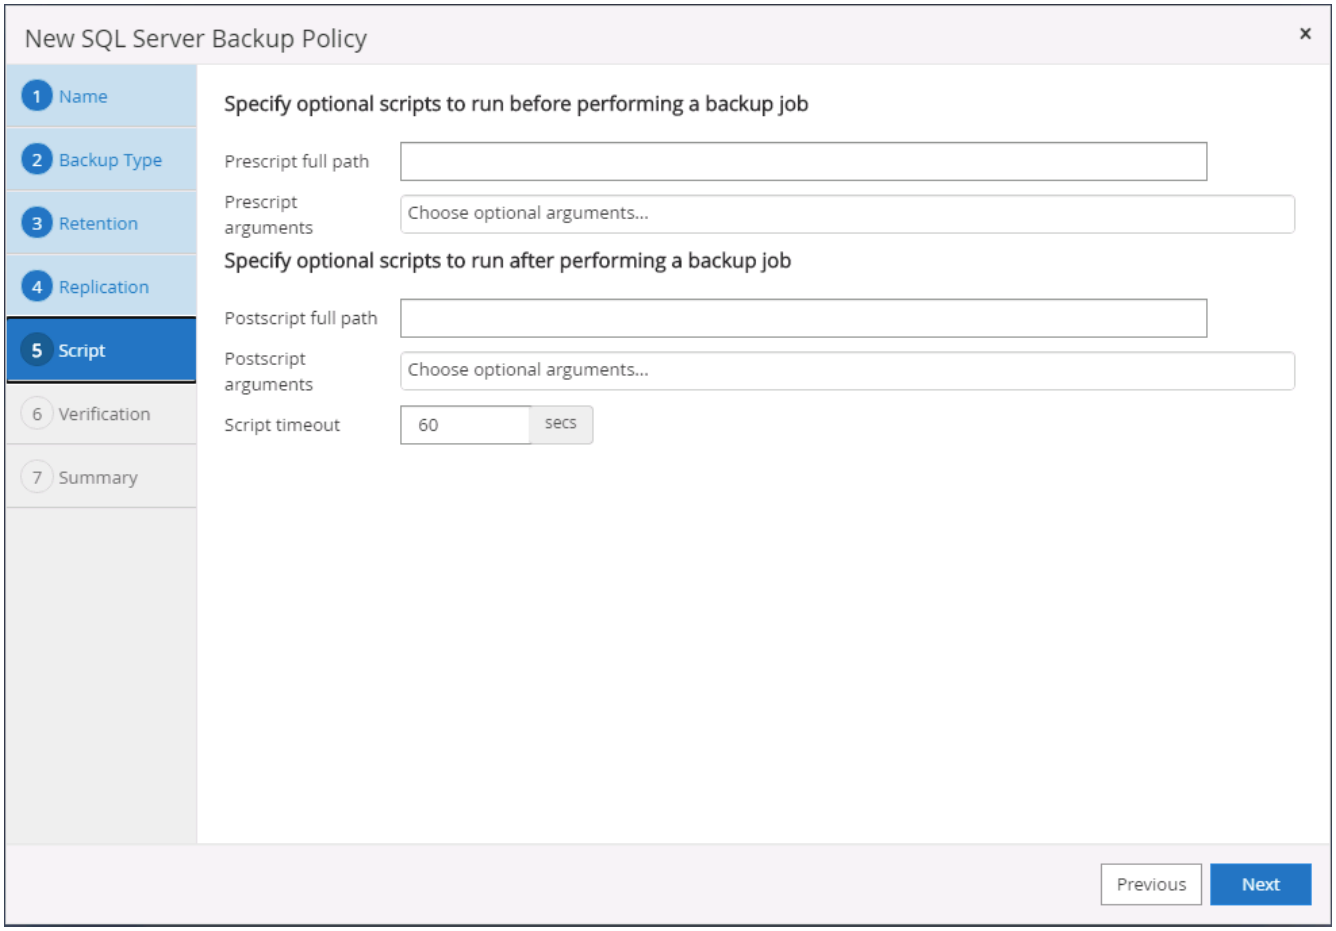
New SQL Server Backup Policy
1 Name
2 Backup Type
3 Retention
4 Replication
5 Script
6 Verification
7 Summary
Specify optional scripts to run before performing a backup job
Prescript full path
Prescript arguments
Choose optional arguments...
Specify optional scripts to run after performing a backup job
Postscript full path
Postscript arguments
Choose optional arguments...
Script timeout
60
secs
Previous
Next

6. Summary.

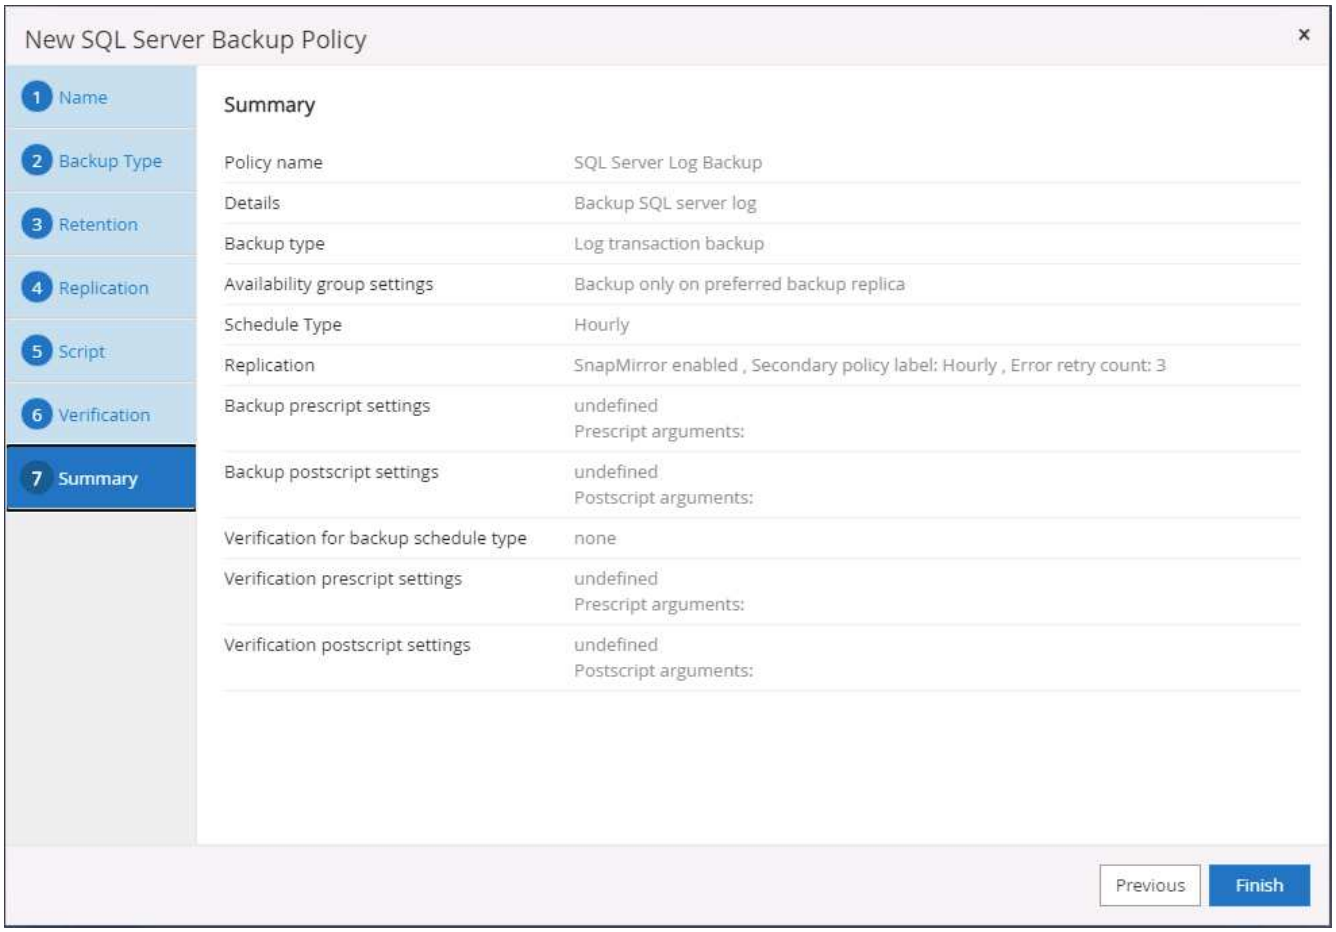

8. Implement backup policy to protect database

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

Create a resource group for full backup of Oracle

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

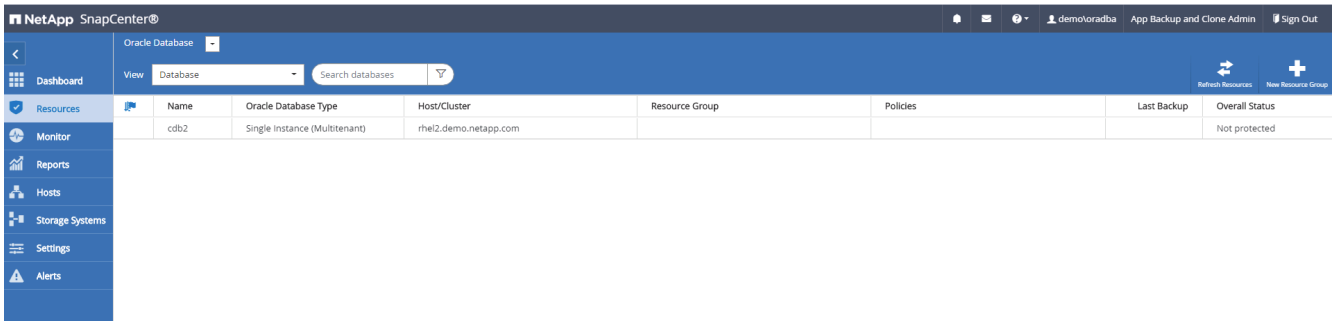

2. Provide a name and tags for the resource group. You can define a naming format for the Snapshot copy and bypass the redundant archive log destination if configured.

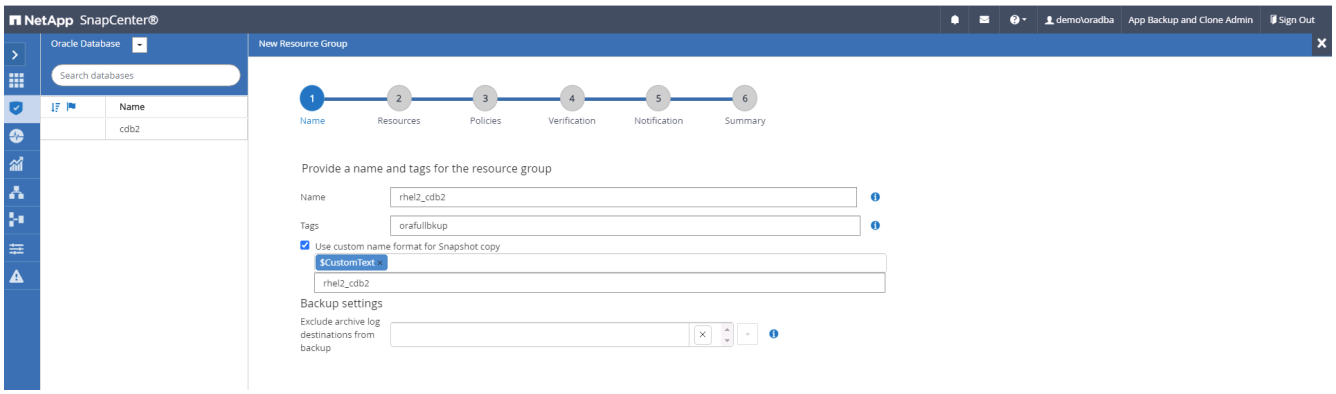

3. Add database resources to the resource group.

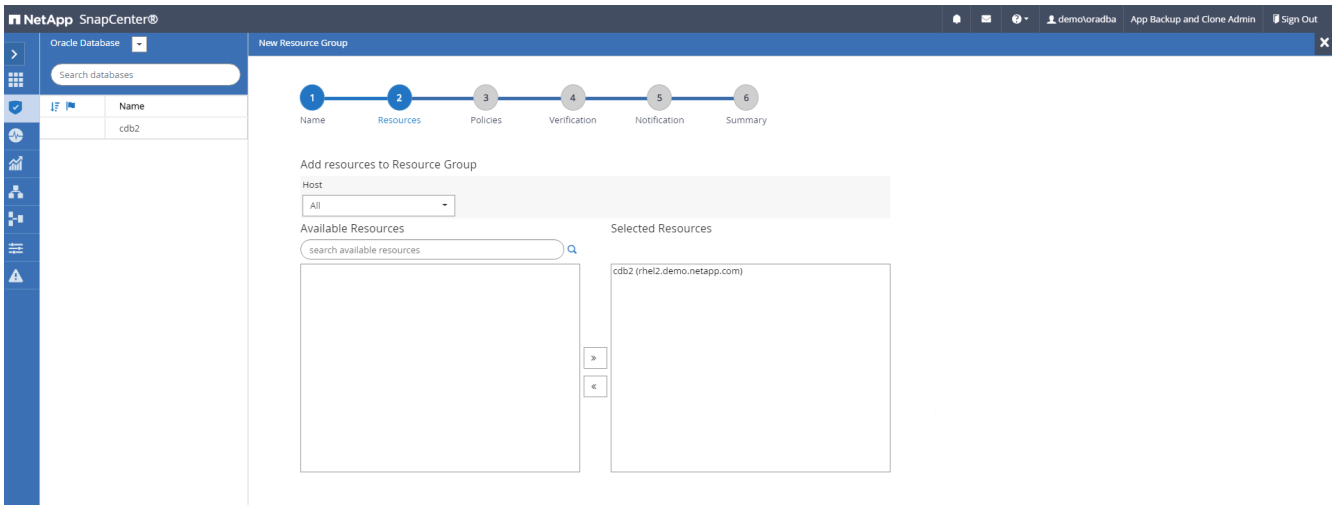

4. Select a full backup policy created in section 7 from the drop-down list.

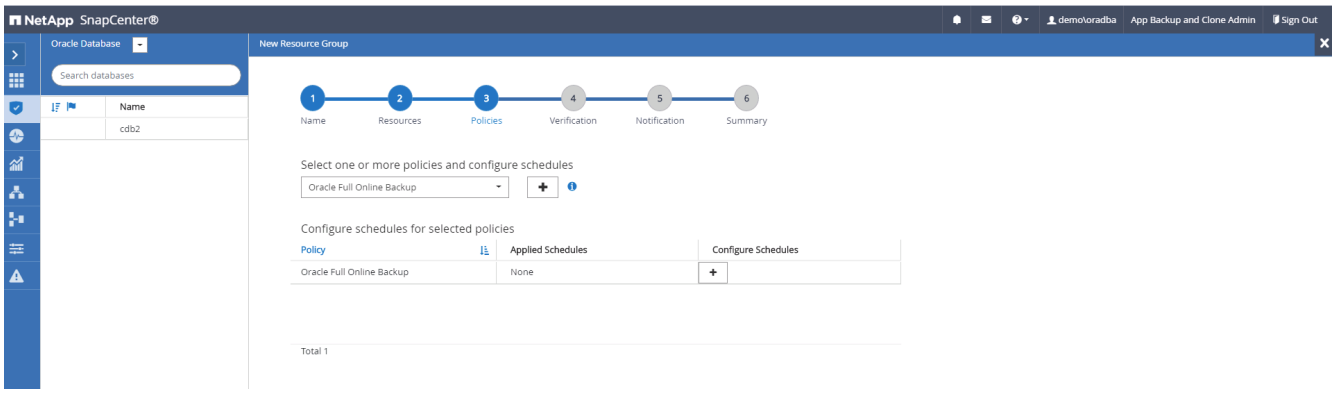

5. Click the (+) sign to configure the desired backup schedule.

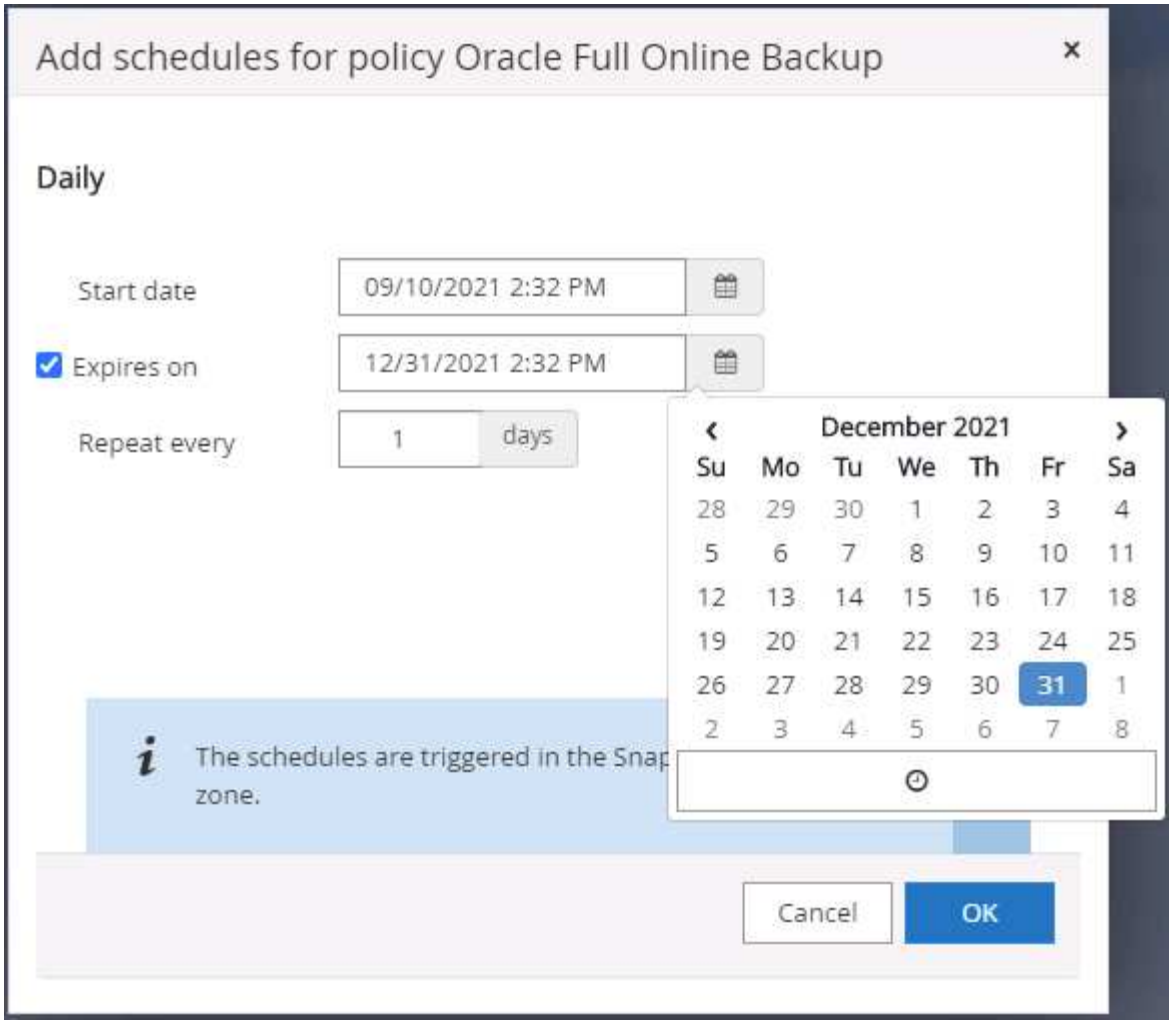

6. Click Load Locators to load the source and destination volume.

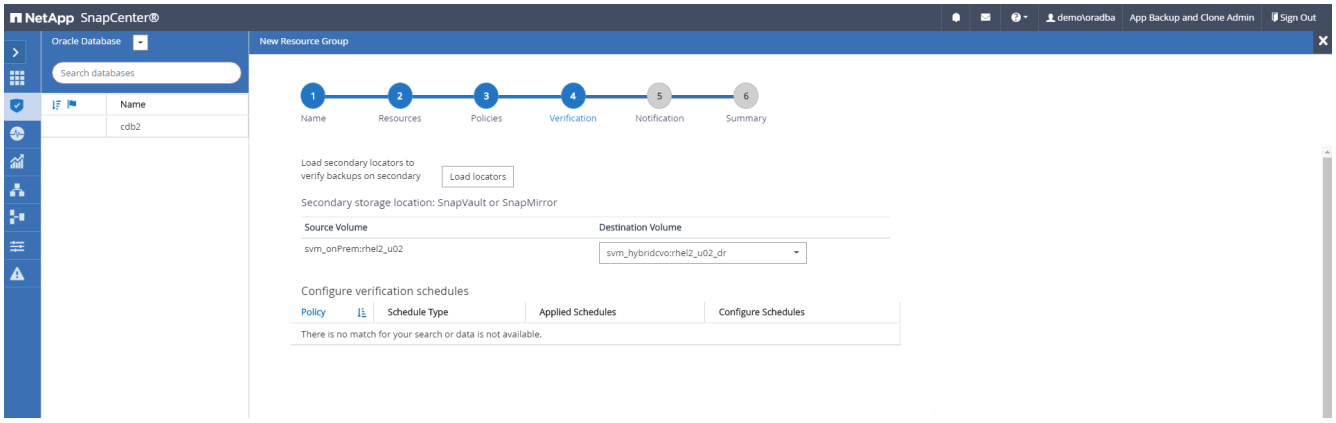

7. Configure the SMTP server for email notification if desired.

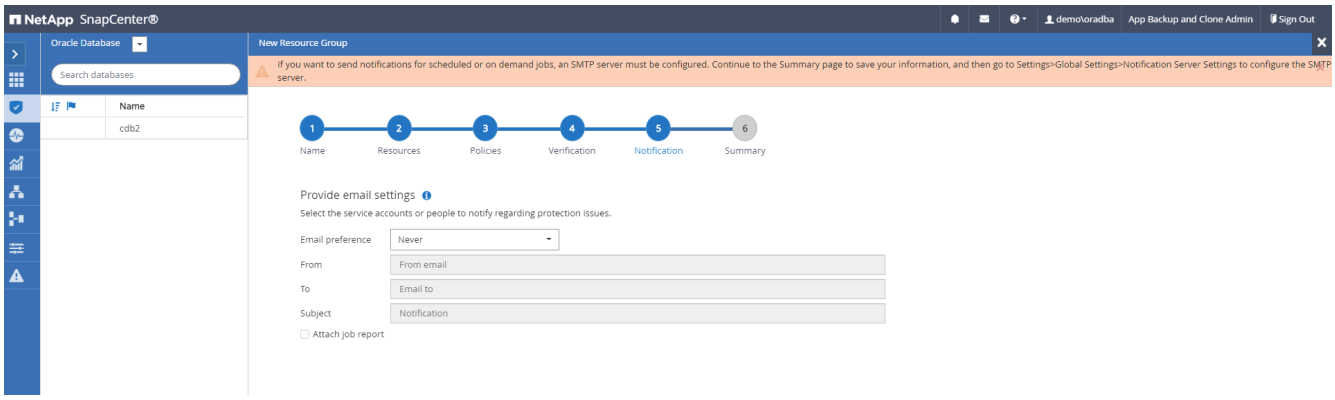

8. Summary.

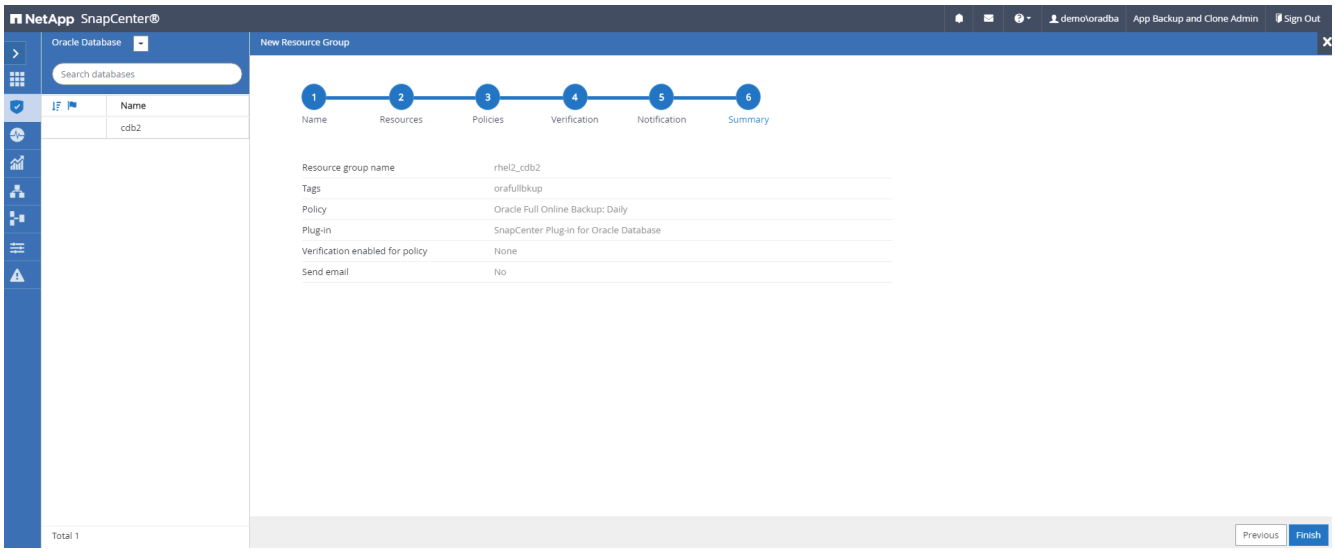

Create a resource group for log backup of Oracle

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

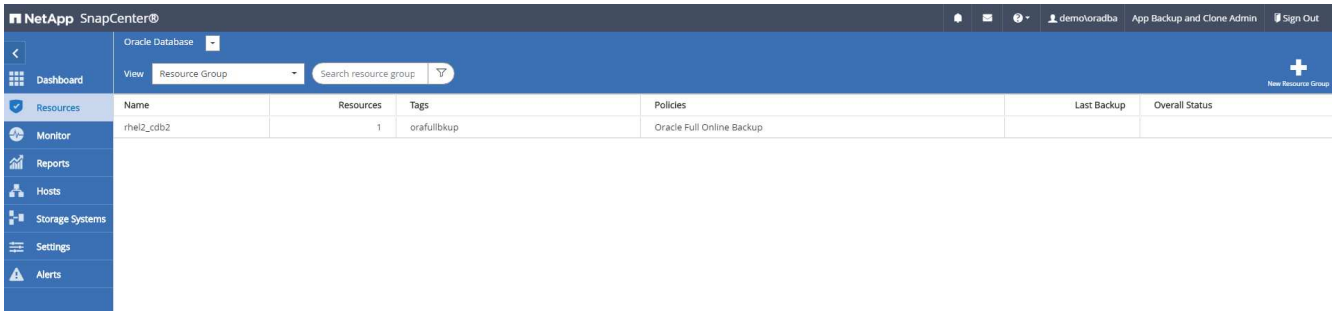

2. Provide a name and tags for the resource group. You can define a naming format for the Snapshot copy and bypass the redundant archive log destination if configured.

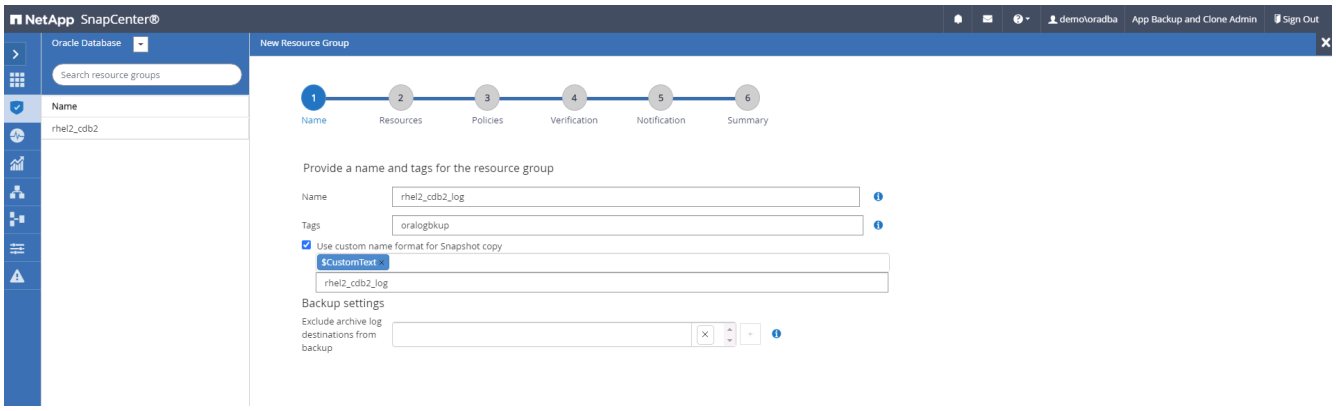

3. Add database resources to the resource group.

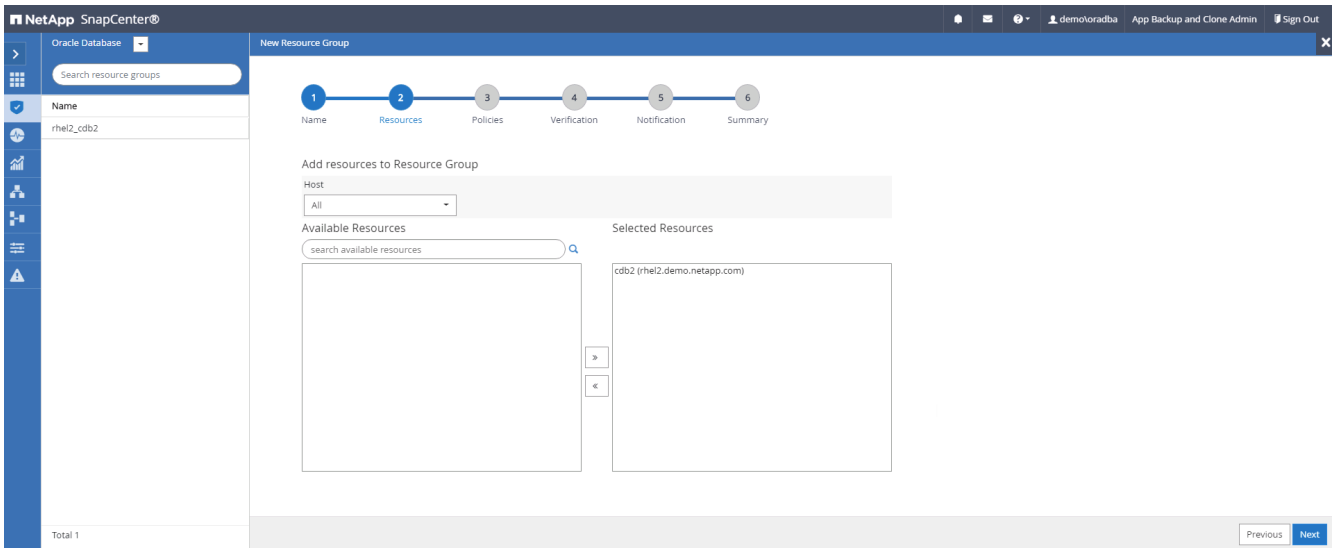

4. Select a log backup policy created in section 7 from the drop-down list.

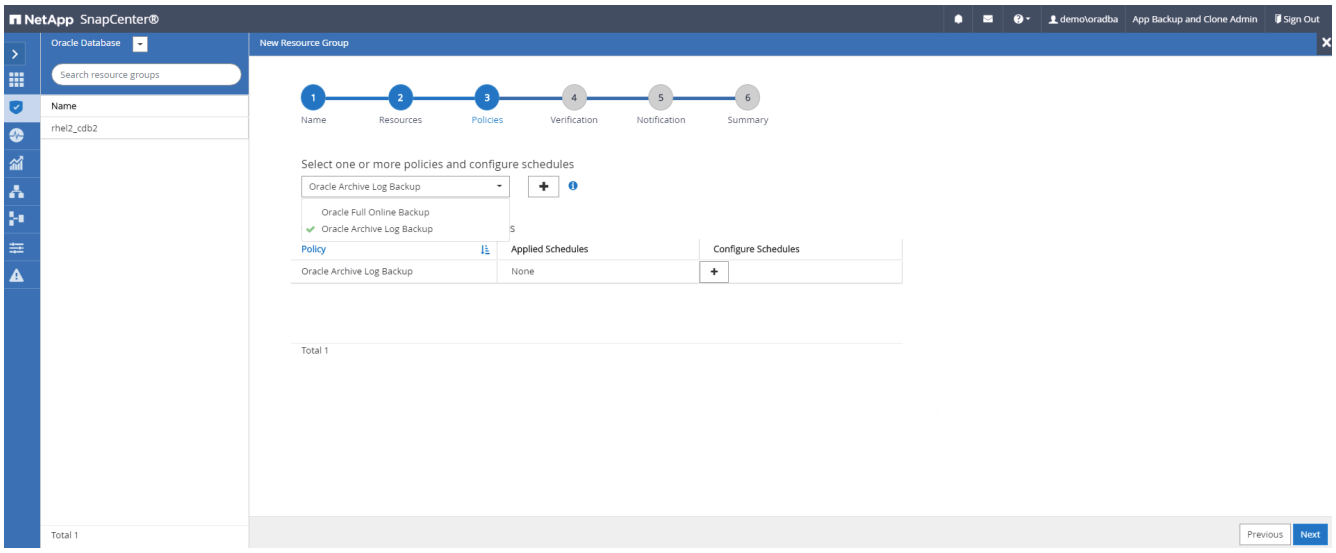

5. Click on the (+) sign to configure the desired backup schedule.

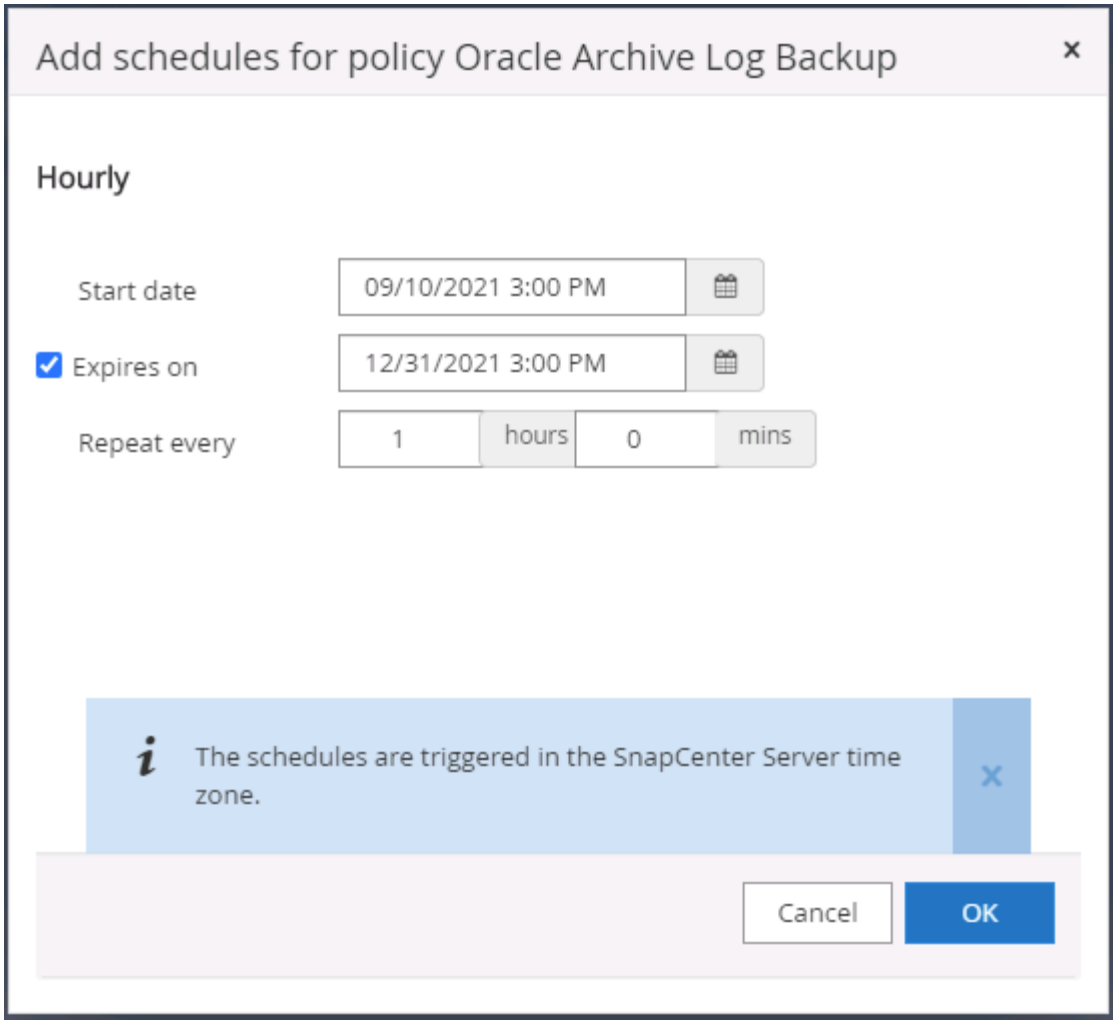

6. If backup verification is configured, it displays here.

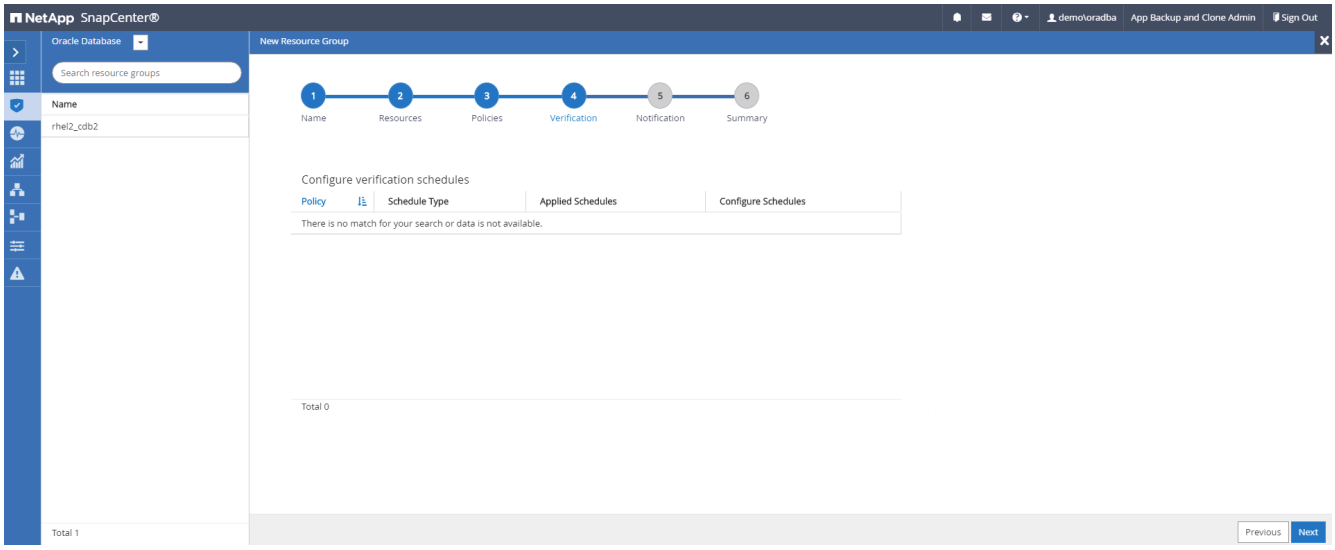

7. Configure an SMTP server for email notification if desired.

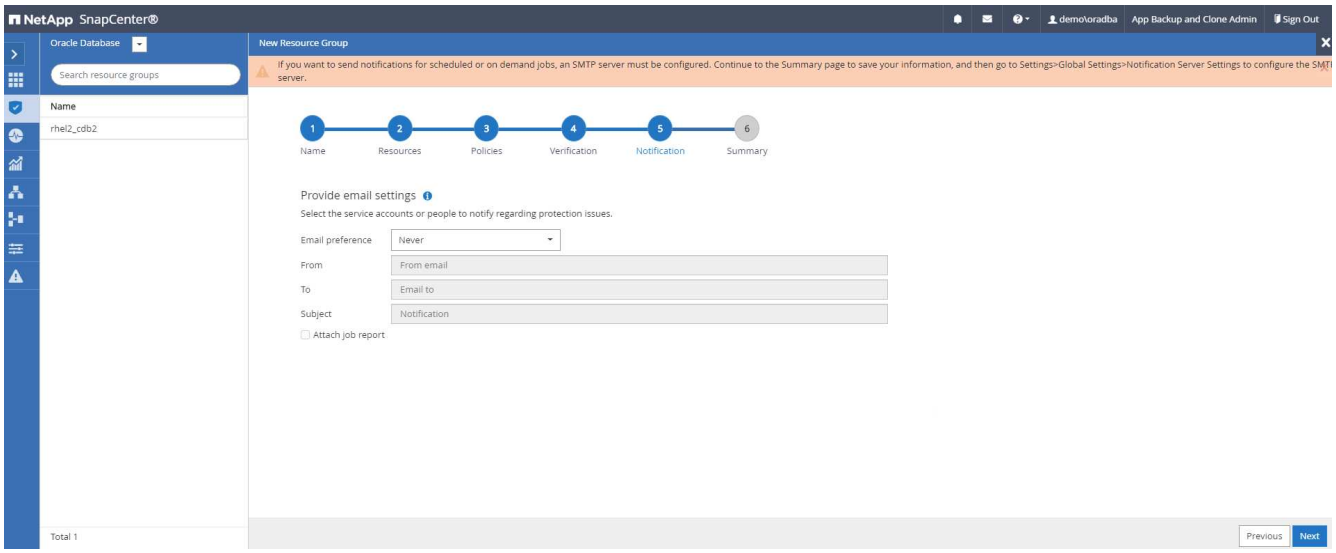

8. Summary.

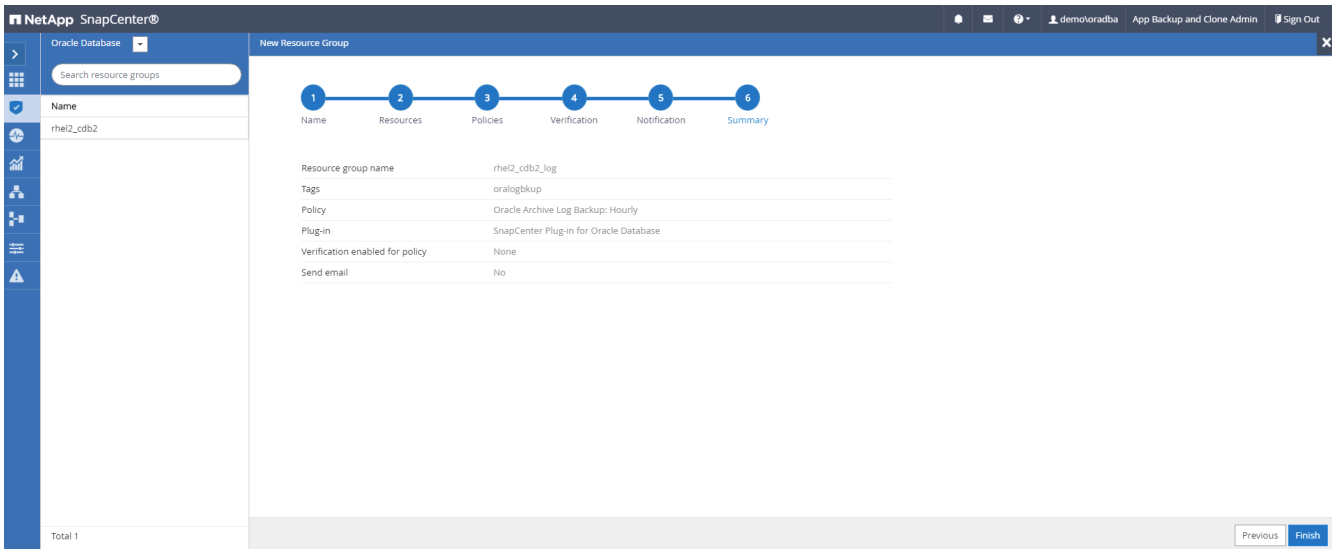

Create a resource group for full backup of SQL Server

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

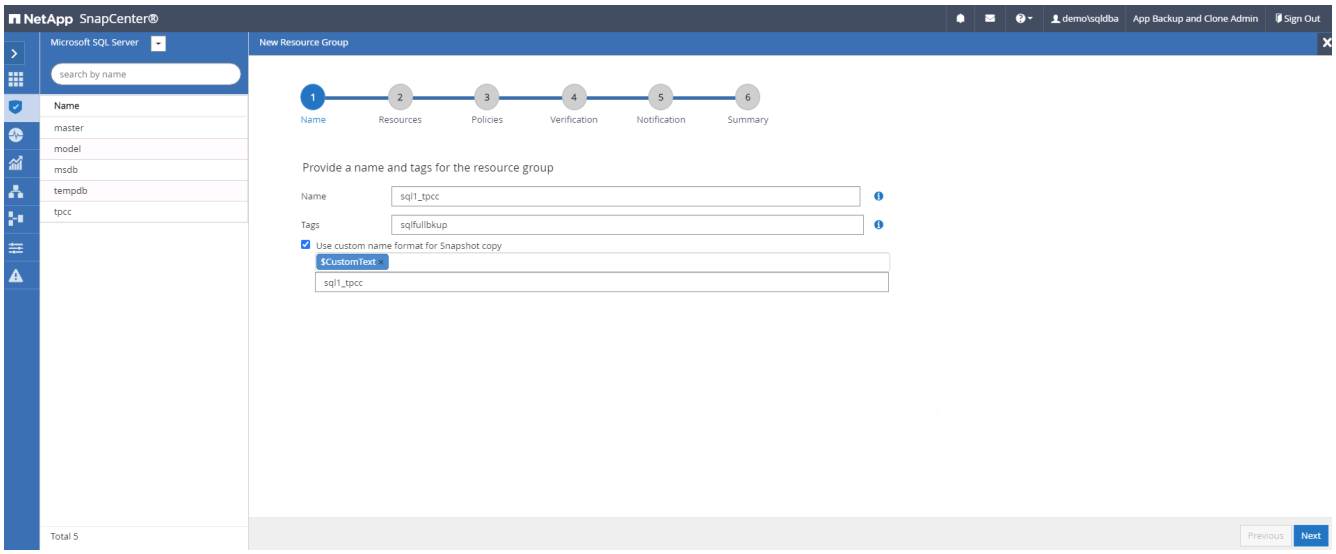

2. Select the database resources to be backed up.

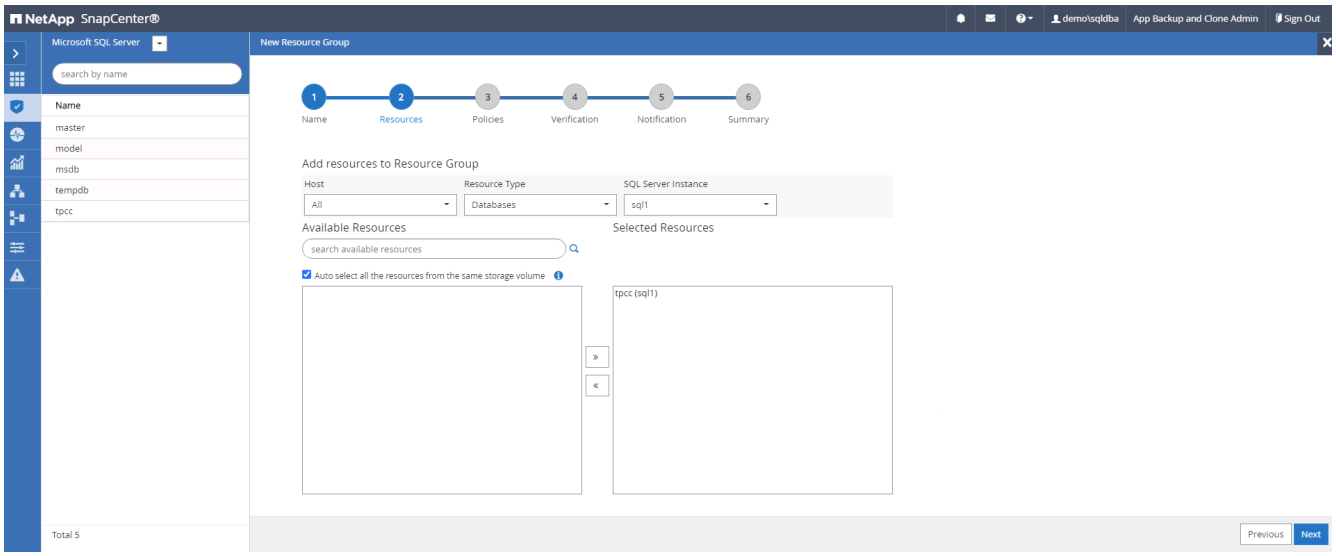

3. Select a full SQL backup policy created in section 7.

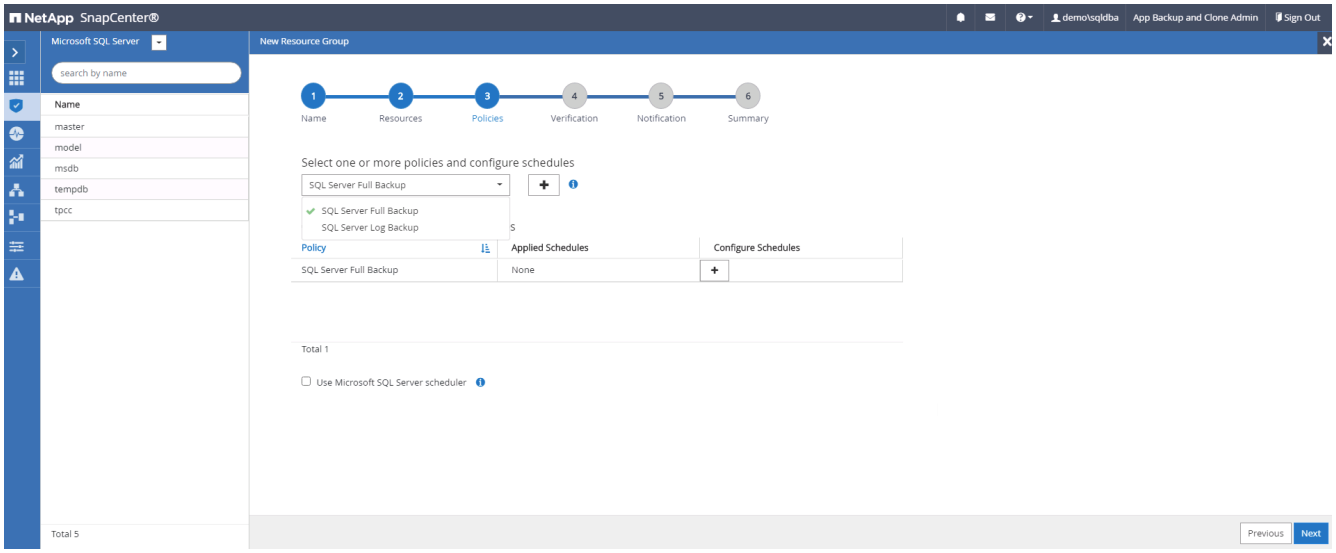

4. Add exact timing for backups as well as the frequency.

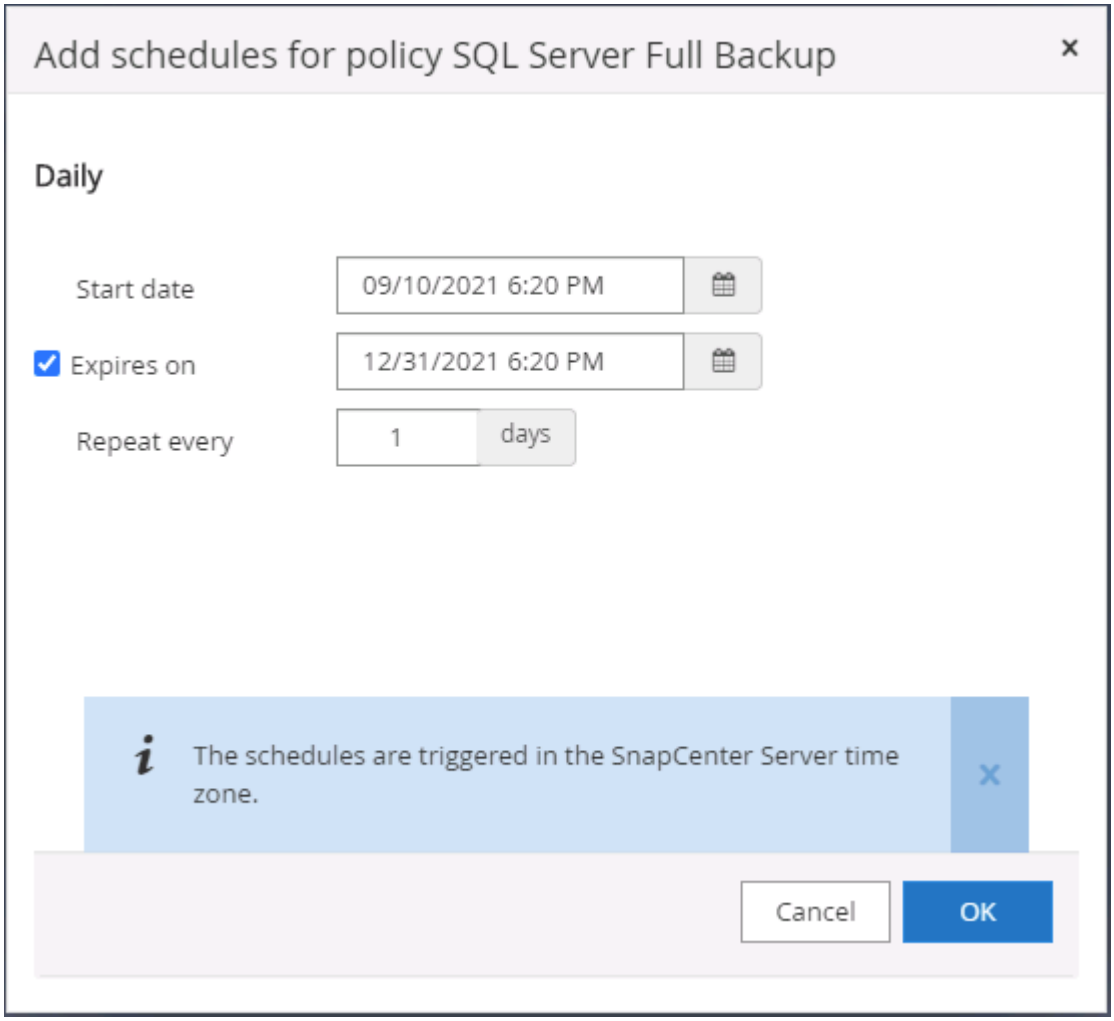

5. Choose the verification server for the backup on secondary if backup verification is to be performed. Click Load Locator to populate the secondary storage location.

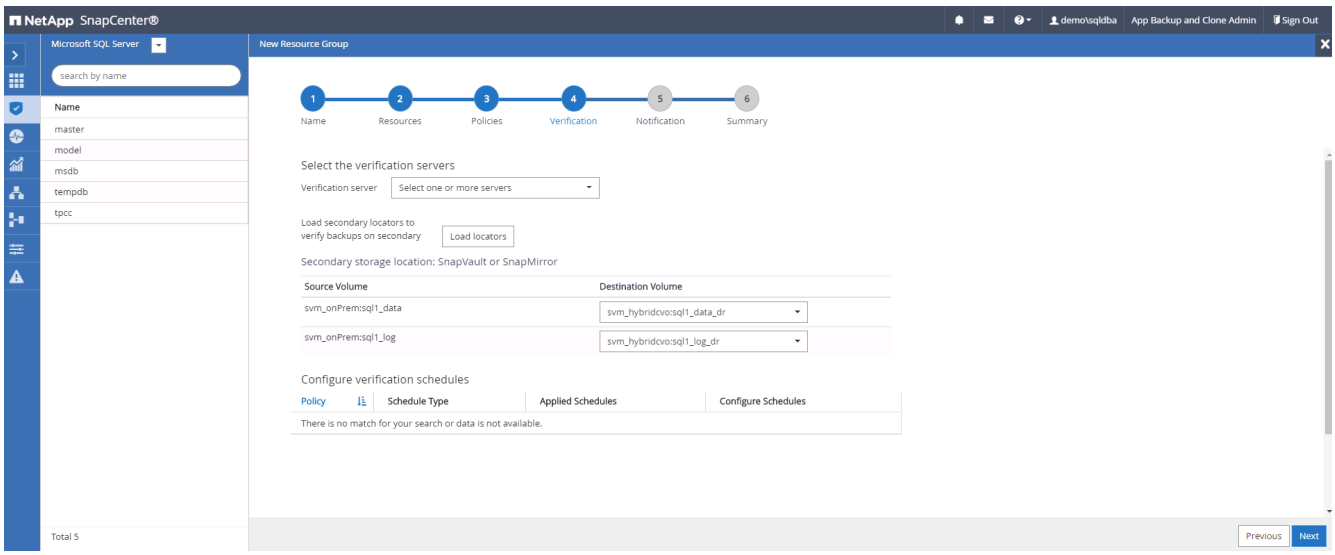

6. Configure the SMTP server for email notification if desired.

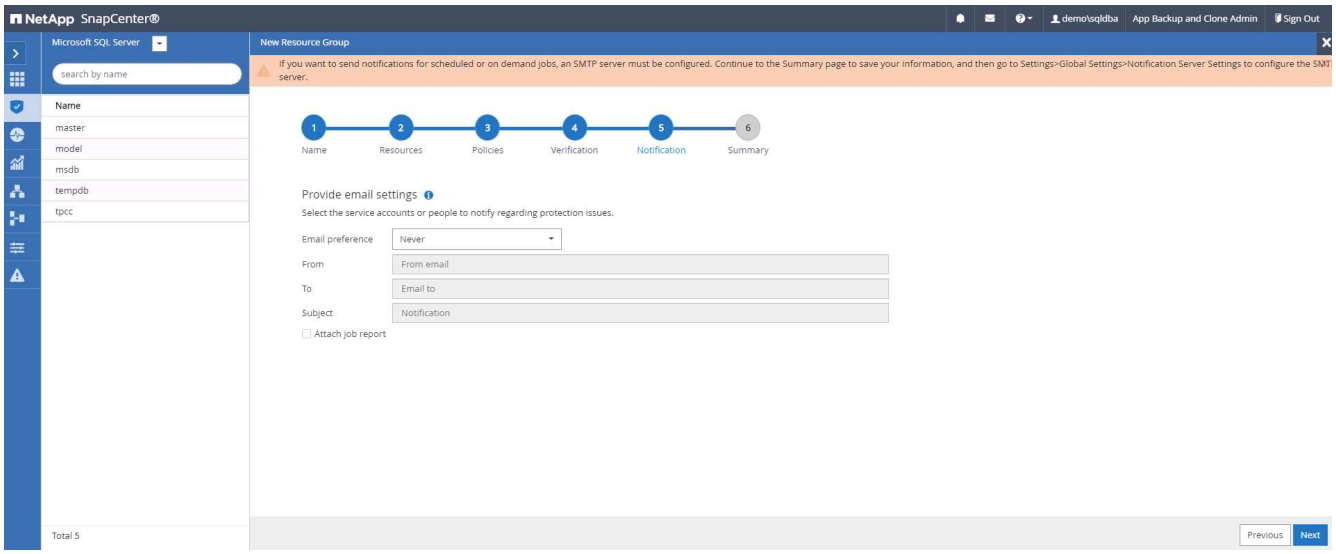

7. Summary.

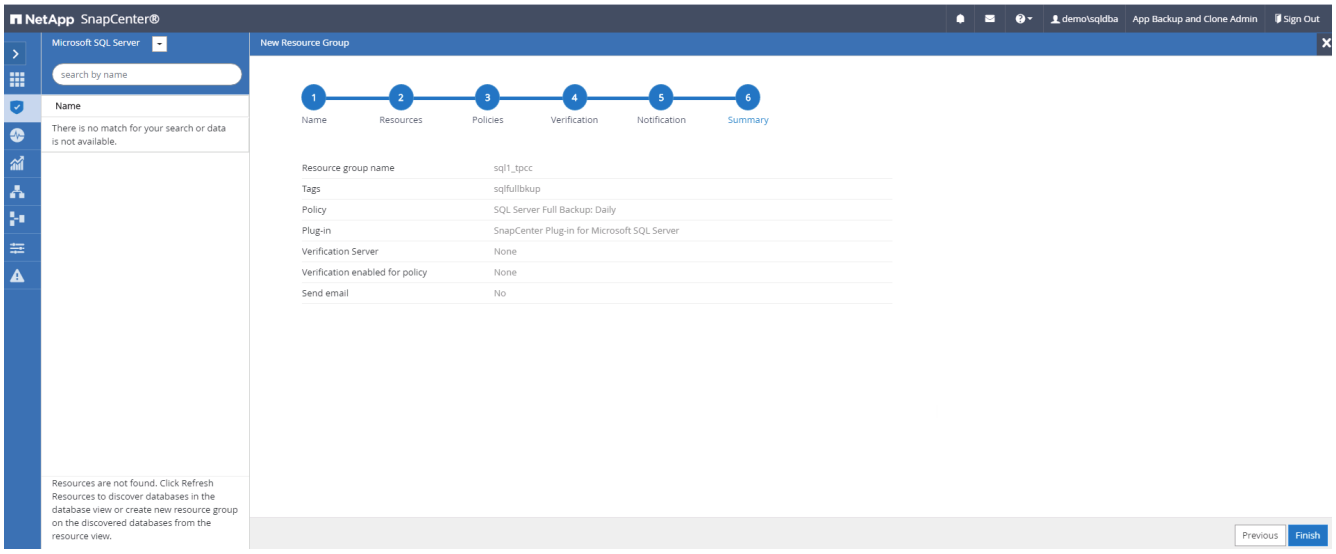

Create a resource group for log backup of SQL Server

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.

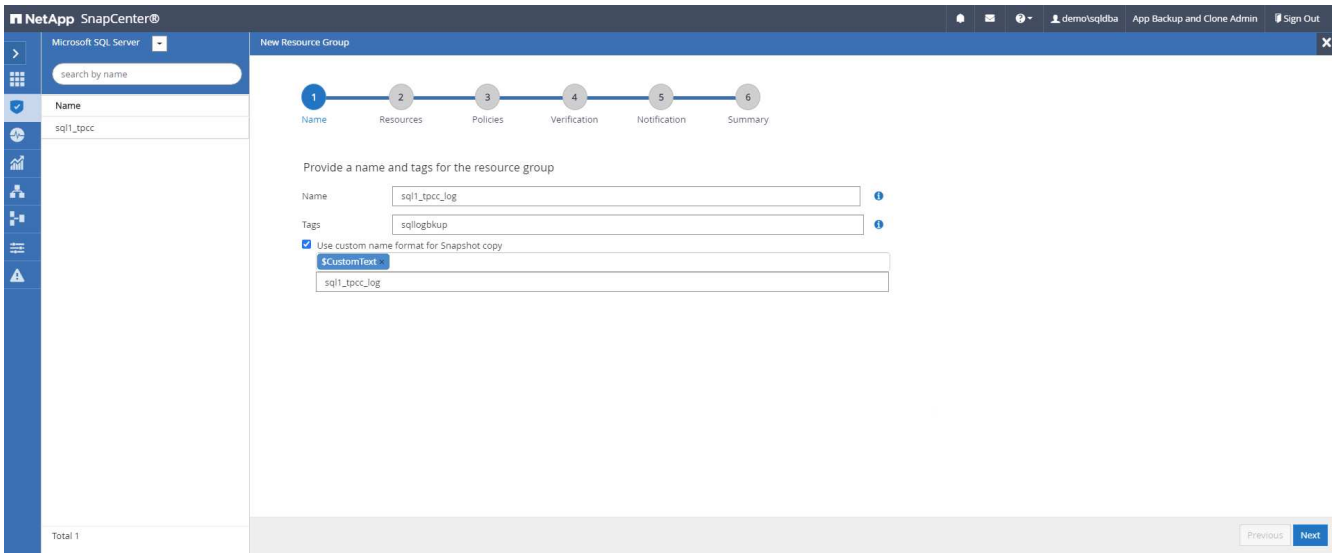

2. Select the database resources to be backed up.

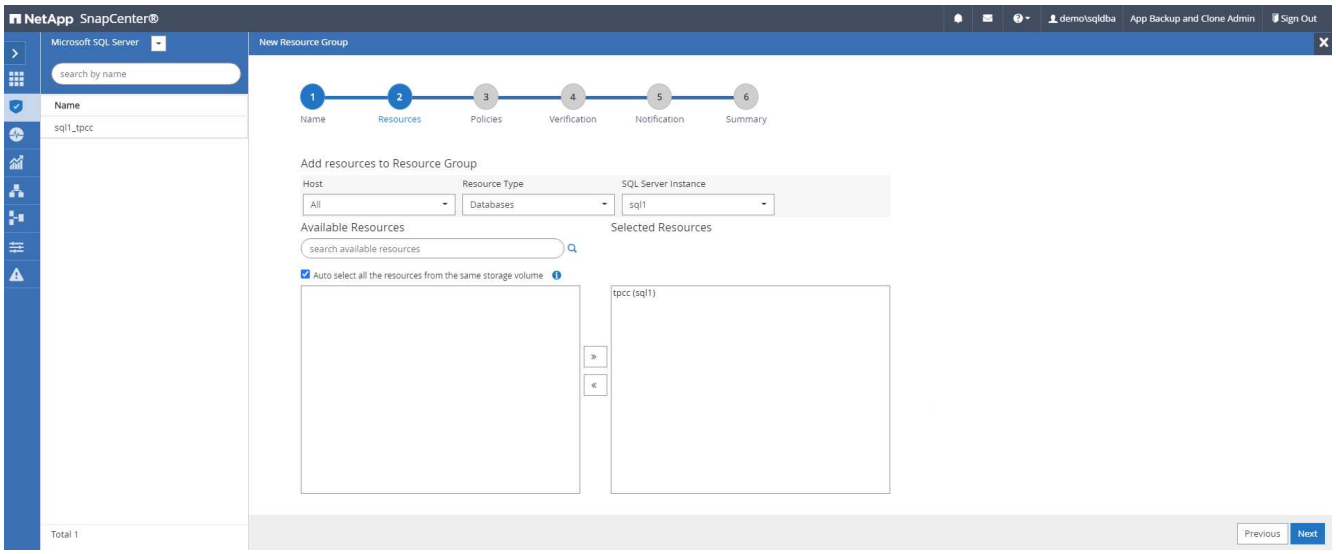

3. Select a SQL log backup policy created in section 7.

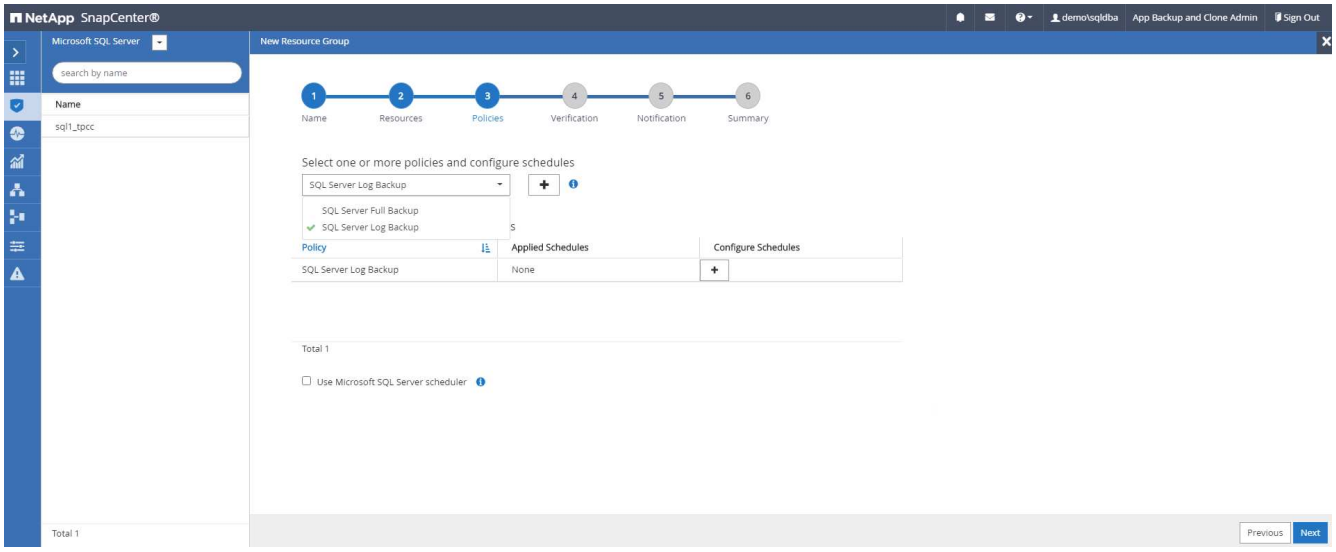

4. Add exact timing for the backup as well as the frequency.

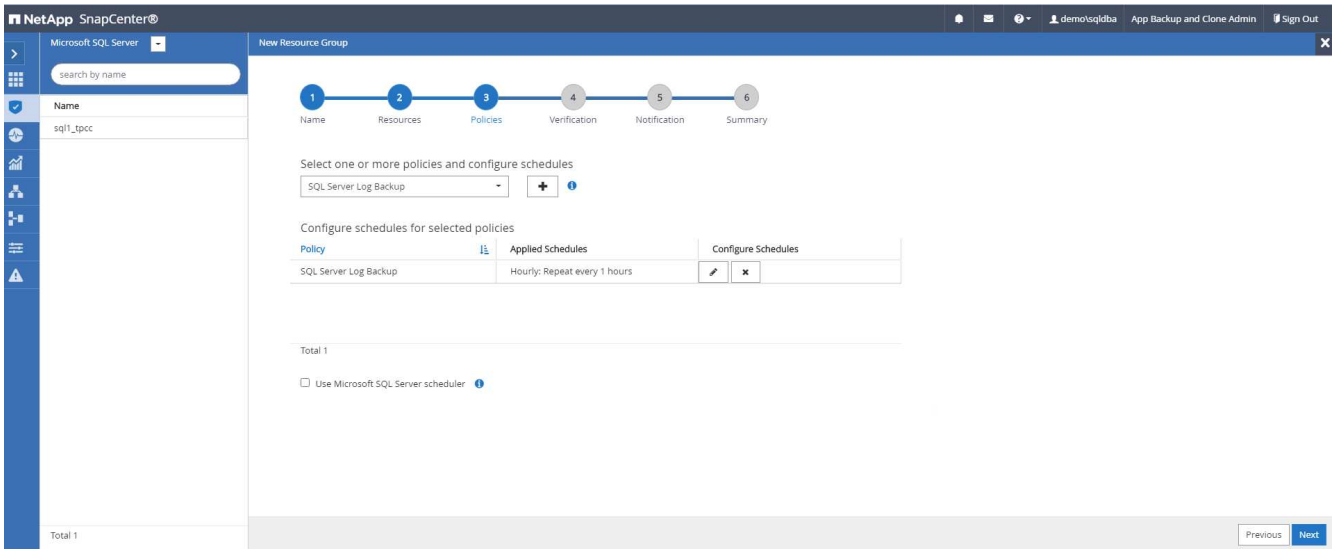

5. Choose the verification server for the backup on secondary if backup verification is to be performed. Click the Load Locator to populate the secondary storage location.

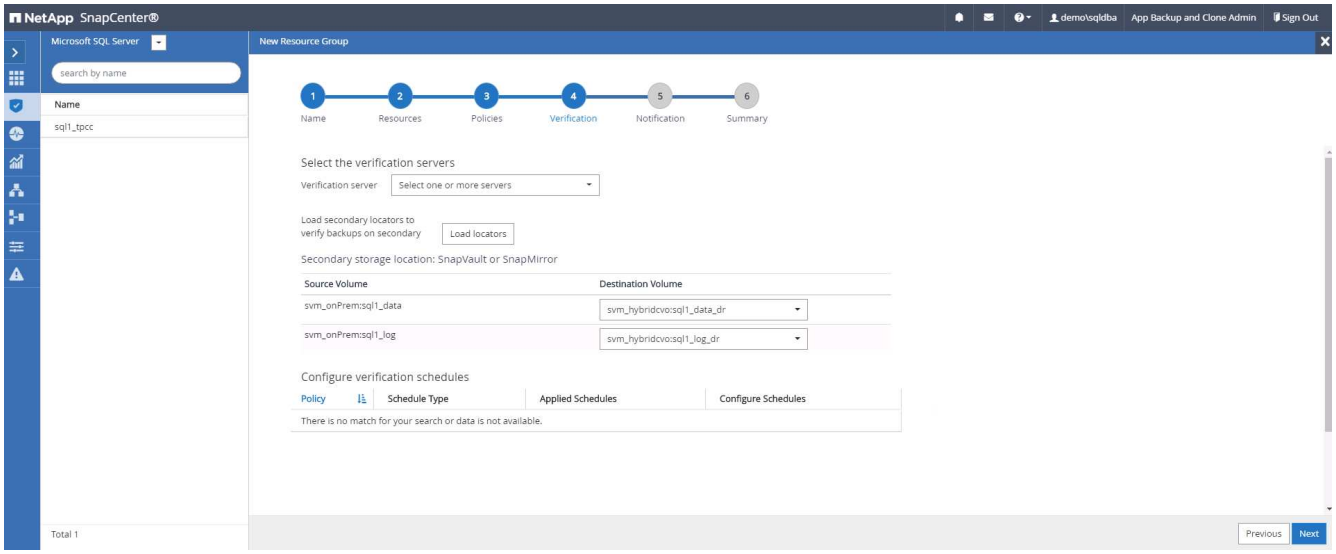

6. Configure the SMTP server for email notification if desired.

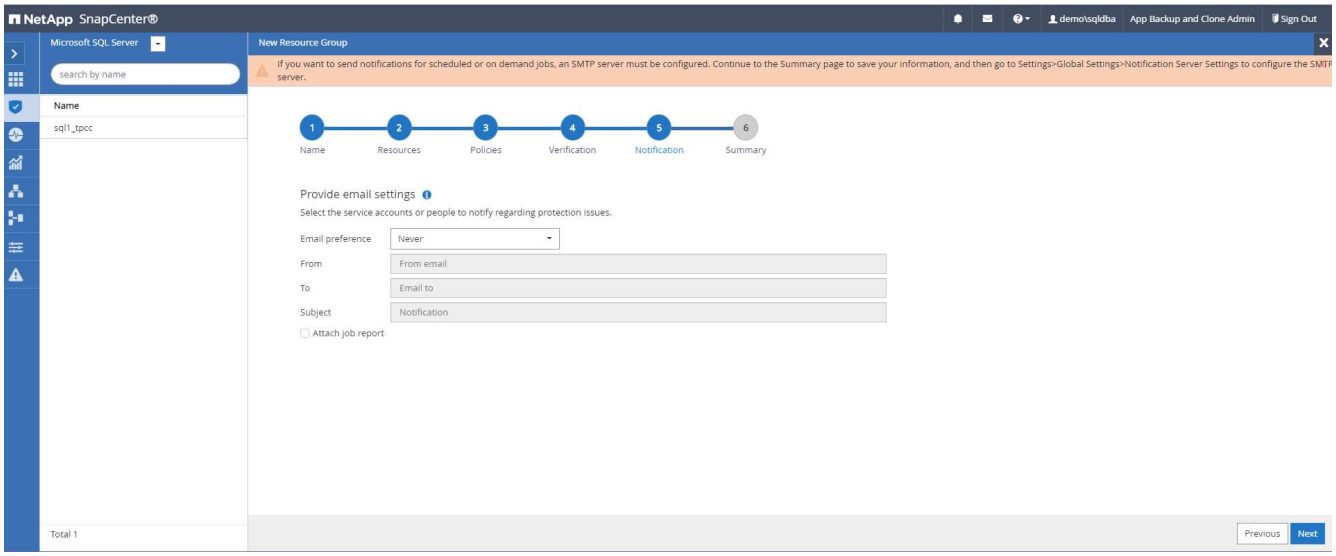

7. Summary.

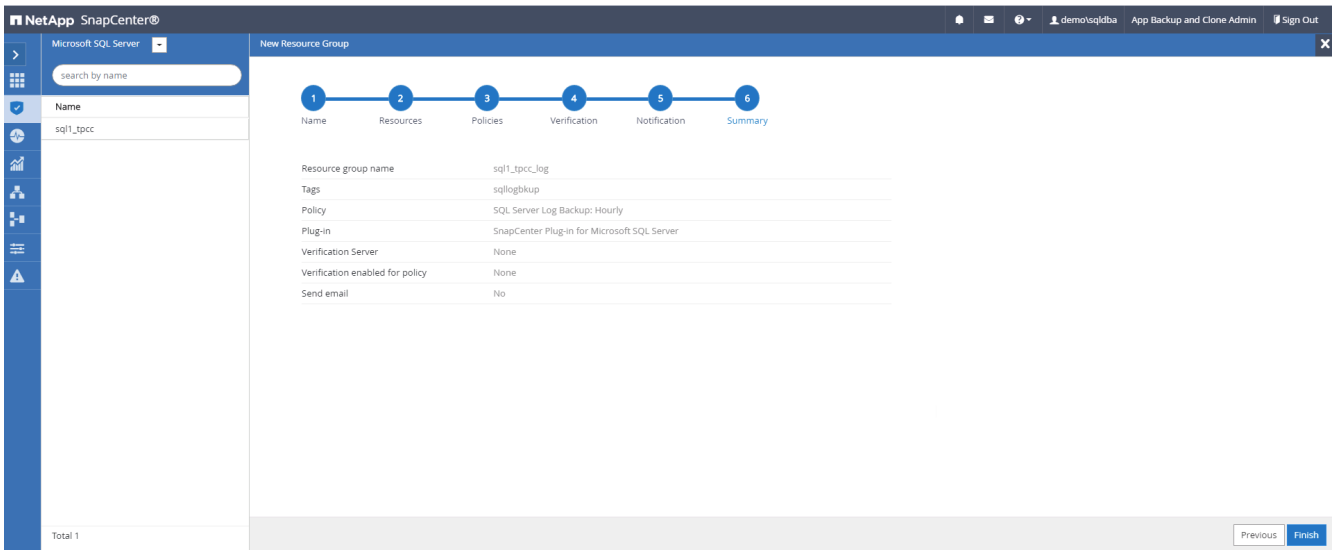

9. Validate backup

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.

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.

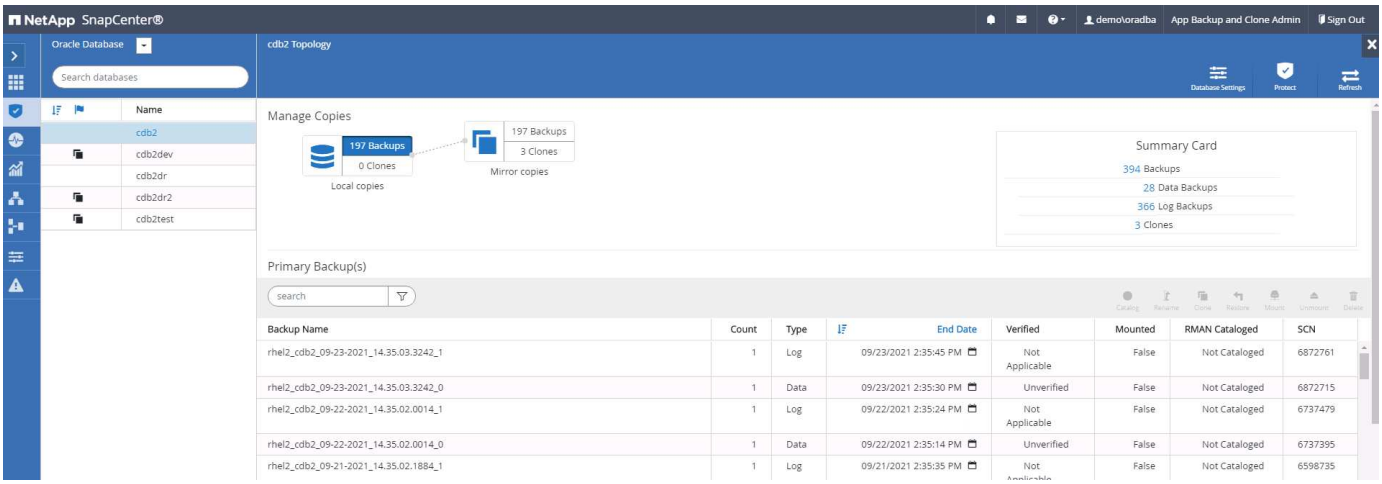

At this point, database backup copies in the cloud are ready to clone to run dev/test processes or for disaster recovery in the event of a primary failure.

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.

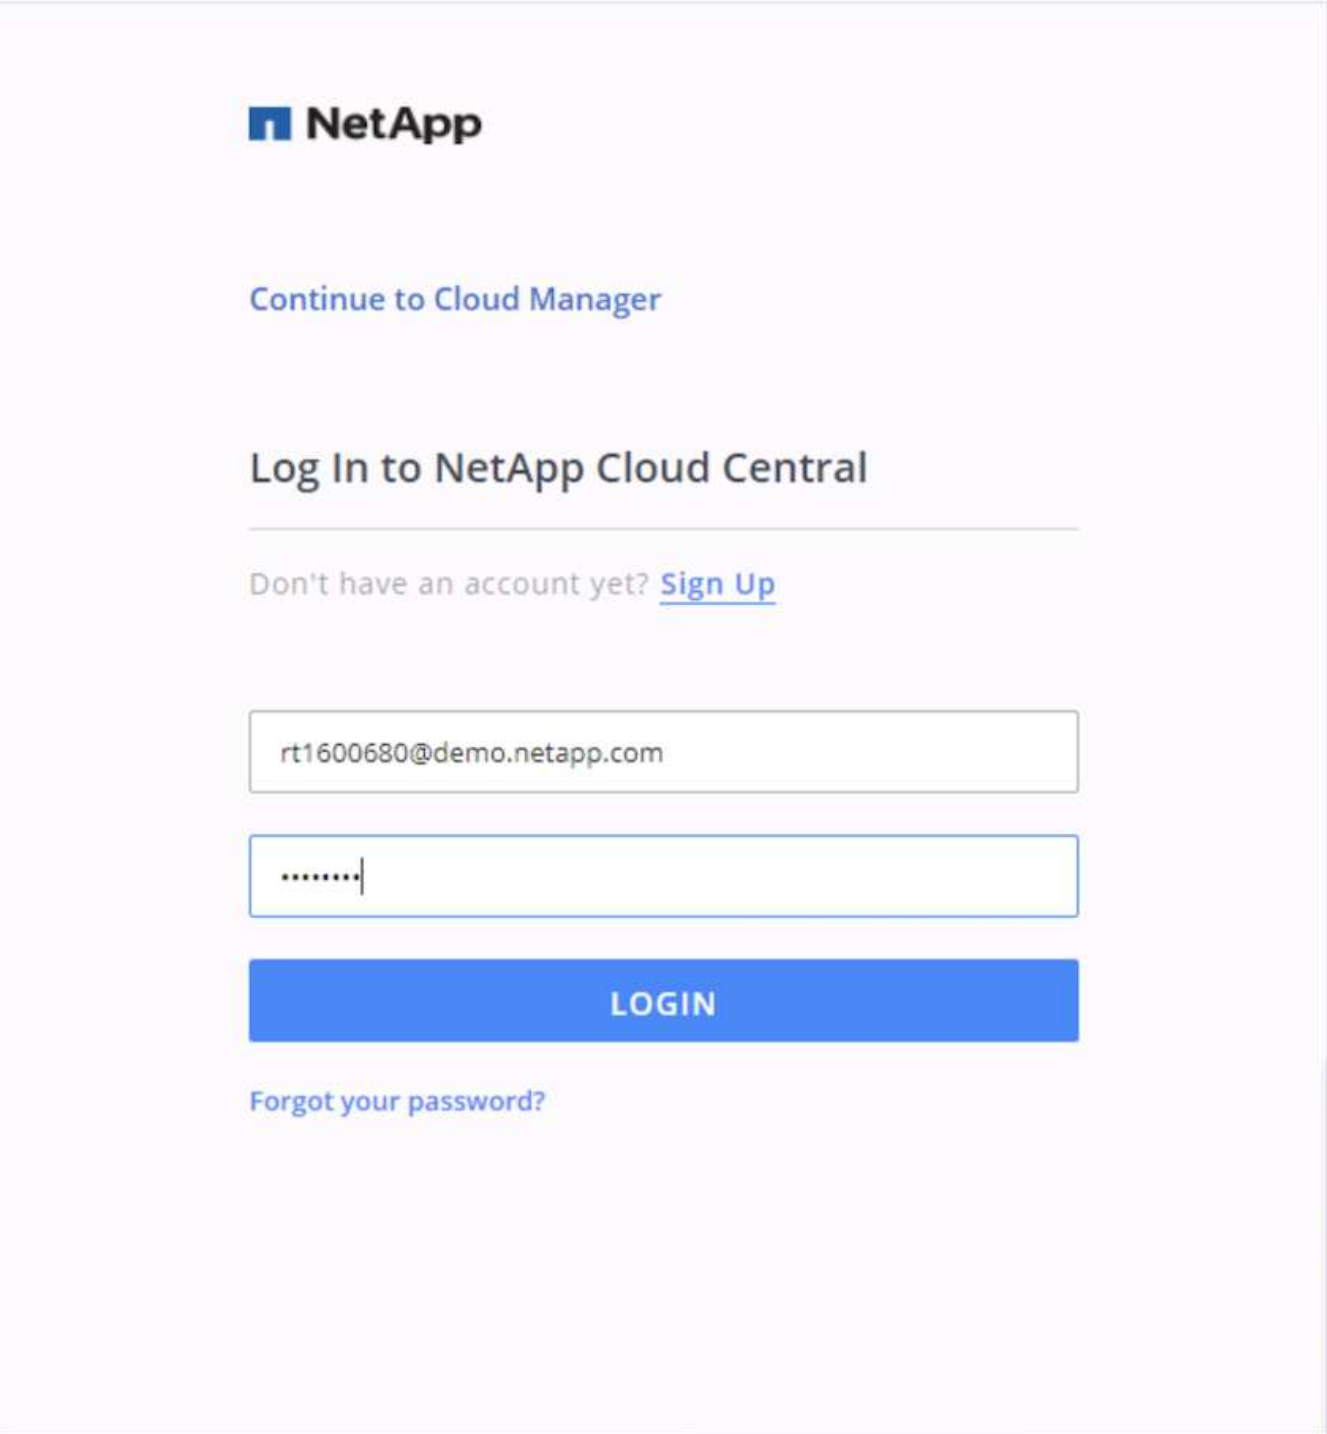

2. After you log in, you should be taken to the Canvas.

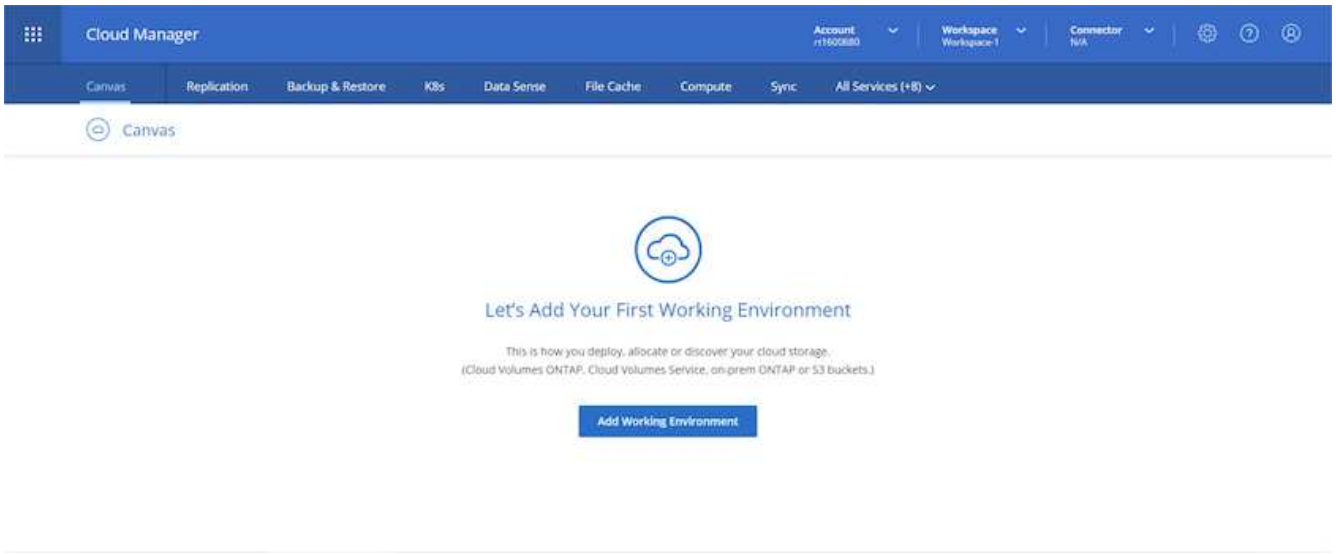

3. Click "Add Working Environment" and choose Cloud Volumes ONTAP in AWS. Here, you also choose whether you want to deploy a single node system or a high availability pair. I have chosen to deploy a high availability pair.

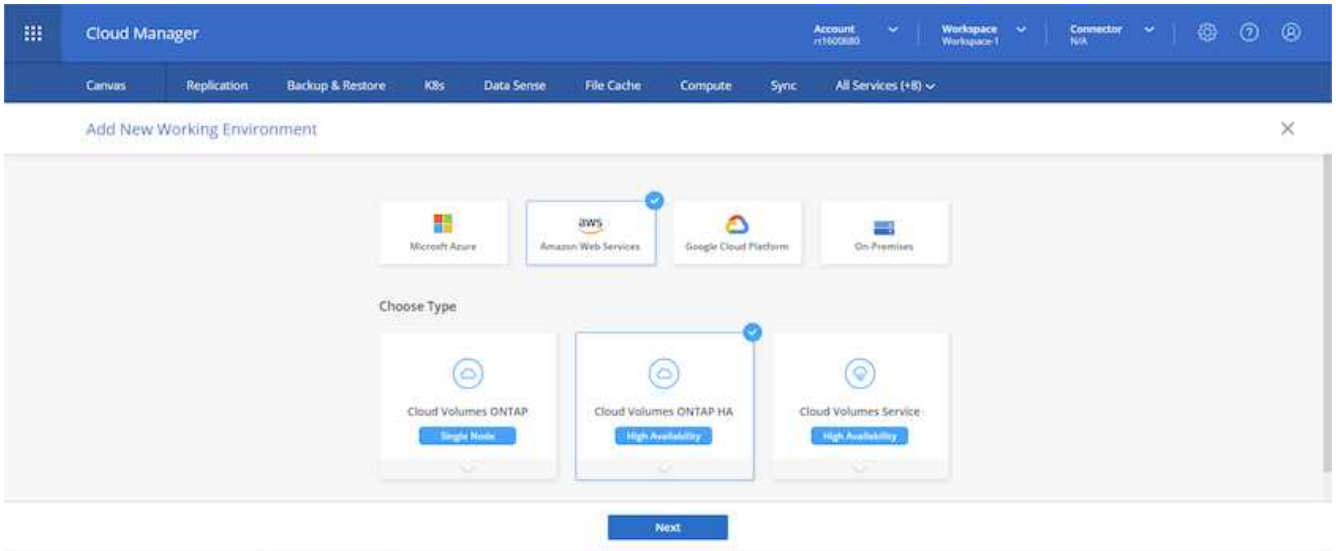

4. If no connector has been created, a pop-up appears asking you to create a connector.

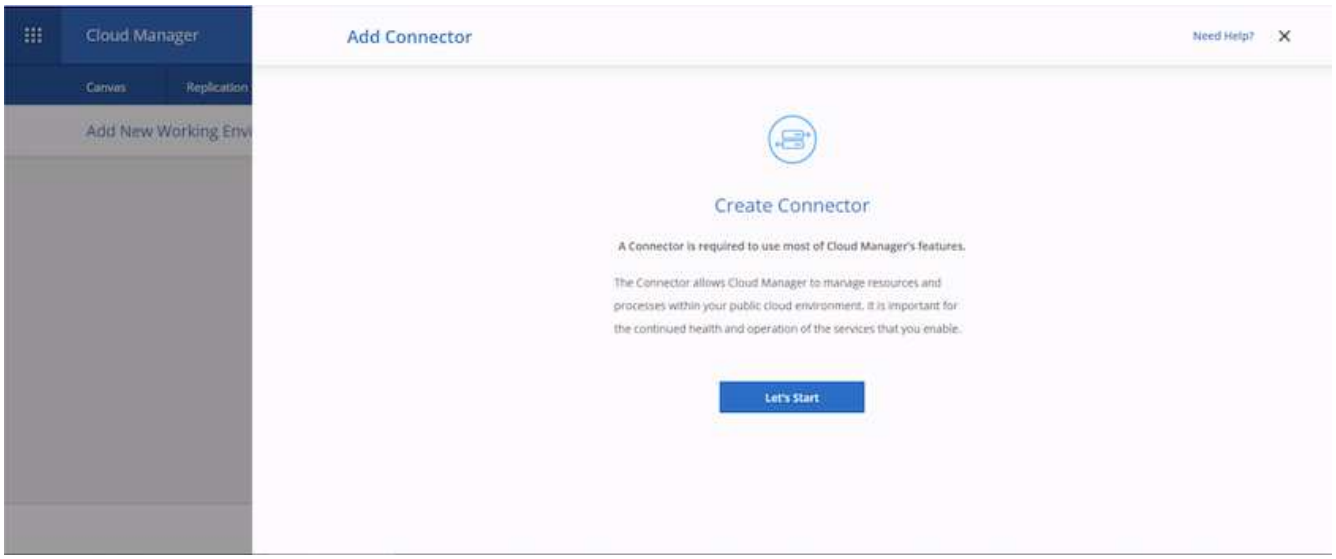

5. Click Lets Start, and then choose AWS.

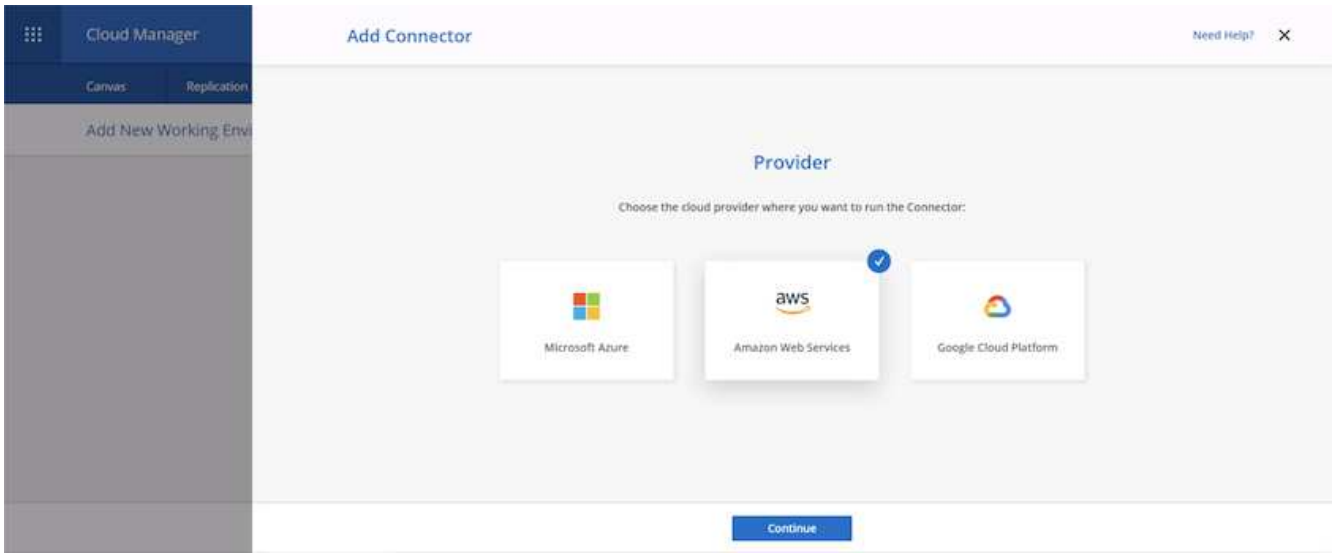

6. Enter your secret key and access key. Make sure that your user has the correct permissions outlined on the NetApp policies page.

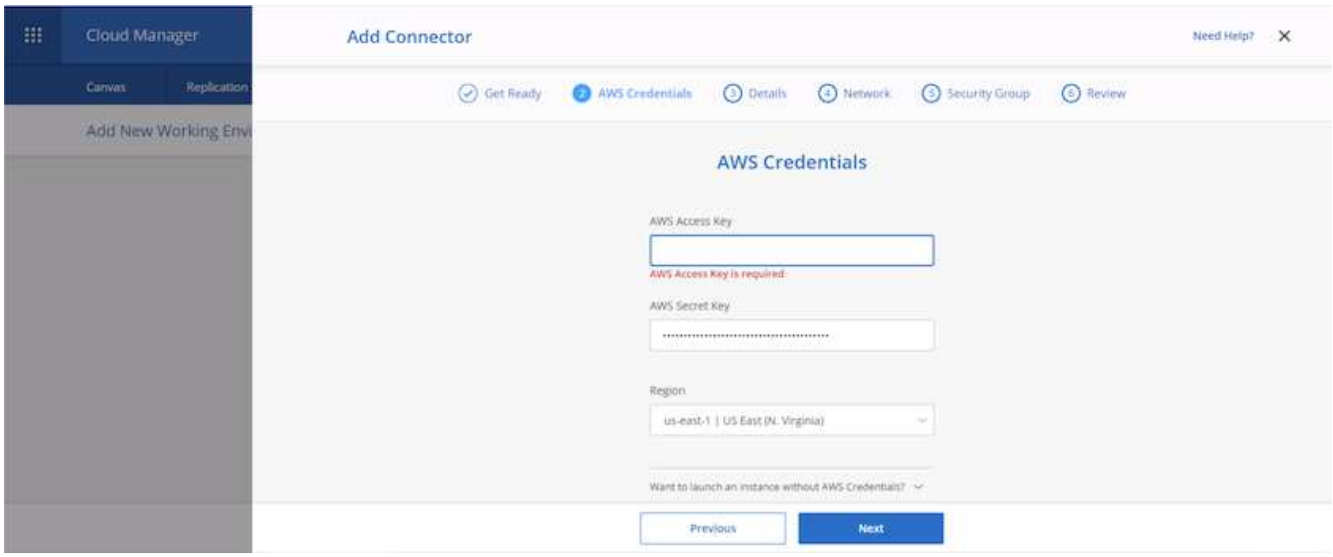

7. Give the connector a name and either use a predefined role as described on the NetApp policies page or ask Cloud Manager to create the role for you.

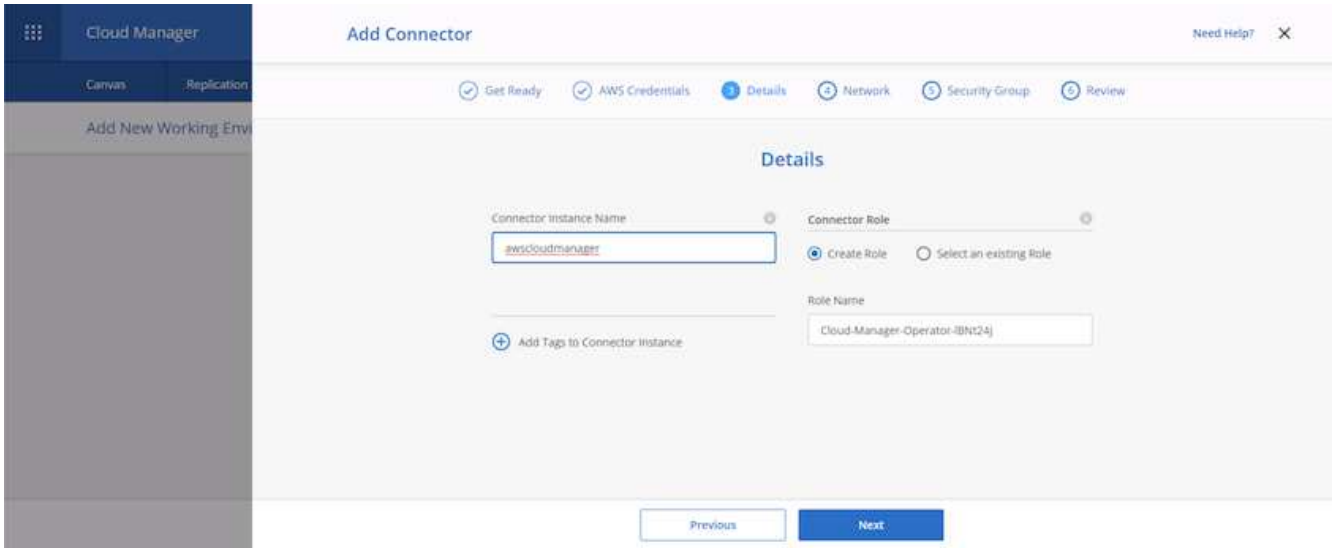

8. Give the networking information needed to deploy the connector. Verify that outbound internet access is enabled by:
   a. Giving the connector a public IP address
   b. Giving the connector a proxy to work through
   c. Giving the connector a route to the public internet through an Internet Gateway

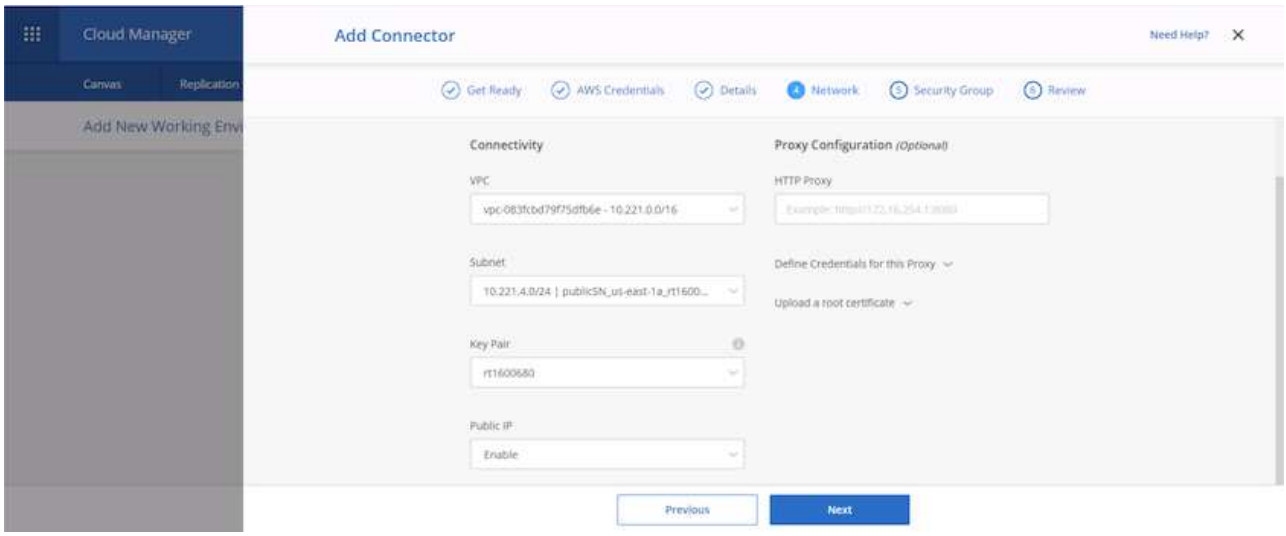

9. Provide communication with the connector via SSH, HTTP, and HTTPs by either providing a security group or creating a new security group. I have enabled access to the connector from my IP address only.

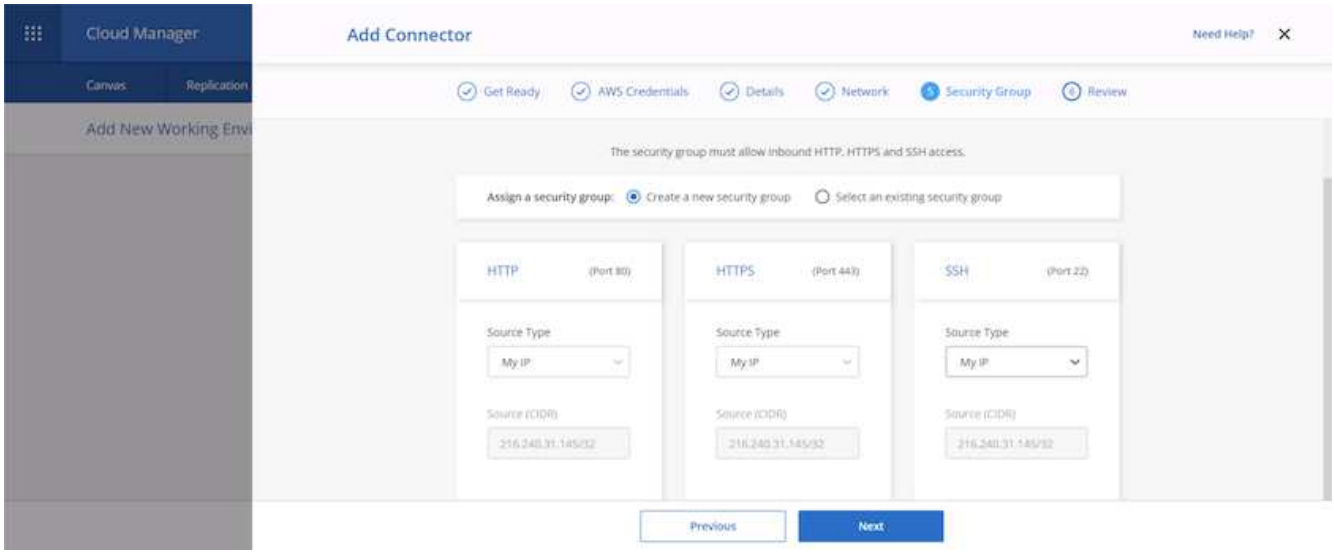

10. Review the information on the summary page and click Add to deploy the connector.

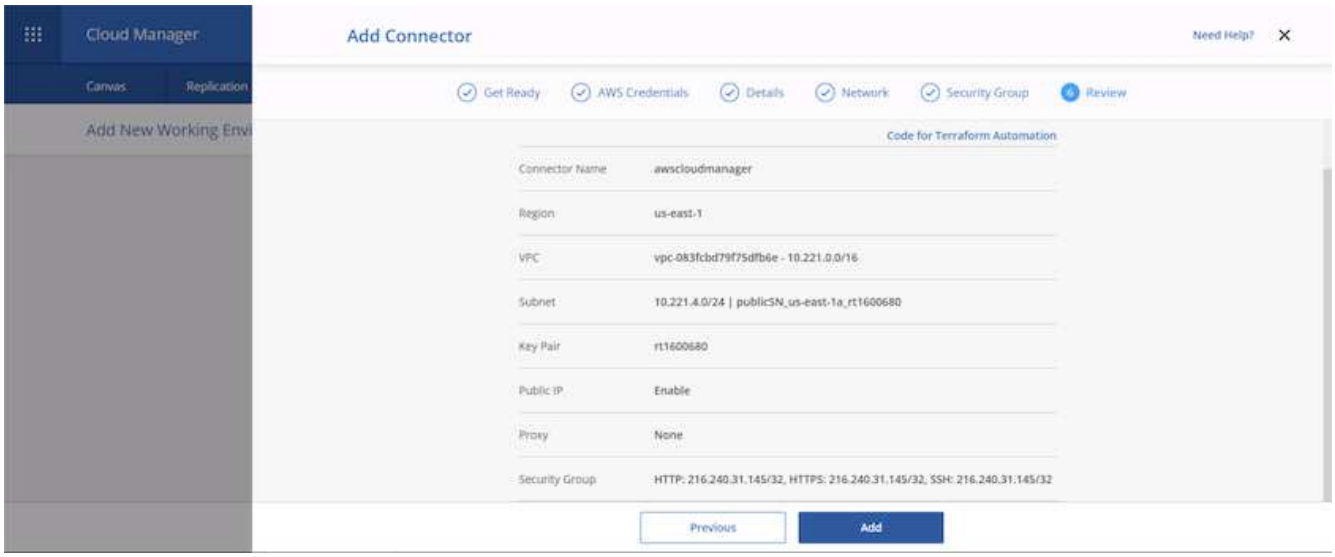

11. The connector now deploys using a cloud formation stack. You can monitor its progress from Cloud Manager or through AWS.

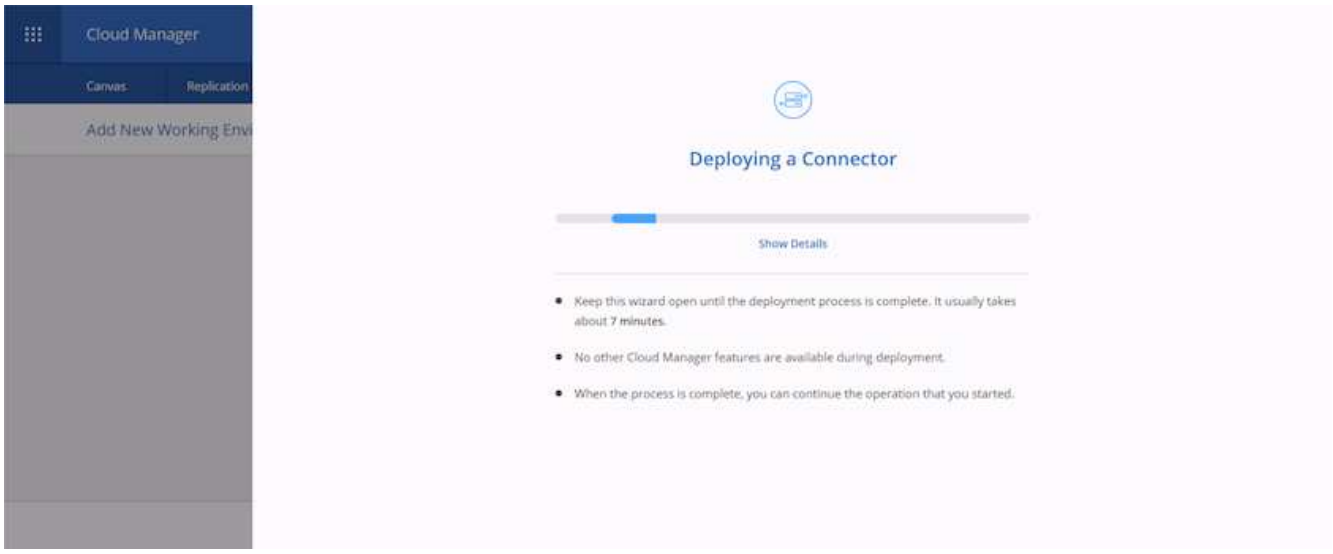

12. When the deployment is complete, a success page appears.

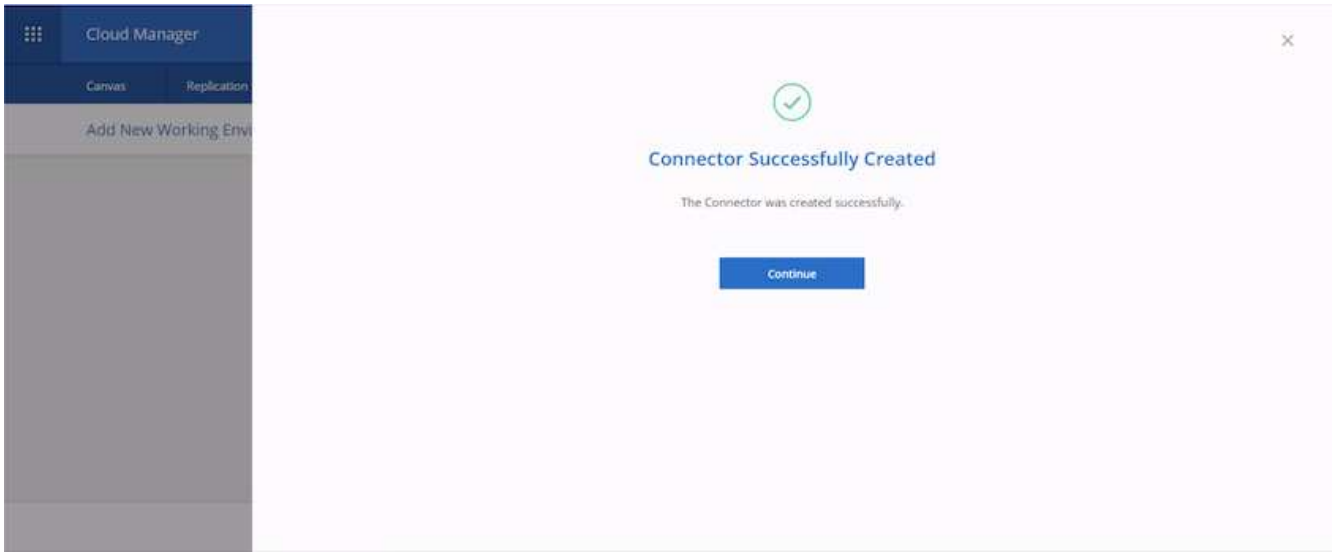

Deploy Cloud Volumes ONTAP

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

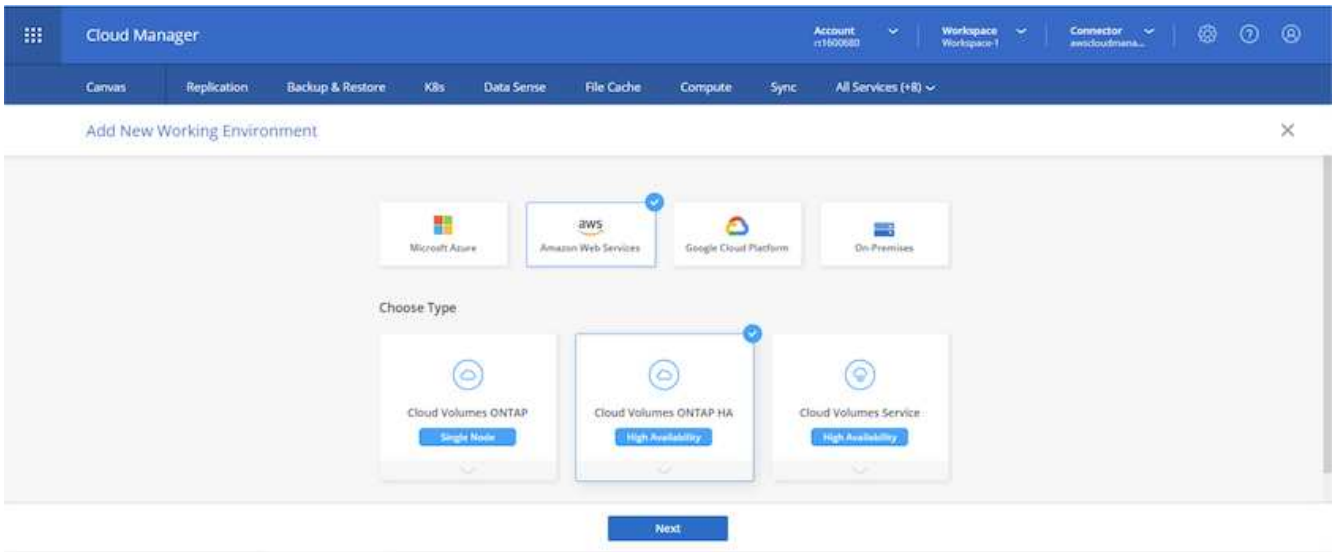

2. If no subscription has been assigned and you wish to purchase with PAYGO, choose Edit Credentials.

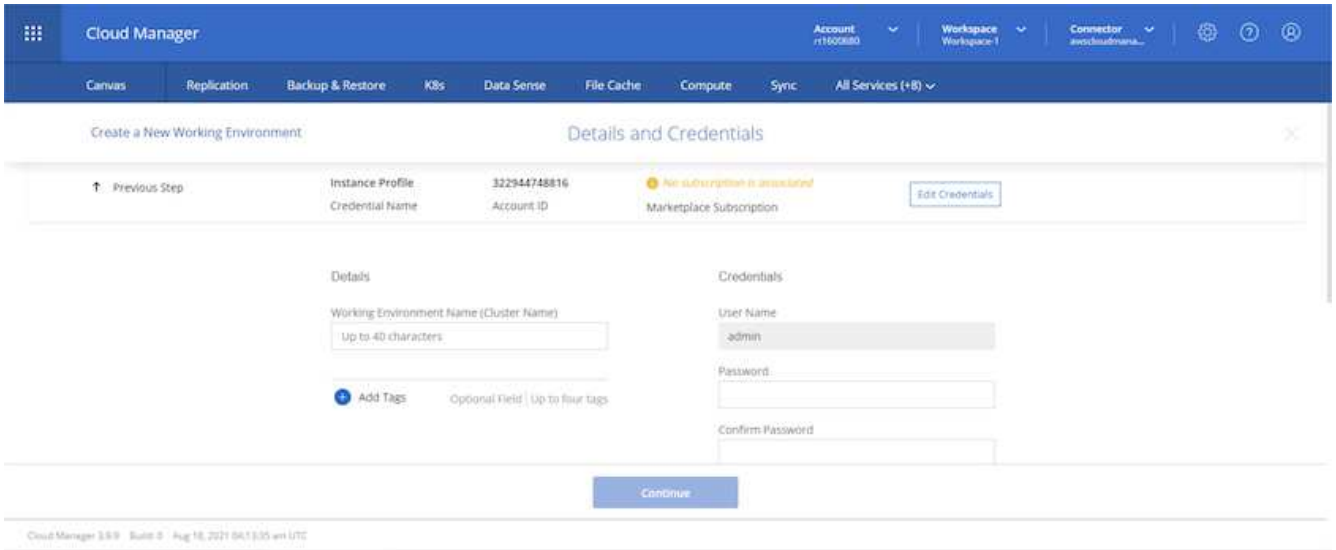

3. Choose Add Subscription.

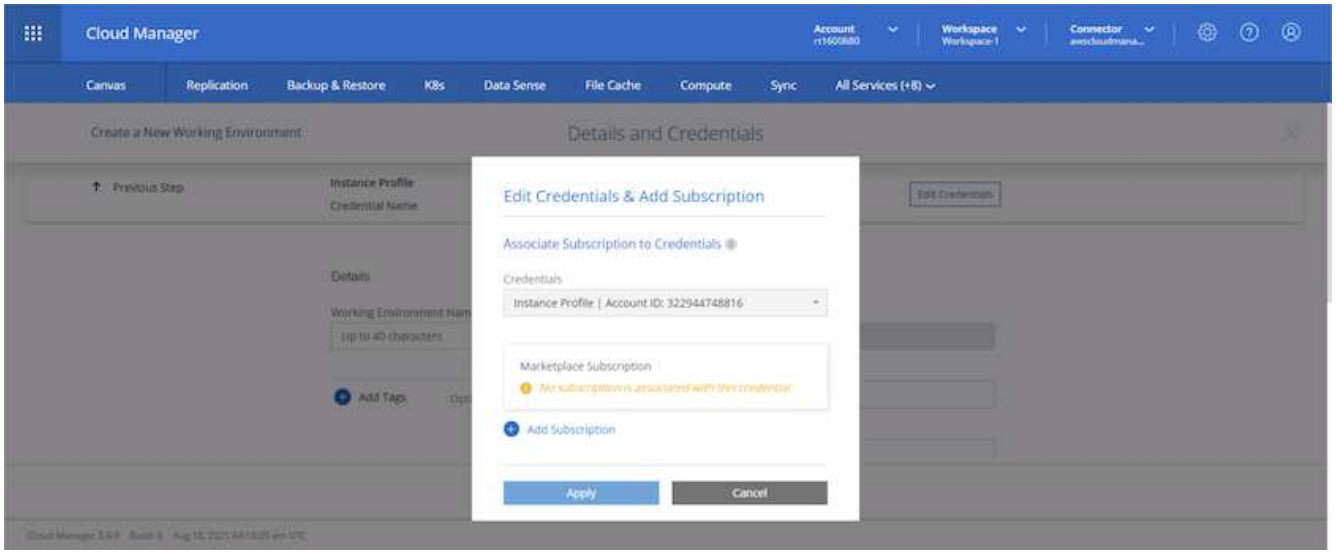

4. Choose the type of contract that you wish to subscribe to. I chose Pay-as-you-go.

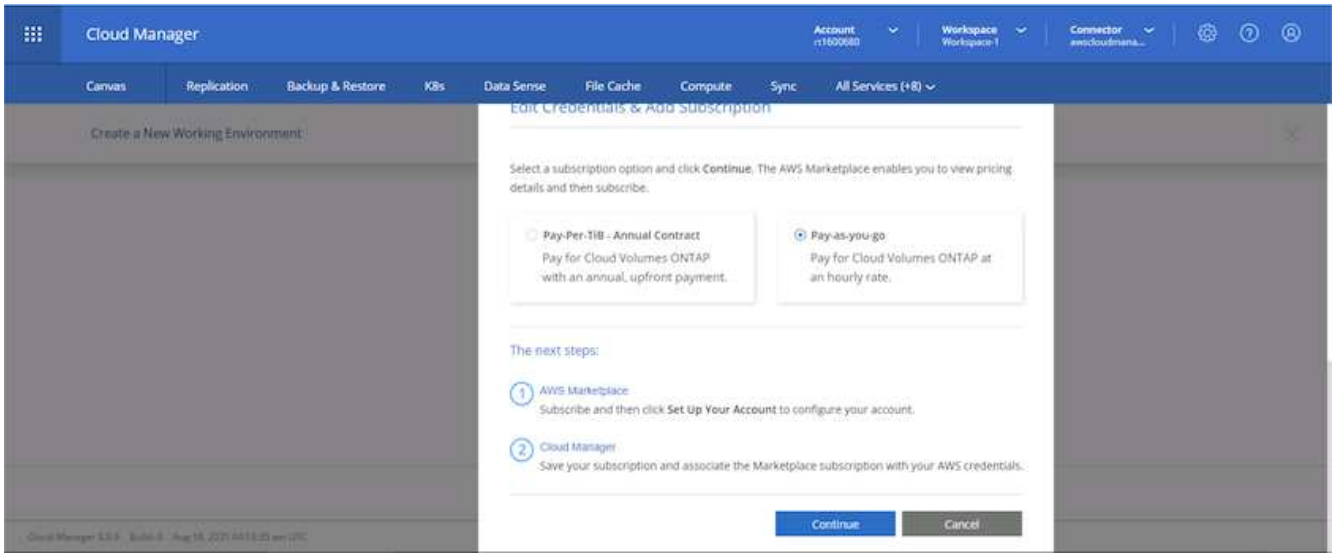

5. You are redirected to AWS; choose Continue to Subscribe.

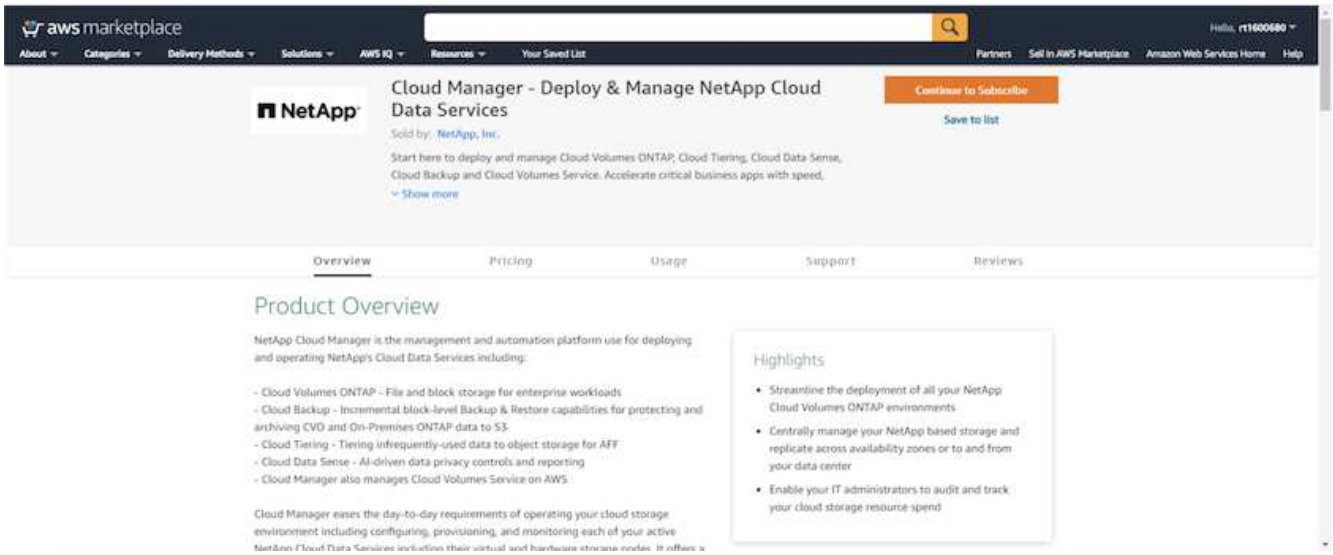

6. Subscribe and you are redirected back to NetApp Cloud Central. If you have already subscribed and don't get redirected, choose the "Click here" link.

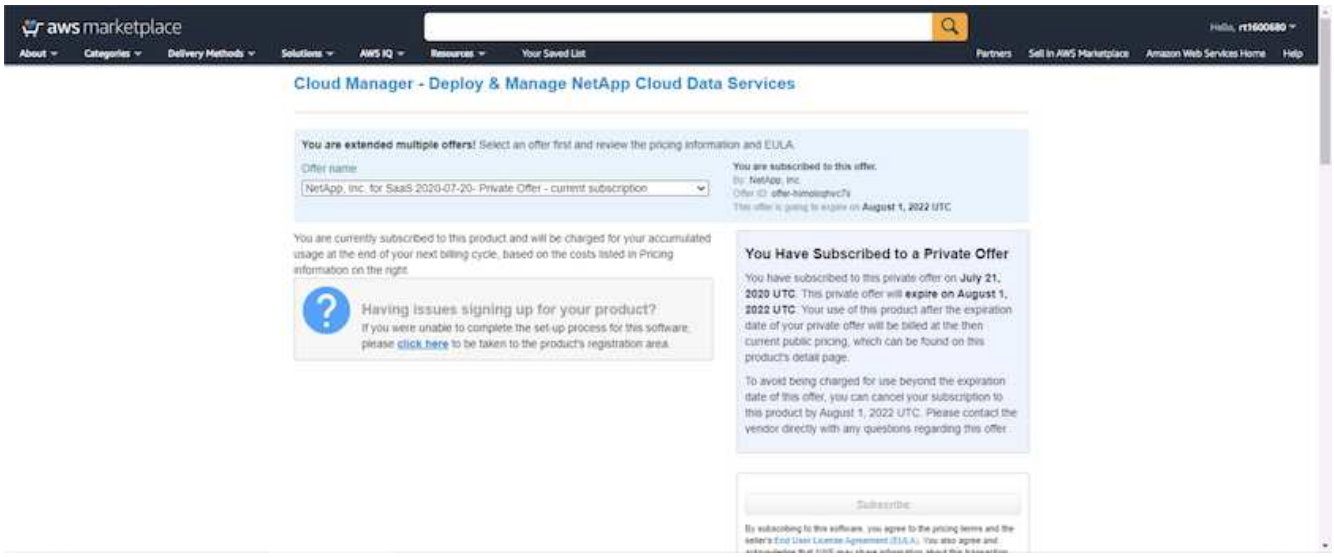

7. You are redirected to Cloud Central where you must name your subscription and assign it to your Cloud Central account.

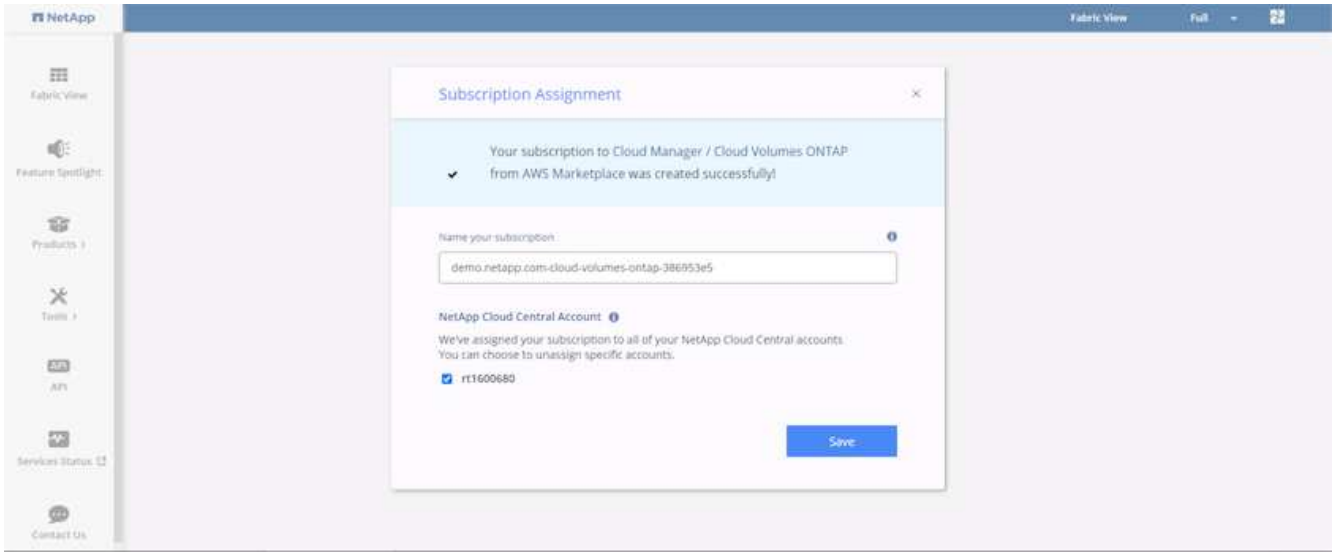

8. When successful, a check mark page appears. Navigate back to your Cloud Manager tab.

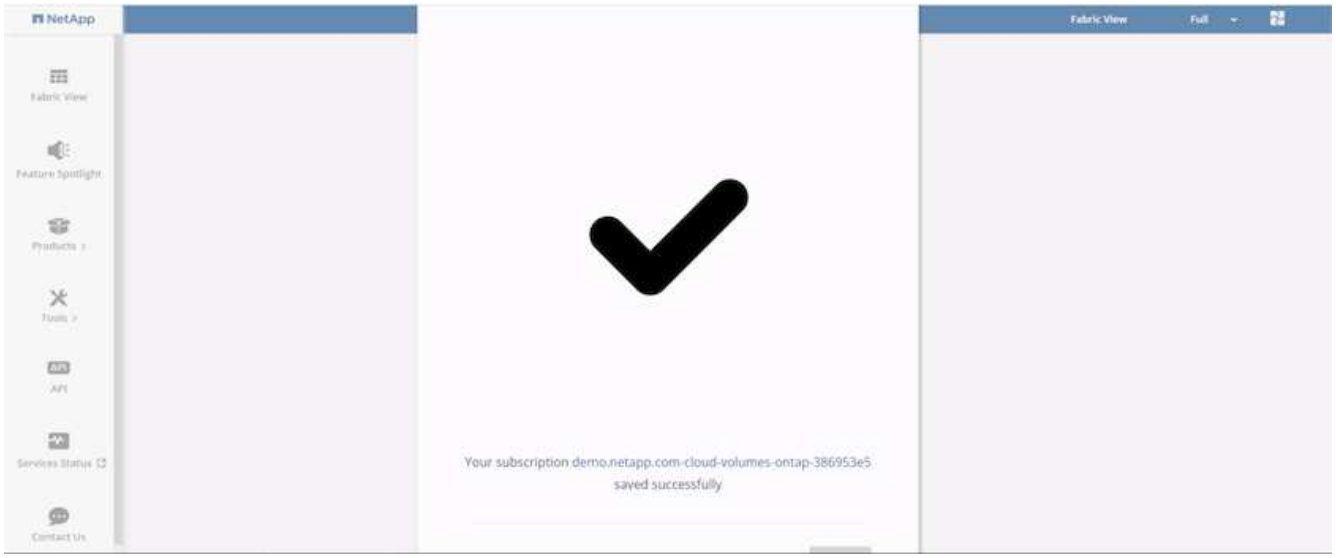

9. The subscription now appears in Cloud Central. Click Apply to continue.

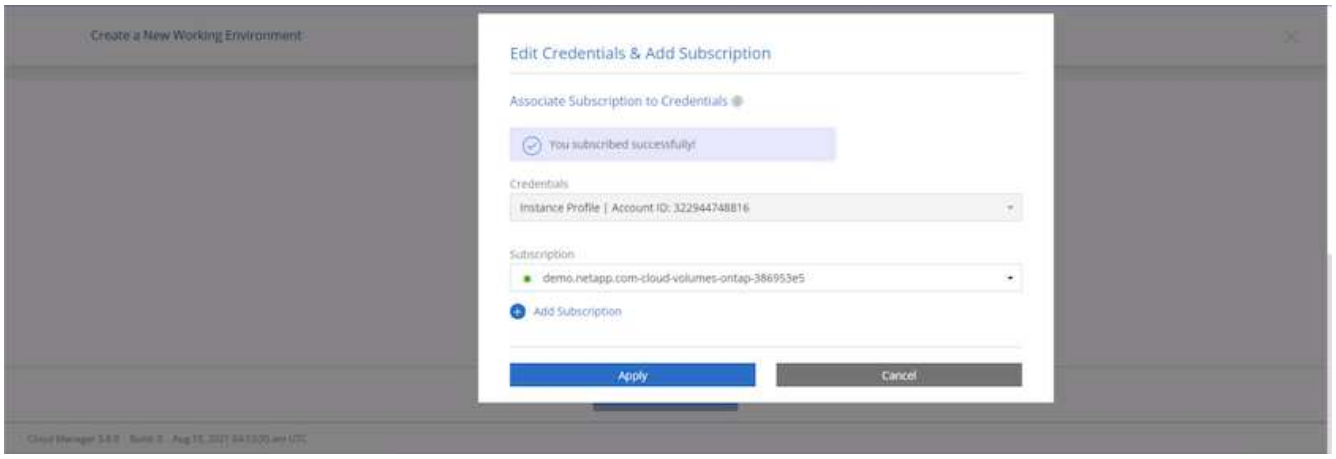

10. Enter the working environment details such as:
    a. Cluster name
    b. Cluster password
    c. AWS tags (Optional)

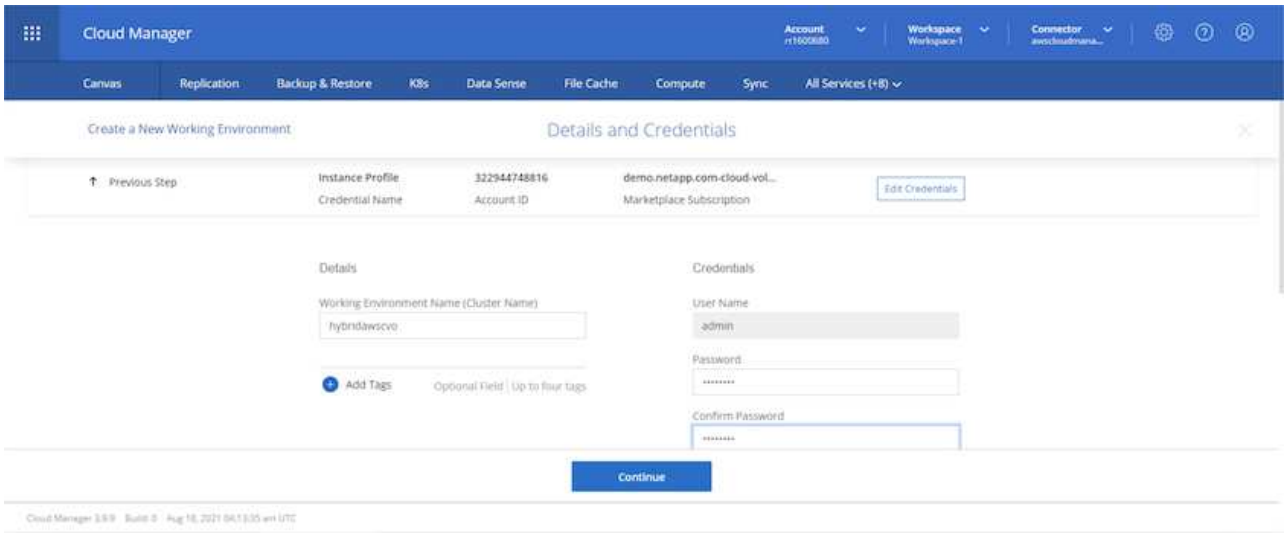

11. Choose which additional services you would like to deploy. To discover more about these services, visit the NetApp Cloud Homepage.

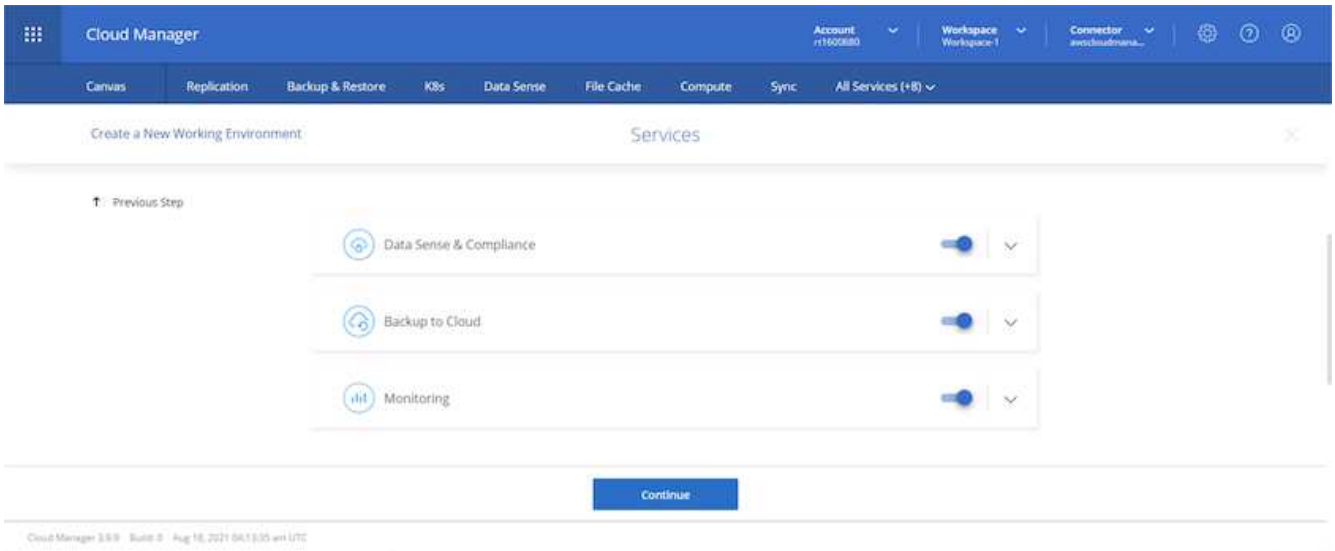

12. Choose whether to deploy in multiple availability zones (reguires three subnets, each in a different AZ), or a single availability zone. I chose multiple AZs.

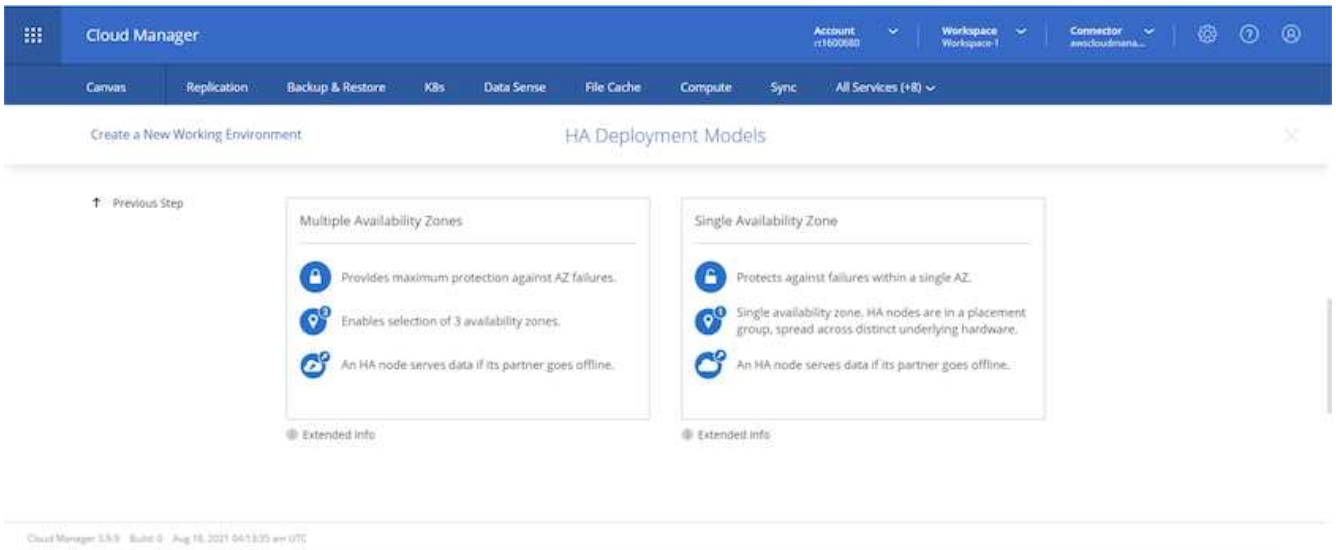

13. Choose the region, VPC, and security group for the cluster to be deployed into. In this section, you also assign the availability zones per node (and mediator) as well as the subnets that they occupy.

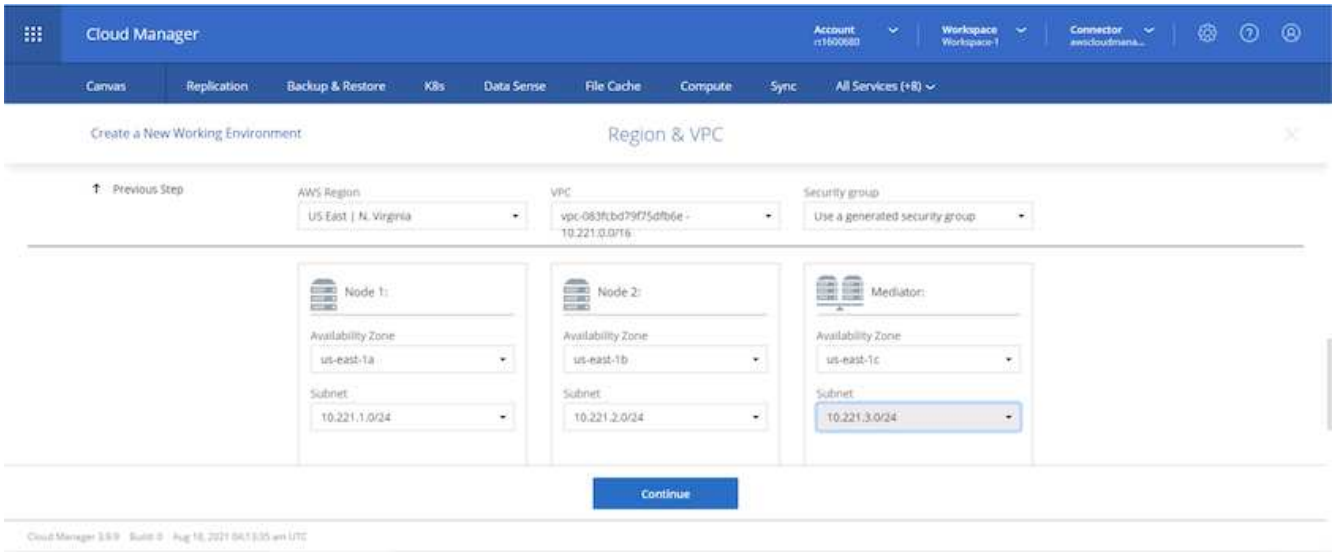

14. Choose the connection methods for the nodes as well as the mediator.

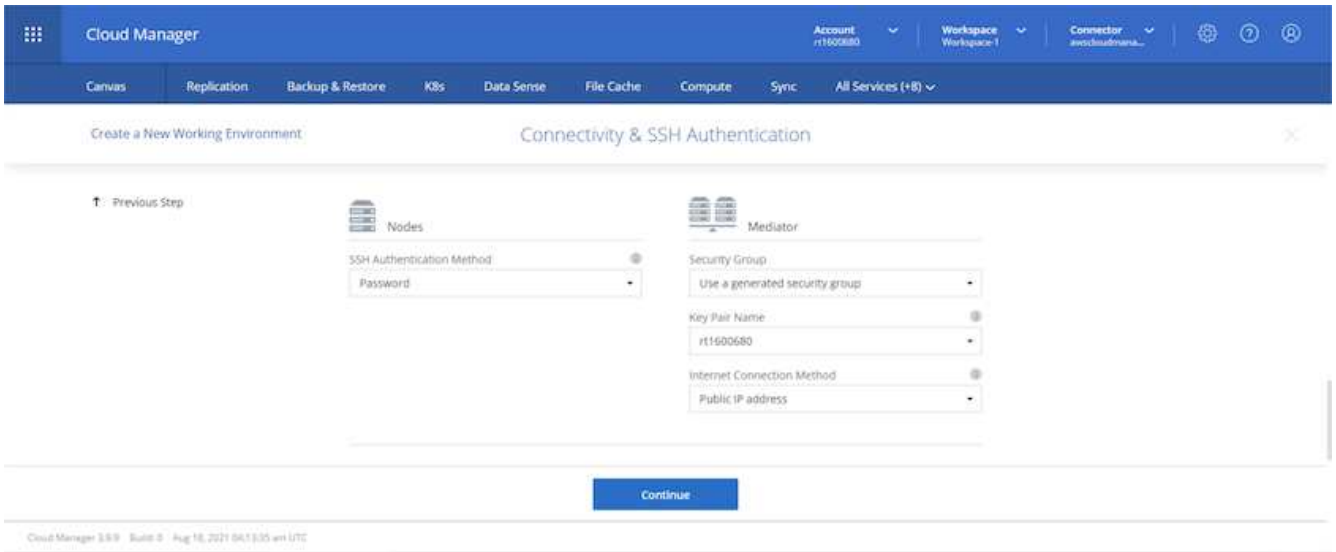

The mediator requires communication with the AWS APIs. A public IP address is not required so long as the APIs are reachable after the mediator EC2 instance has been deployed.

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 Documenation.

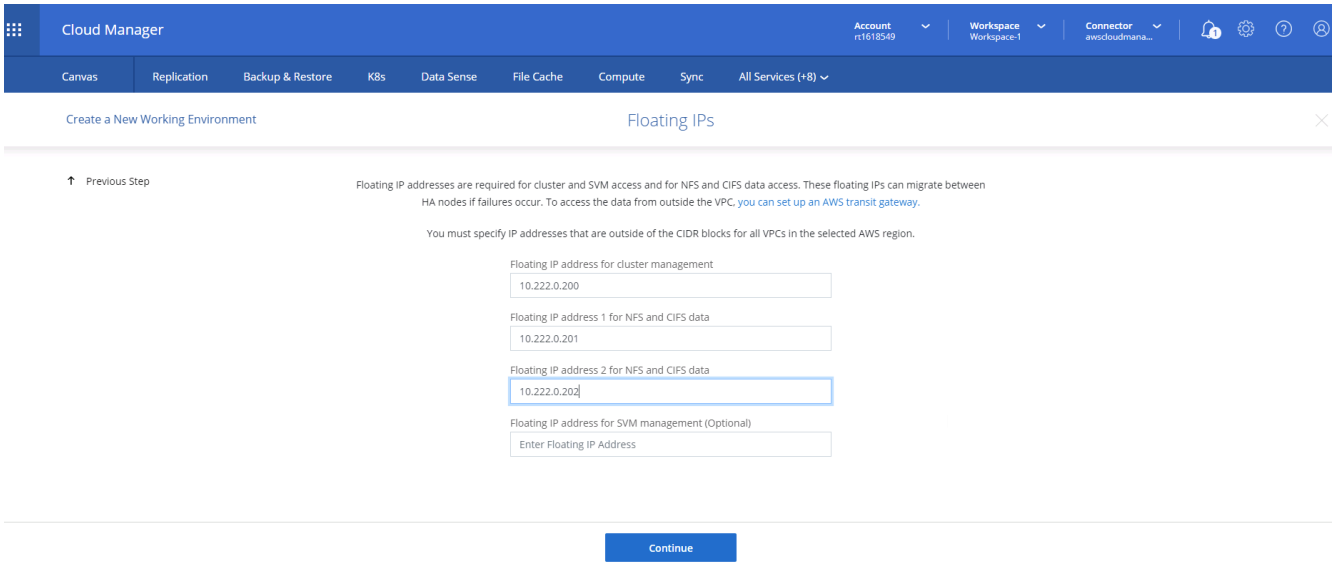

2. Select which route tables the floating IP addresses are added to. These route tables are used by clients to communicate with Cloud Volumes ONTAP.

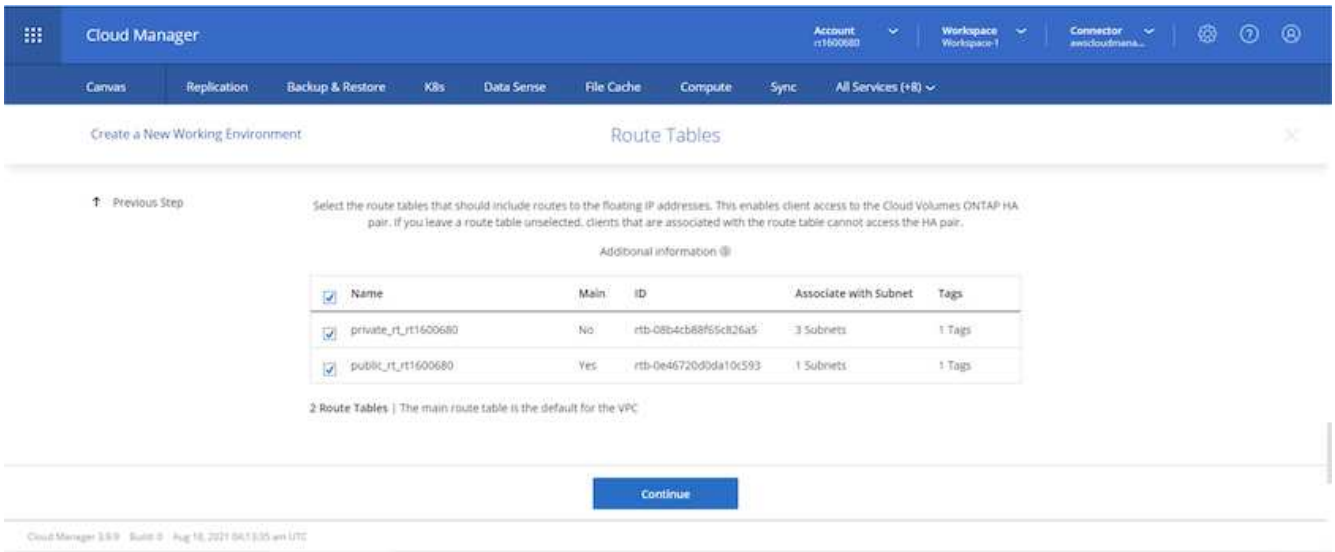

3. Choose whether to enable AWS managed encryption or AWS KMS to encrypt the ONTAP root, boot, and data disks.

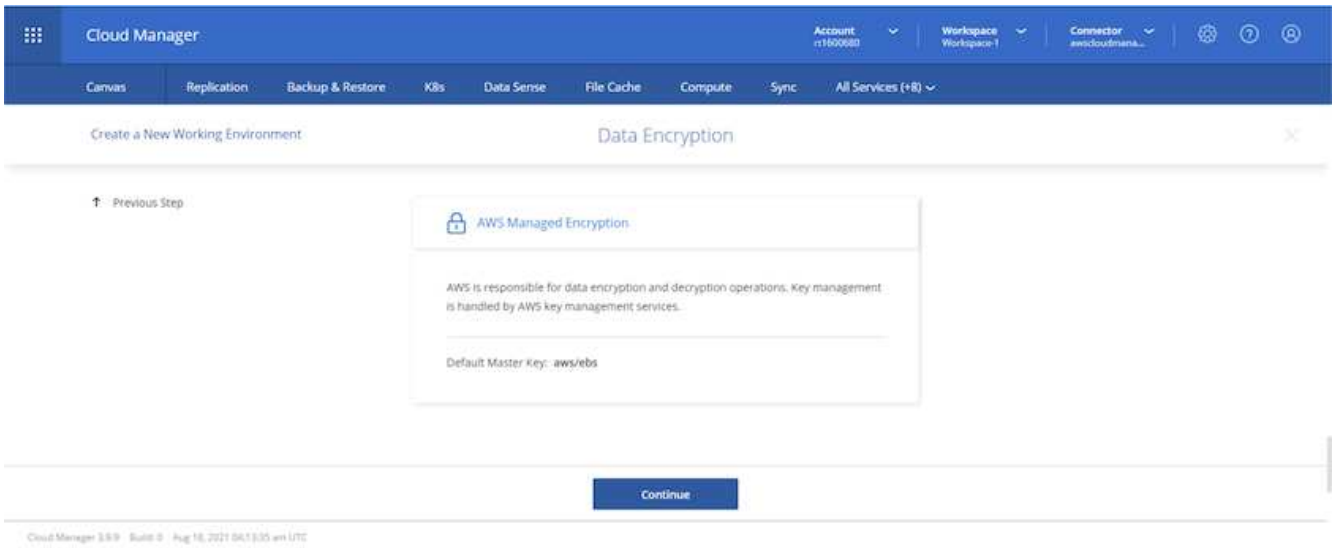

4. Choose your licensing model. If you don't know which to choose, contact your NetApp representative.

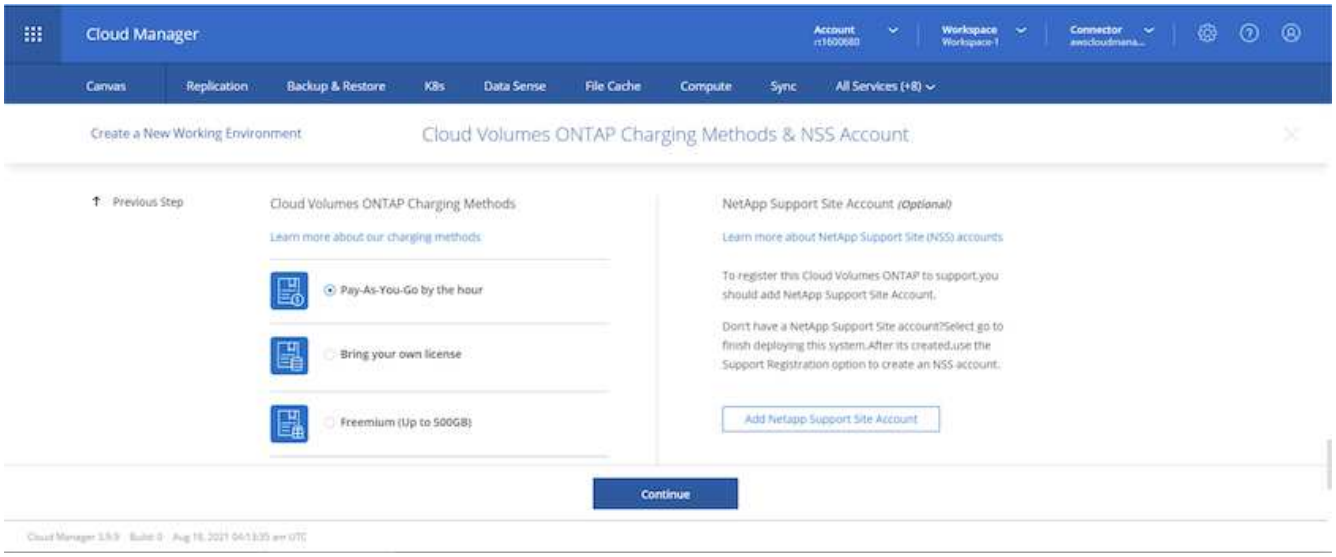

5. Select which configuration best suits your use case. This is related to the sizing considerations covered in the prerequisites page.

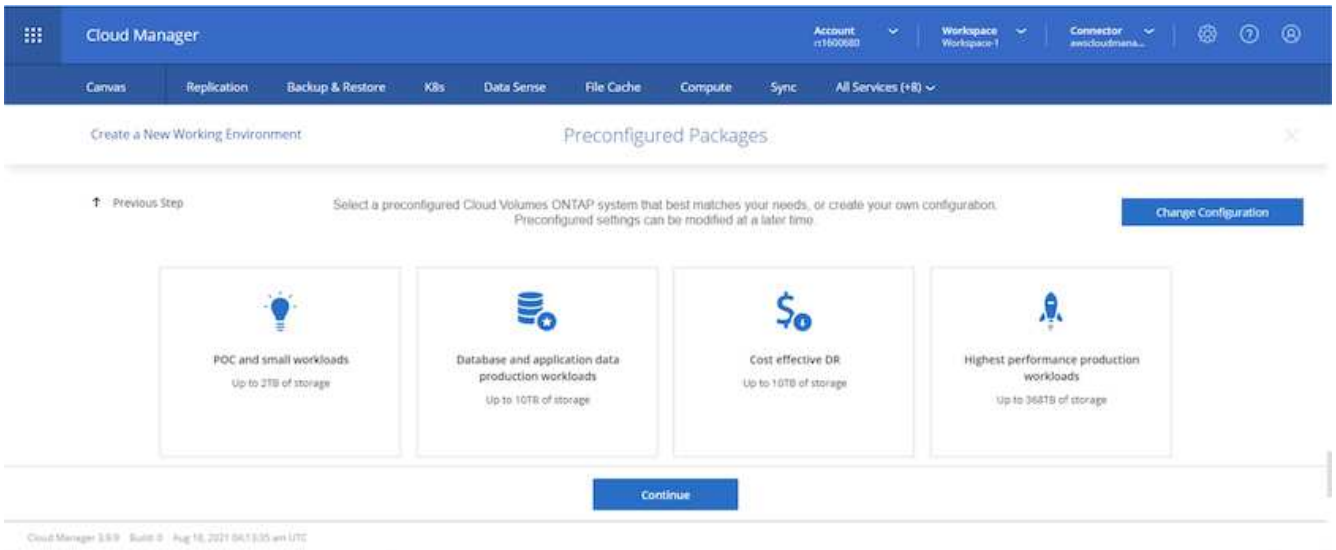

6. Optionally, create a volume. This is not required, because the next steps use SnapMirror, which creates the volumes for us.

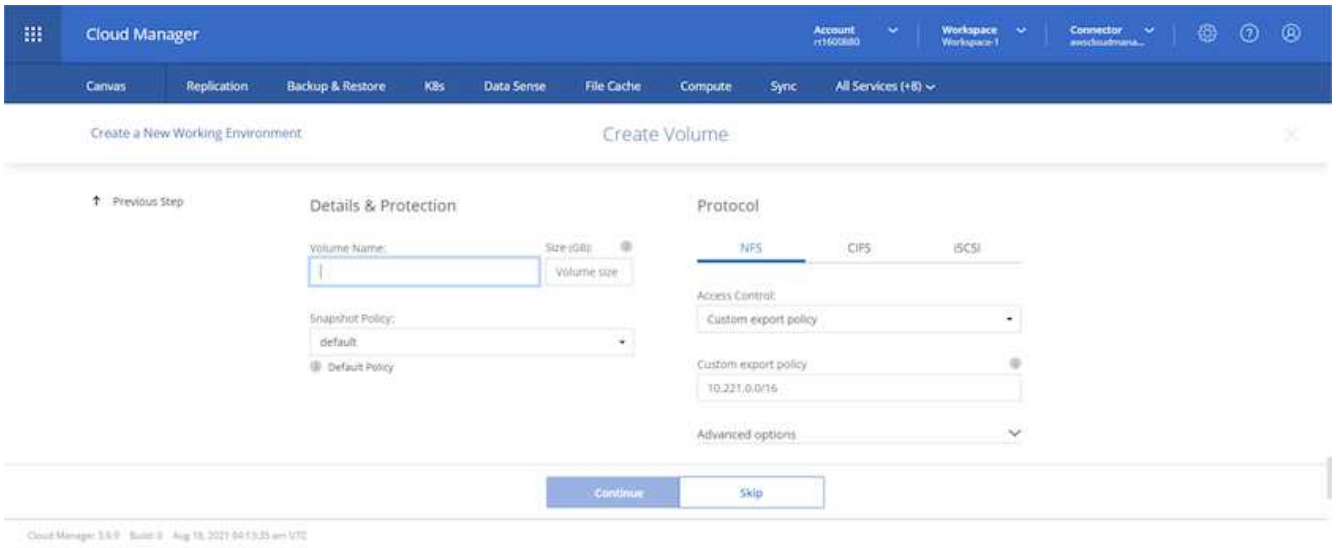

7. Review the selections made and tick the boxes to verify that you understand that Cloud Manager deploys resources into your AWS environment. When ready, click Go.

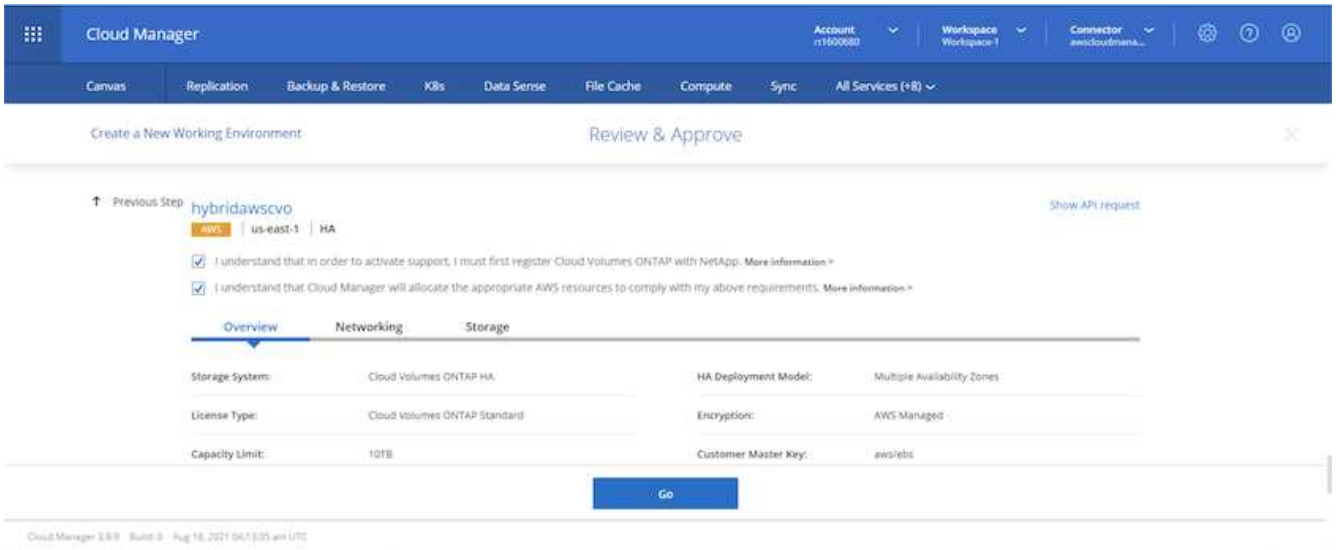

8. Cloud Volumes ONTAP now starts its deployment process. Cloud Manager uses AWS APIs and cloud formation stacks to deploy Cloud Volumes ONTAP. It then configures the system to your specifications, giving you a ready-to-go system that can be instantly utilized. The timing for this process varies depending on the selections made.

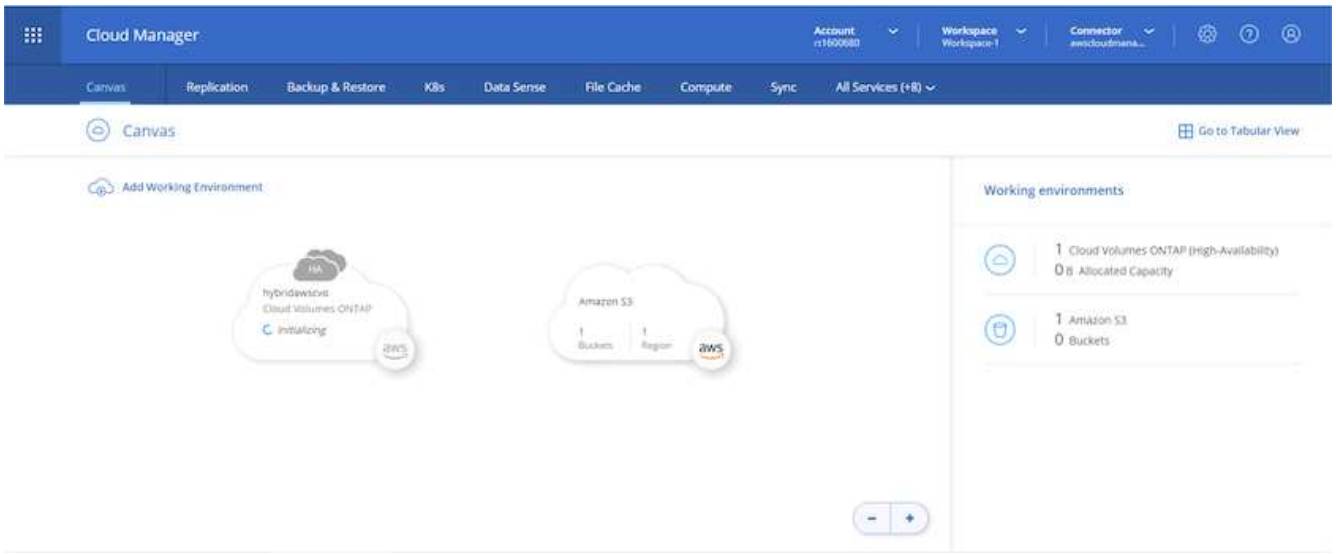

9. You can monitor the progress by navigating to the Timeline.

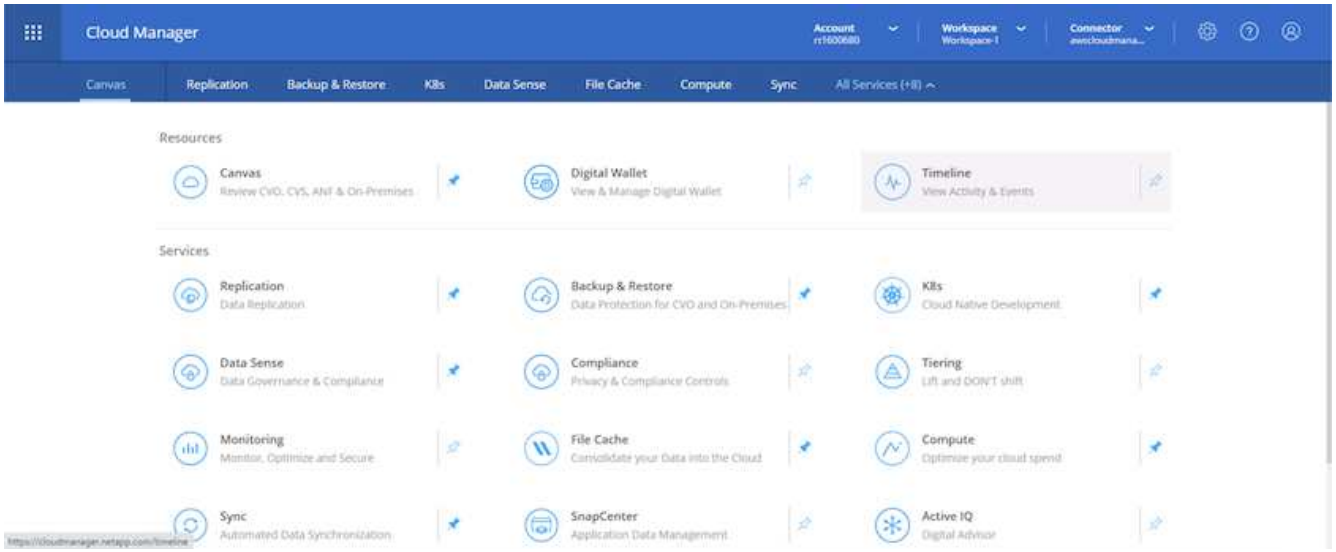

10. The Timeline acts as an audit of all actions performed in Cloud Manager. You can view all of the API calls that are made by Cloud Manager during setup to both AWS as well as the ONTAP cluster. This can also be effectively used to troubleshoot any issues that you face.

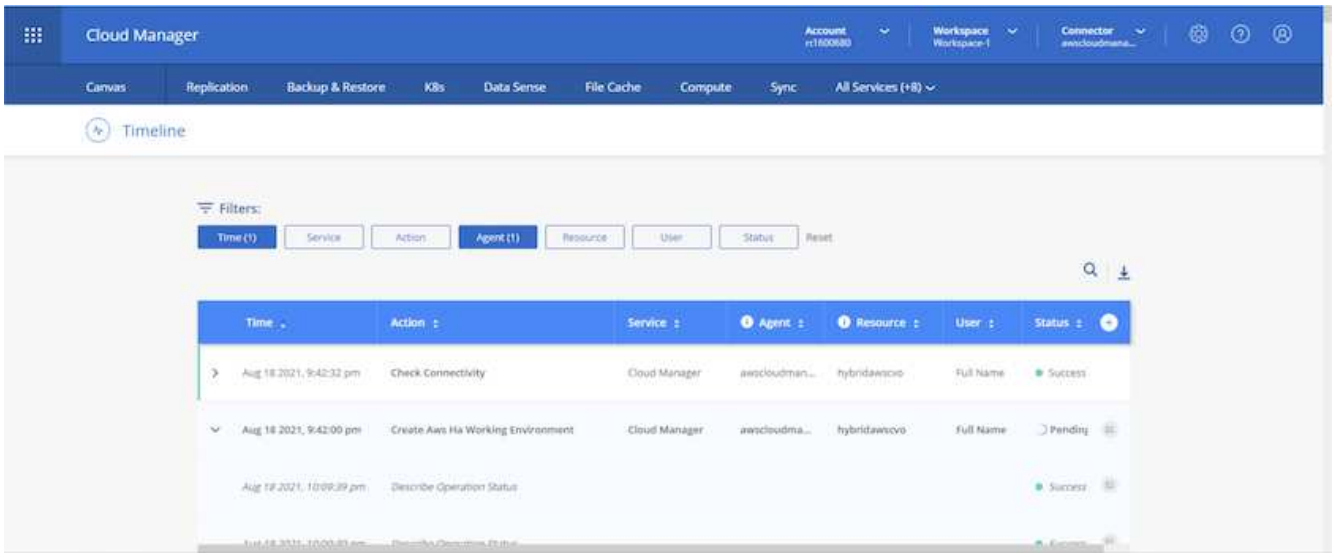

11. After deployment is complete, the CVO cluster appears on the Canvas, which the current capacity. The ONTAP cluster in its current state is fully configured to allow a true, out-of-the-box experience.

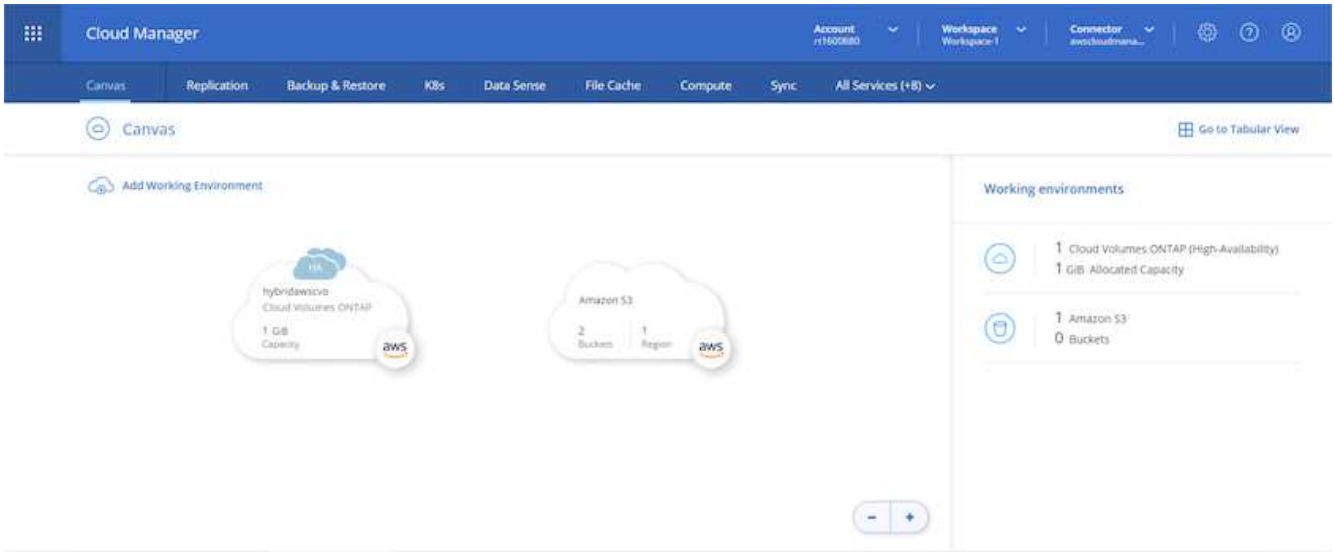

Configure SnapMirror from on-premises to cloud

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.

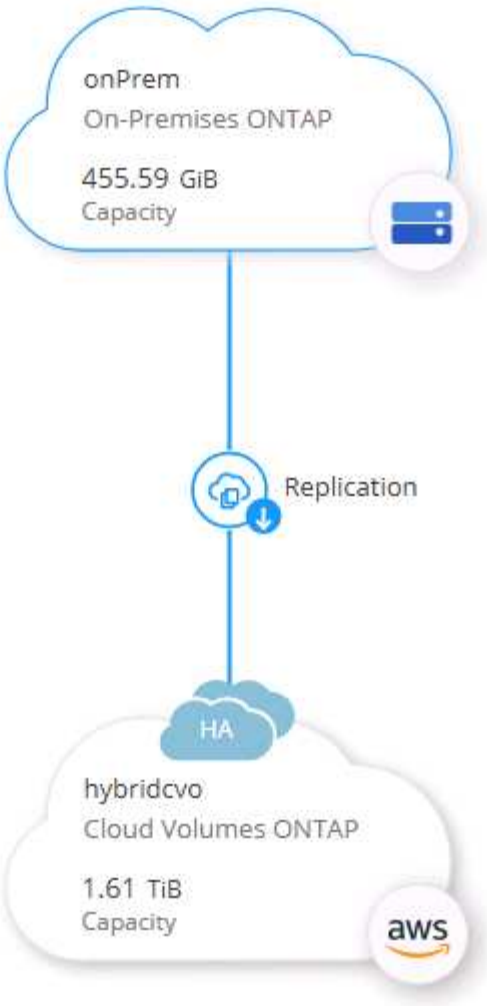

Select Enable.

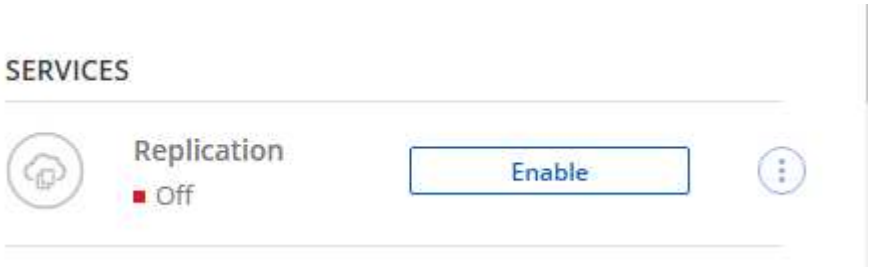

Or Options.

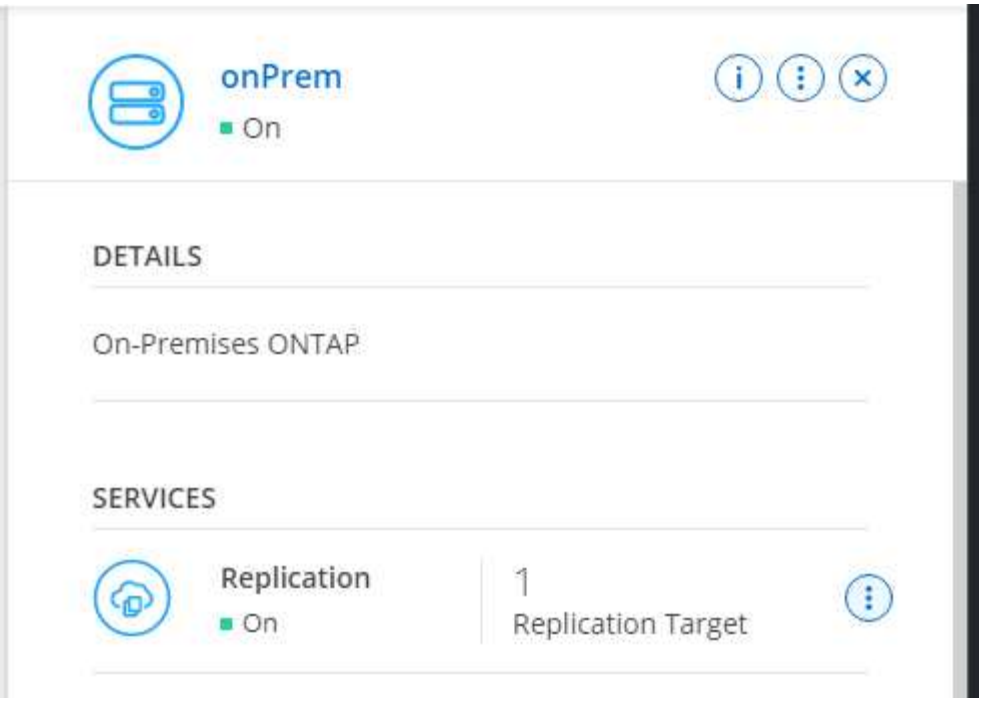

Replicate.

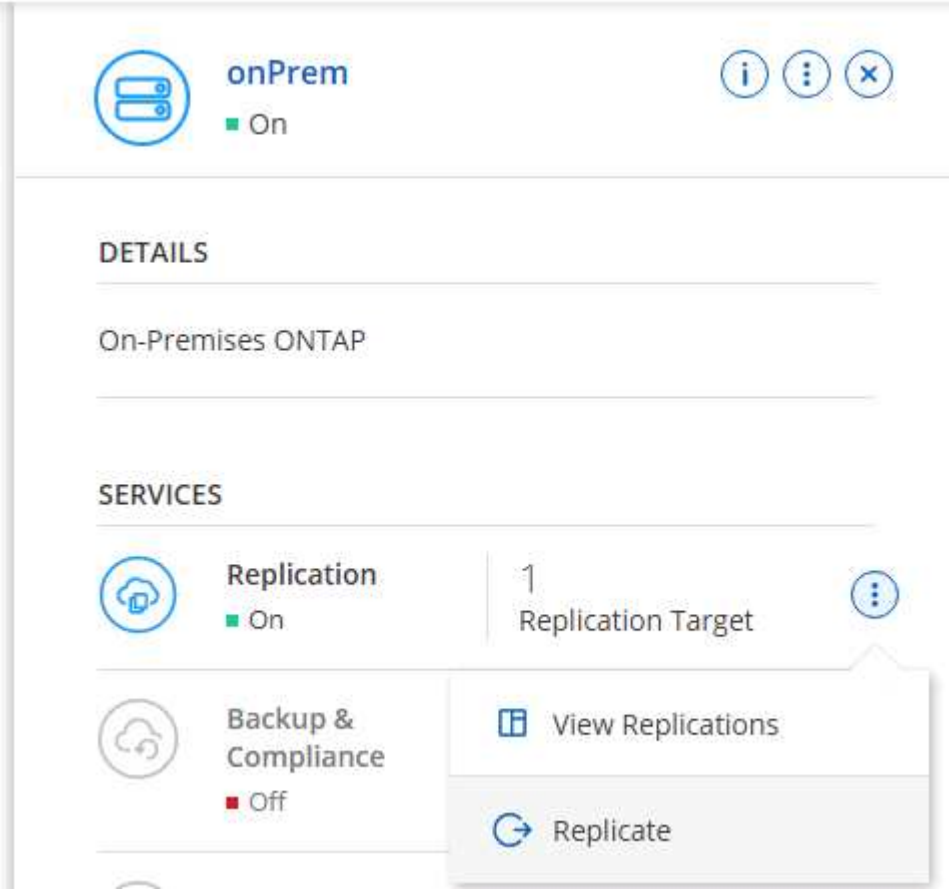

2. If you did not drag and drop, choose the destination cluster to replicate to.

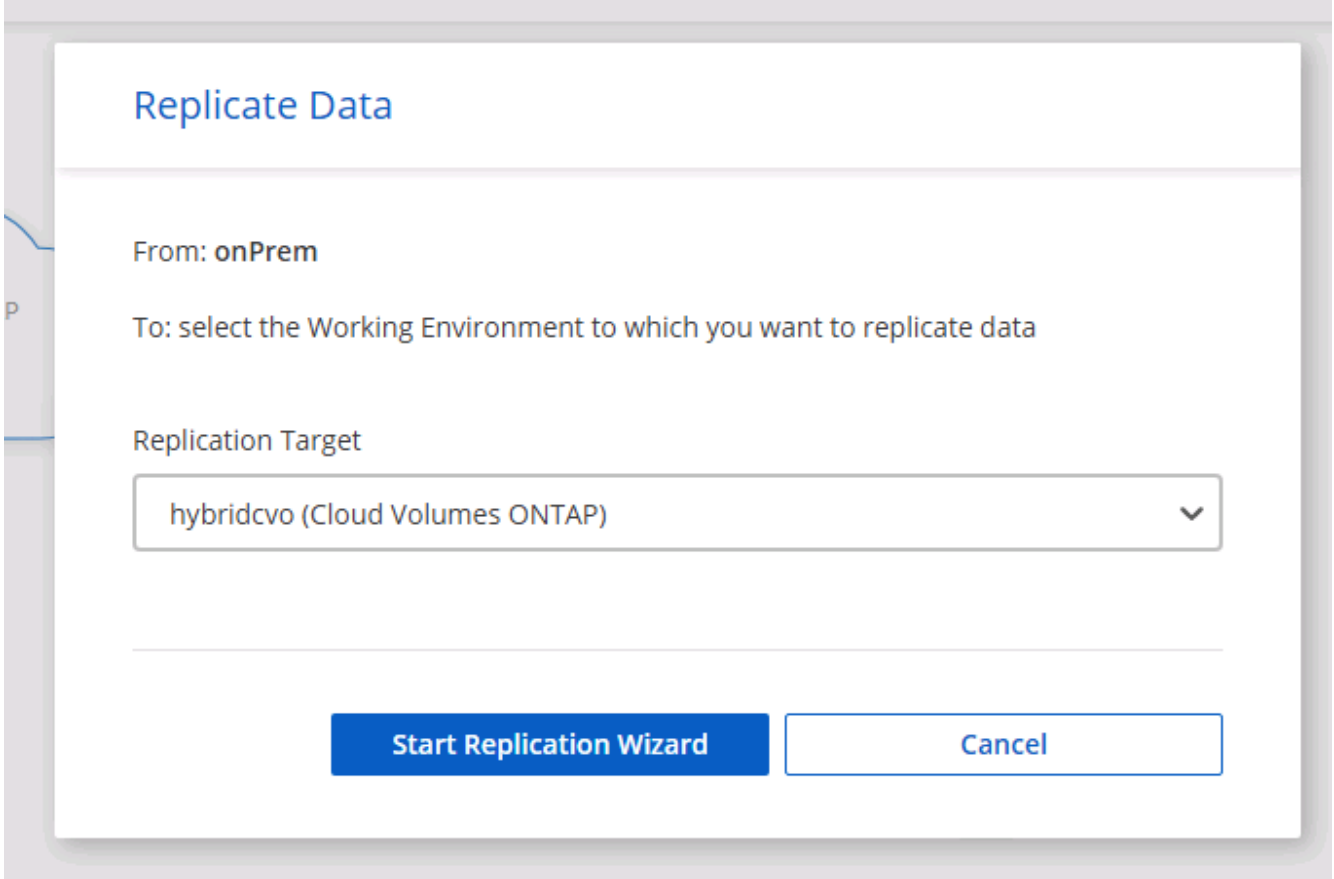

3. Choose the volume that you'd like to replicate. We replicated the data and all log volumes.

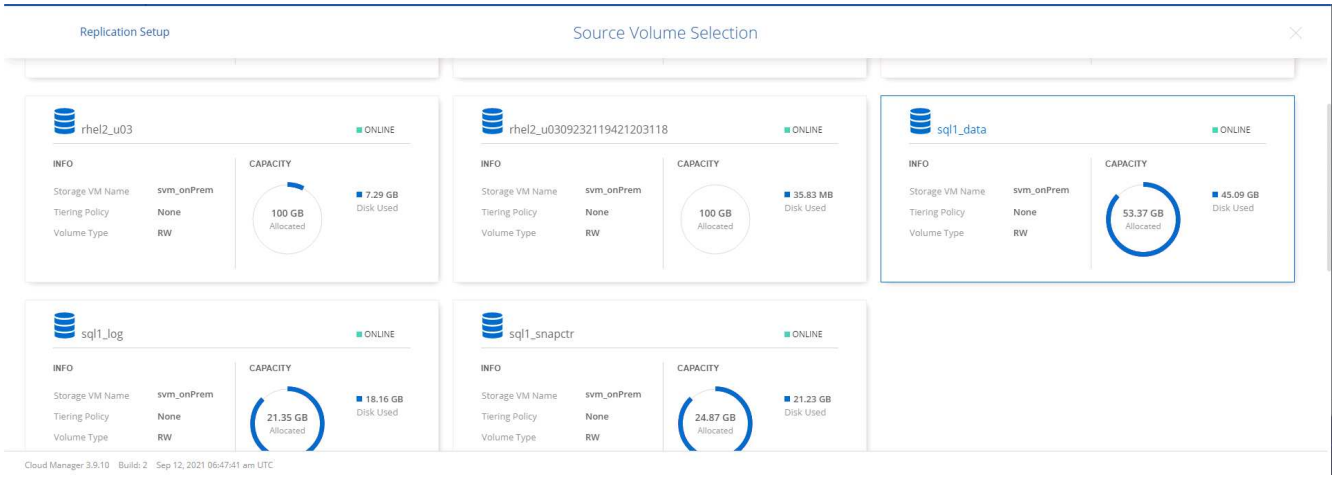

4. Choose the destination disk type and tiering policy. For disaster recovery, we recommend an SSD as the disk type and to maintain data tiering. Data tiering tiers the mirrored data into low-cost object storage and saves you money on local disks. When you break the relationship or clone the volume,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

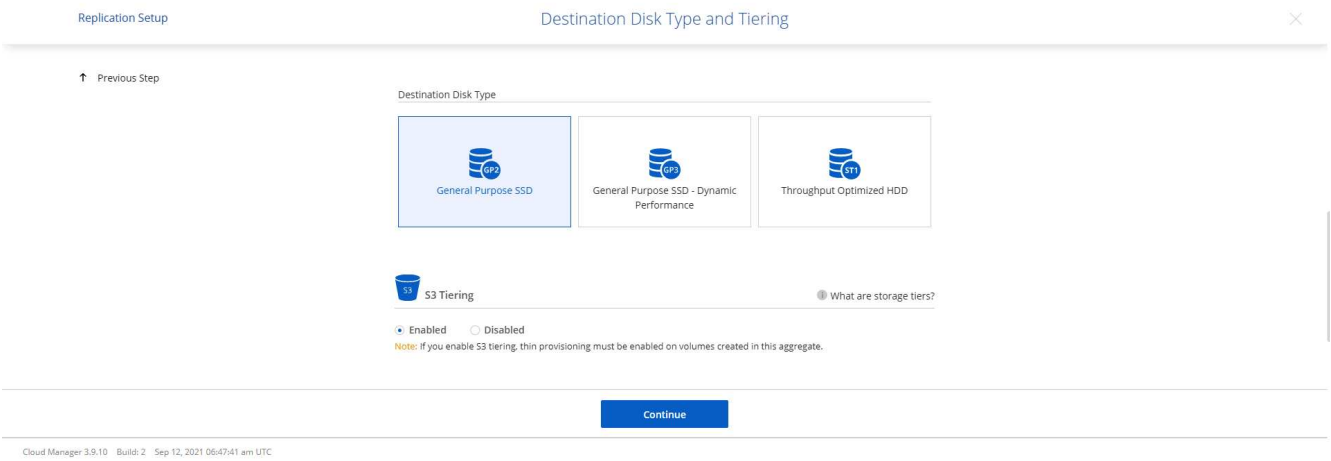

5. Select the destination volume name: we chose `[source_volume_name]_dr`.

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.

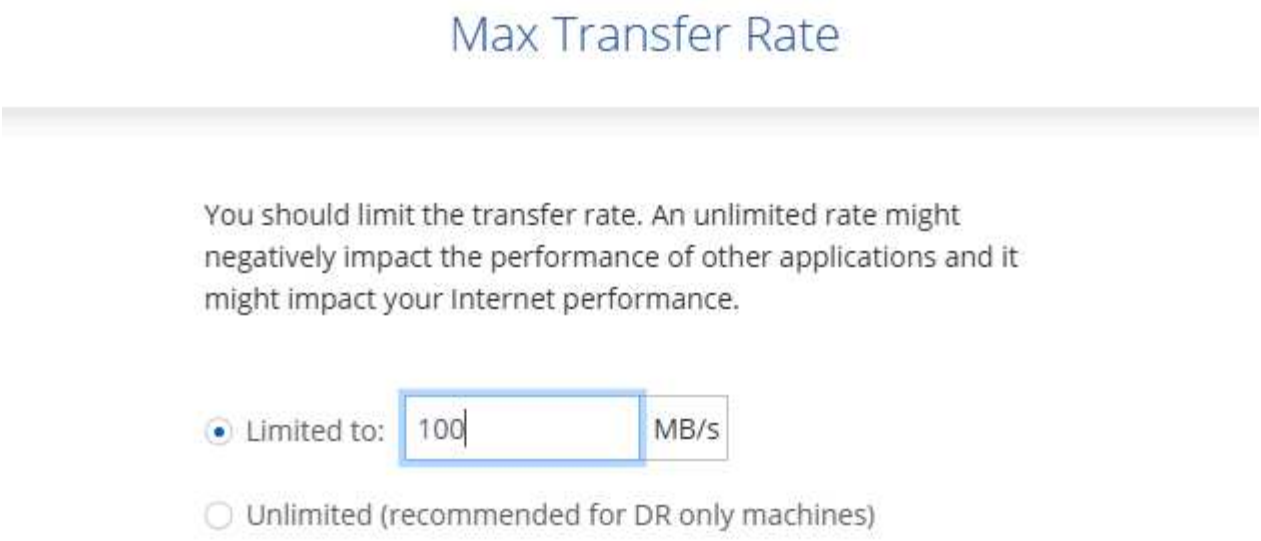

7. Define the replication policy. We chose a Mirror, which takes the most recent dataset and replicates that into the destination volume. You could also choose a different policy based on your requirements.

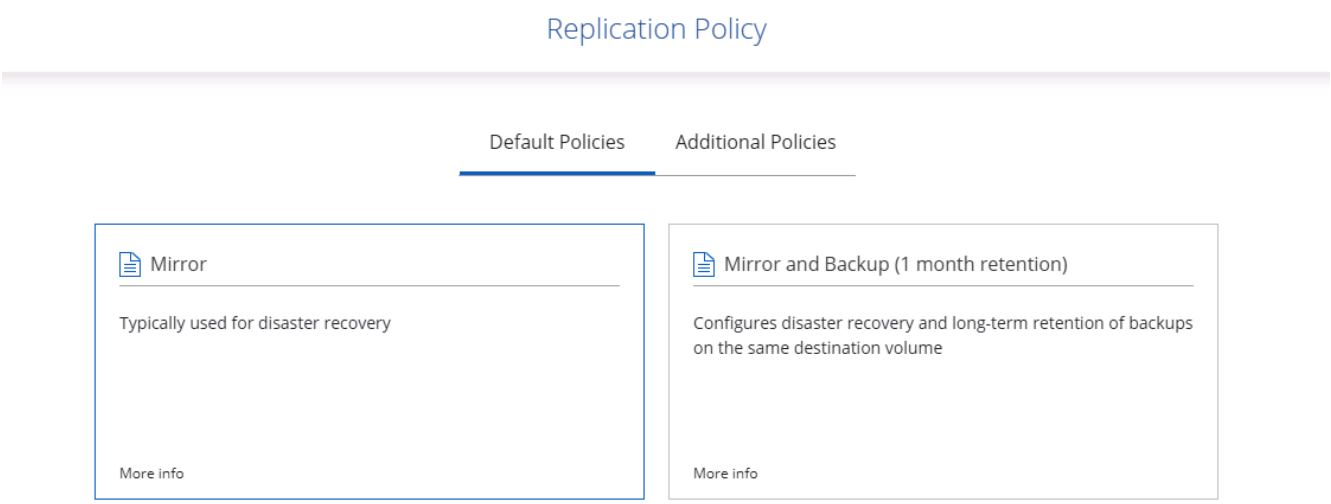

8. Choose the schedule for triggering replication. NetApp recommends setting a "daily" schedule of for the data volume and an "hourly" schedule for the log volumes, although this can be changed based on requirements.

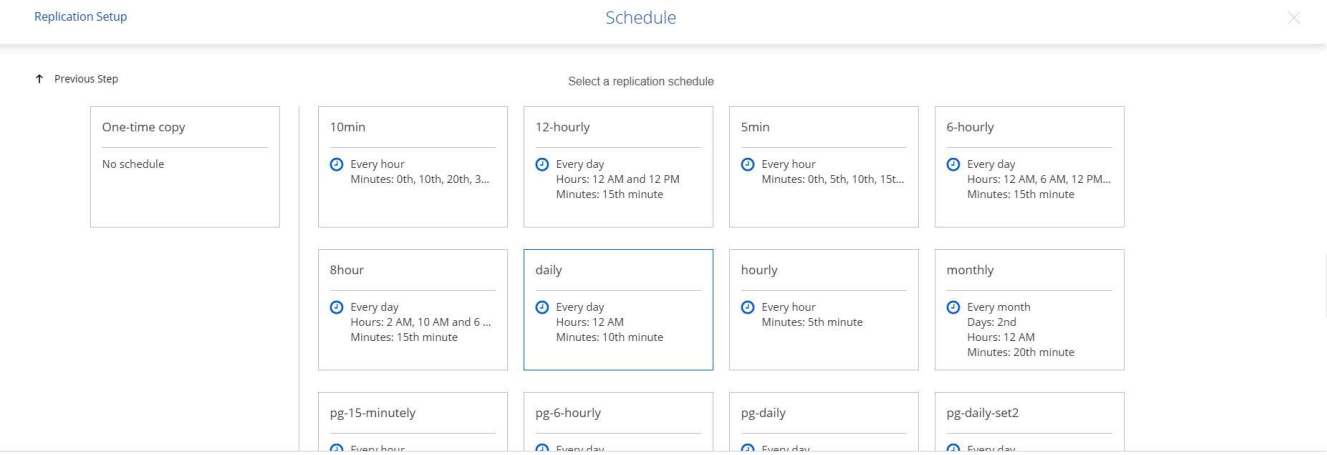

9. Review the information entered, click Go to trigger the cluster peer and SVM peer (if this is your first time replicating between the two clusters), and then implement and initialize the SnapMirror relationship.

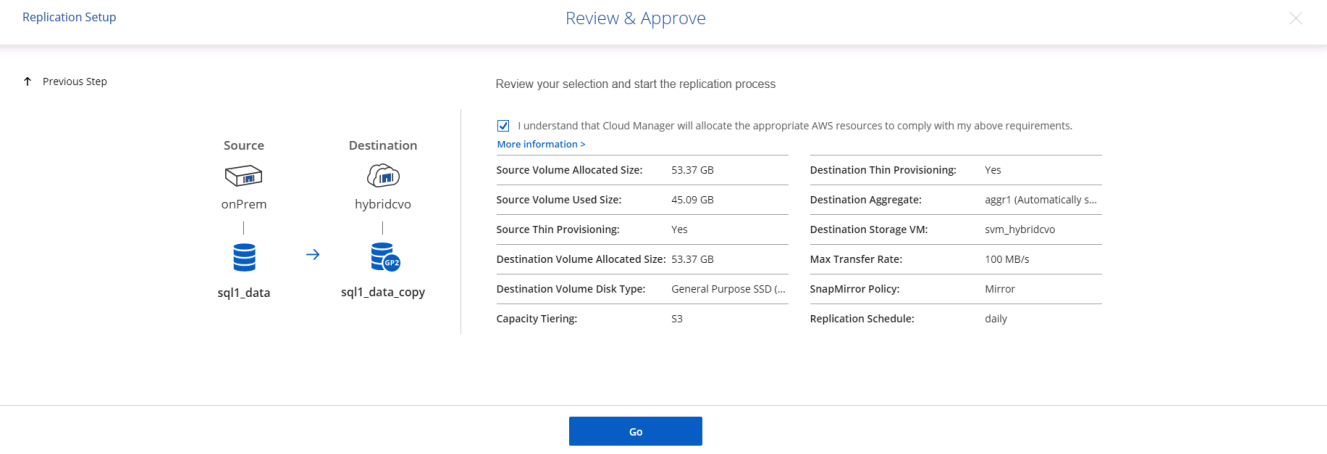

10. Continue this process for data volumes and log volumes.
11. To check all of your relationships, navigate to the Replication tab inside Cloud Manager. Here you can manage your relationships and check on their status.

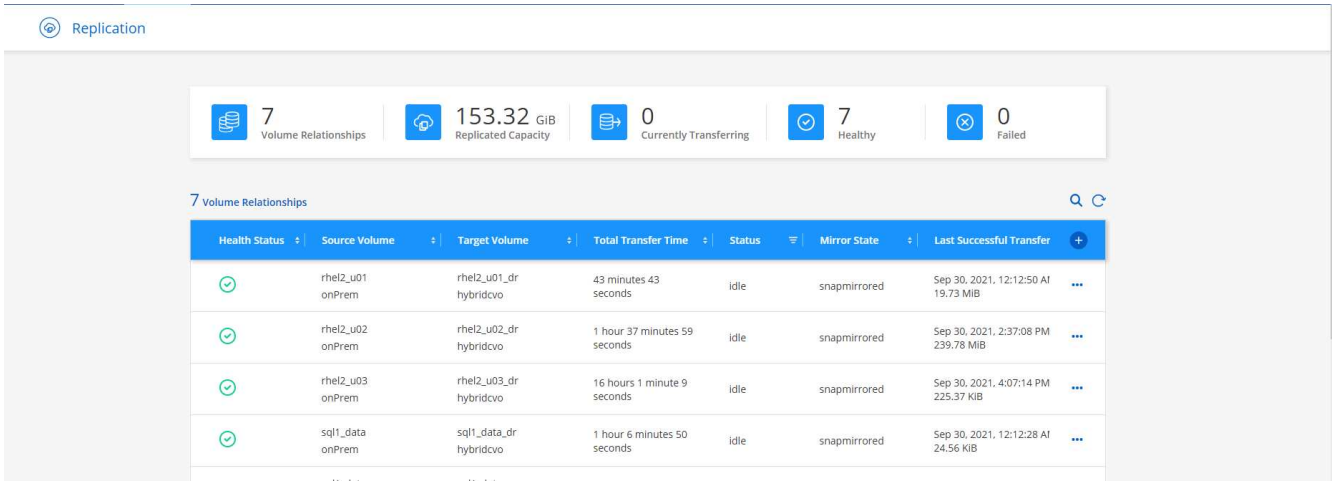

12. After all the volumes have been replicated, you are in a steady state and ready to move on to the disaster recovery and dev/test workflows.

3. Deploy EC2 compute instance for database workload

AWS has preconfigured EC2 compute instances for various workloads. The choice of instance type determines the number of CPU cores, memory capacity, storage type and capacity, and network performance. For the use cases, with the exception of the OS partition, the main storage to run database workload is allocated from CVO or the FSx ONTAP storage engine. Therefore, the main factors to consider are the choice of CPU cores, memory, and network performance level. Typical AWS EC2 instance types can be found here: EC2 Instance Type.

Sizing the compute instance

1. Select the right instance type based on the required workload. Factors to consider include the number of business transactions to be supported, the number of concurrent users, data set sizing, and so on.
2. EC2 instance deployment can be launched through the EC2 Dashboard. The exact deployment procedures are beyond the scope of this solution. See Amazon EC2 for details.

Linux instance configuration for Oracle workload

This section contain additional configuration steps after an EC2 Linux instance is deployed.

1. Add an Oracle standby instance to the DNS server for name resolution within the SnapCenter management domain.
2. Add a Linux management user ID as the SnapCenter OS credentials with sudo permissions without a password. Enable the ID with SSH password authentication on the EC2 instance. (By default, SSH password authentication and passwordless sudo is turned off on EC2 instances.)
3. Configure Oracle installation to match with on-premises Oracle installation such as OS patches, Oracle versions and patches, and so on.
4. NetApp Ansible DB automation roles can be leveraged to configure EC2 instances for database dev/test and disaster recovery use cases. The automation code can be download from the NetApp public GitHub site: Oracle 19c Automated Deployment. The goal is to install and configure a database software stack on an EC2 instance to match on-premises OS and database configurations.

Windows instance configuration for SQL Server workload

This section lists additional configuration steps after an EC2 Windows instance is initially deployed.

1. Retrieve the Windows administrator password to log in to an instance via RDP.
2. Disable the Windows firewall, join the host to Windows SnapCenter domain, and add the instance to the DNS server for name resolution.
3. Provision a SnapCenter log volume to store SQL Server log files.
4. Configure iSCSI on the Windows host to mount the volume and format the disk drive.
5. Again, many of the previous tasks can be automated with the NetApp automation solution for SQL Server. Check the NetApp automation public GitHub site for newly published roles and solutions: NetApp Automation.

Workflow for dev/test bursting to cloud

The agility of the public cloud, the time to value, and the cost savings are all meaningful value propositions for enterprises adopting the public cloud for database application development and testing effort. There is no better tool than SnapCenter to make this a

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

Clone an Oracle Database for dev/test from a replicated snapshot backup

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

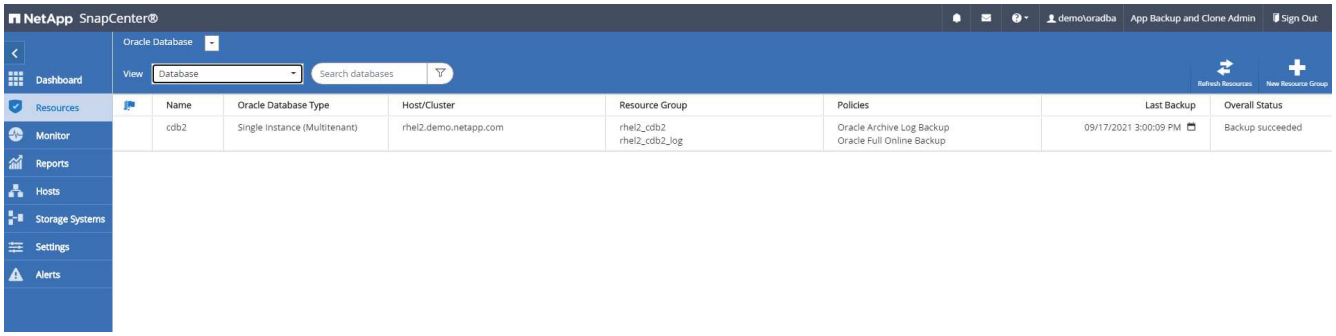

2. Click the intended on-premises database name for the backup topology and the detailed view. If a secondary replicated location is enabled, it shows linked mirror backups.

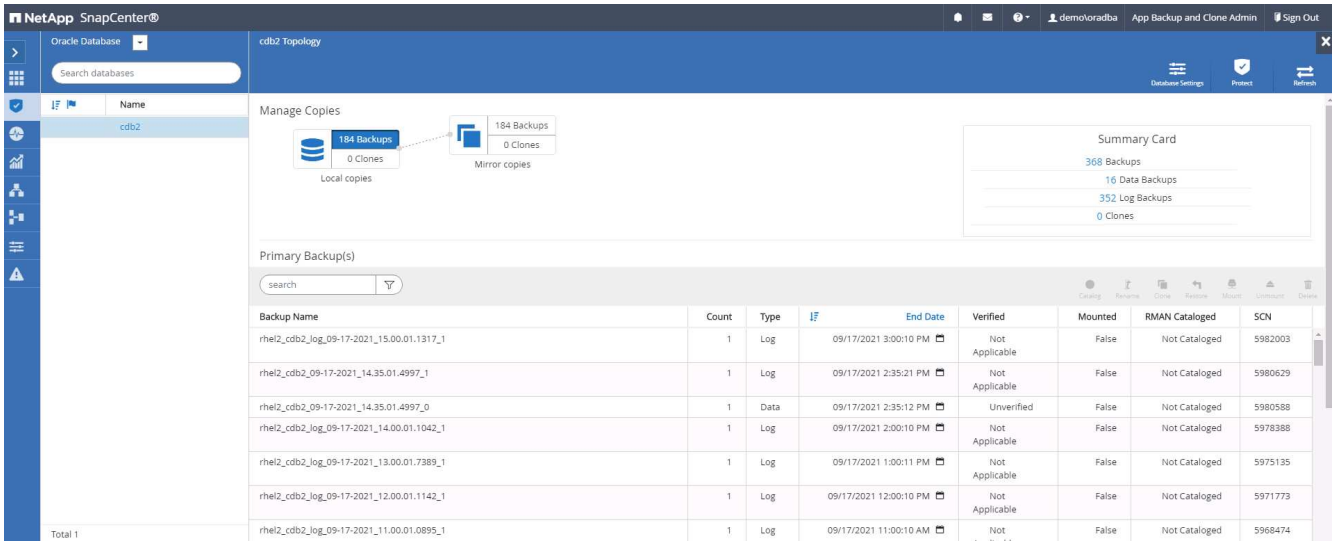

3. Toggled to the mirrored backups view by clicking mirrored backups. The secondary mirror backup(s) is then displayed.

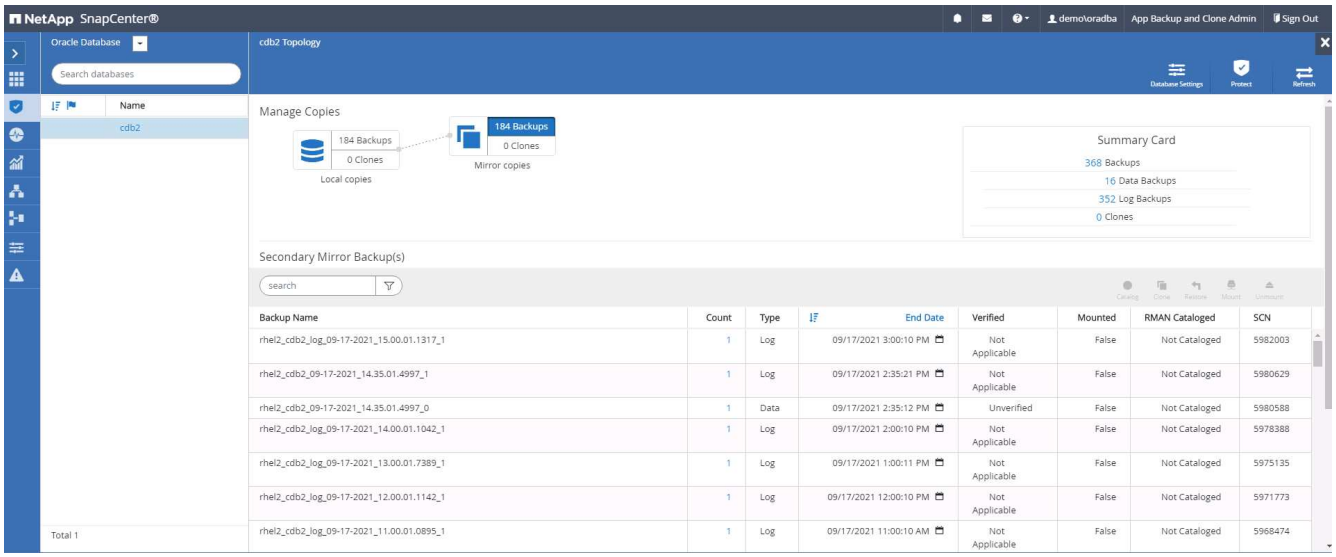

4. Choose a mirrored secondary database backup copy to be cloned and determine a recovery point either by time and system change number or by SCN. Generally, the recovery point should be trailing the full database backup time or SCN to be cloned. After a recovery point is decided, the required log file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

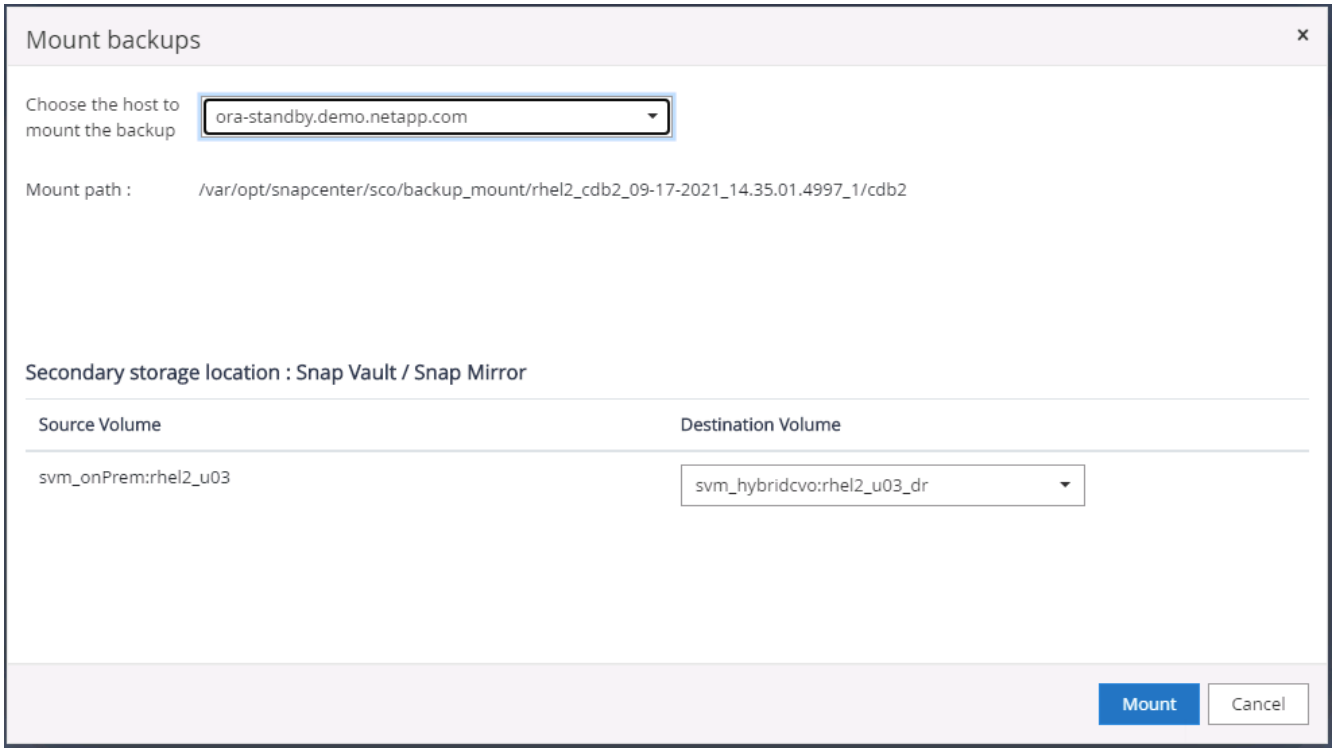

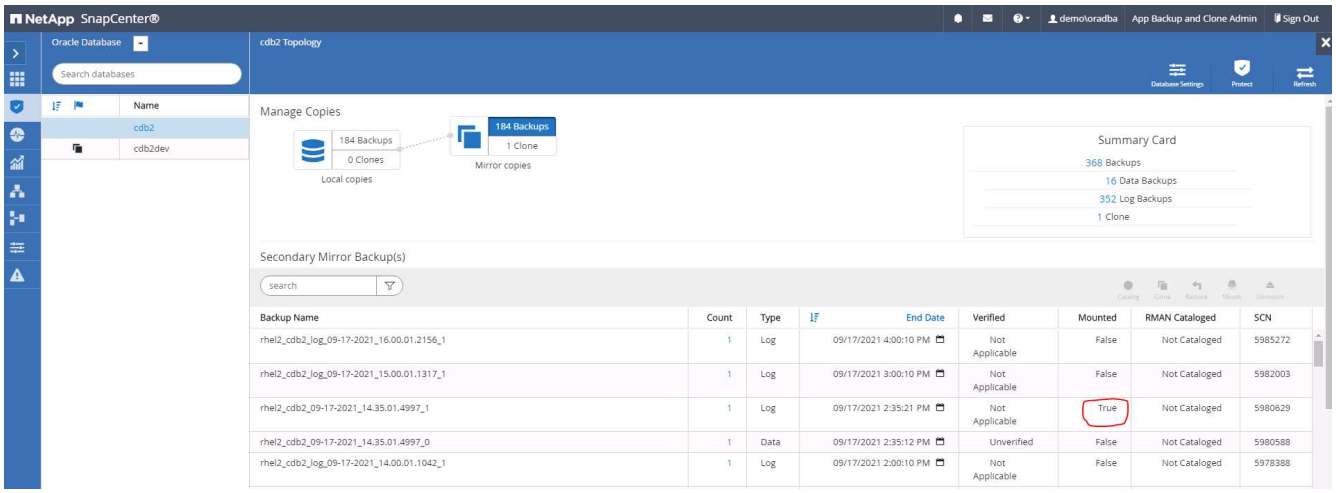

> If log pruning is enabled and the recovery point is extended beyond the last log pruning, multiple archive log backups might need to be mounted.

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

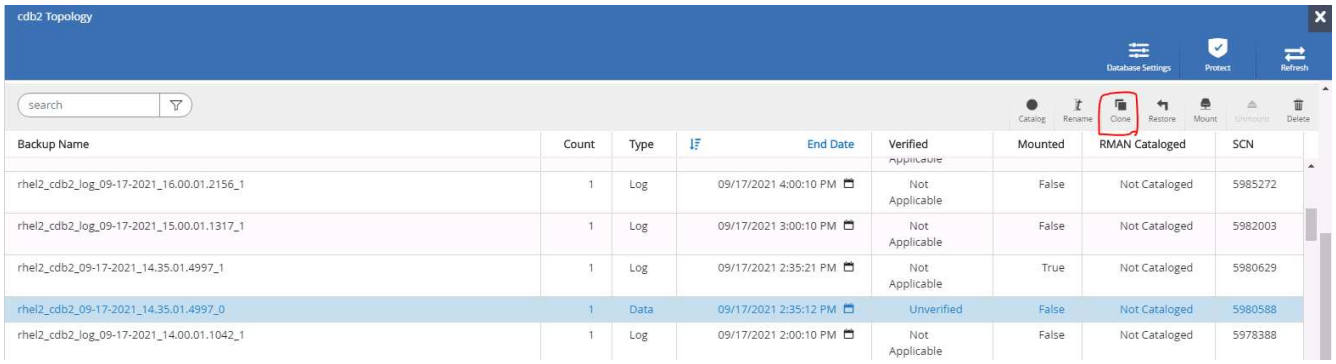

6. Choose a proper clone DB SID for a complete container database or CDB clone.

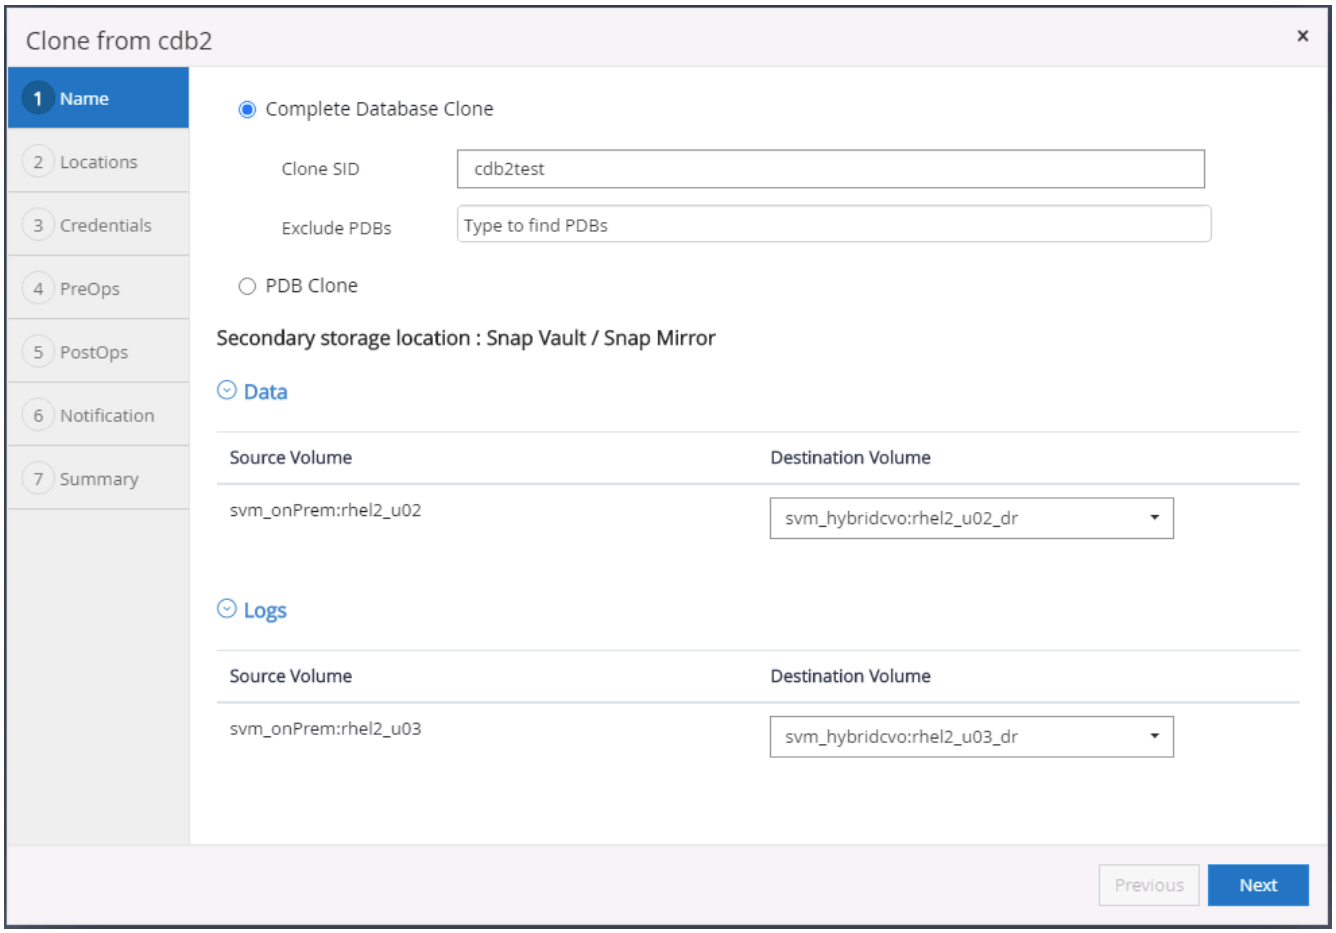

7. Select the target clone host in the cloud, and datafile, control file, and redo log directories are created by the clone workflow.

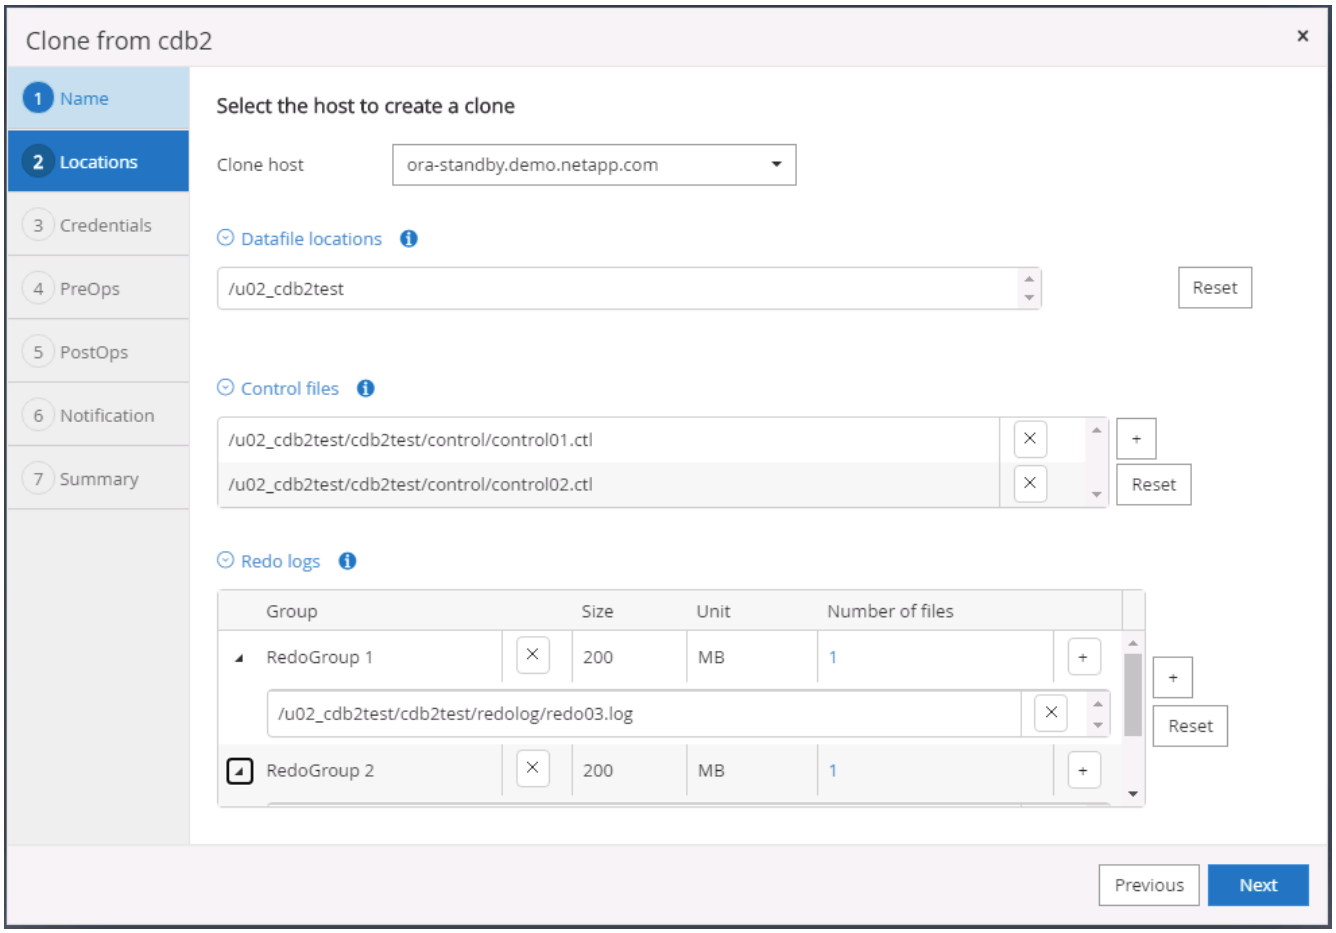

8. The None credential name is used for OS-based authentication, which renders the database port irrelevant. Fill in the proper Oracle Home, Oracle OS User, and Oracle OS Group as configured in the target clone DB server.

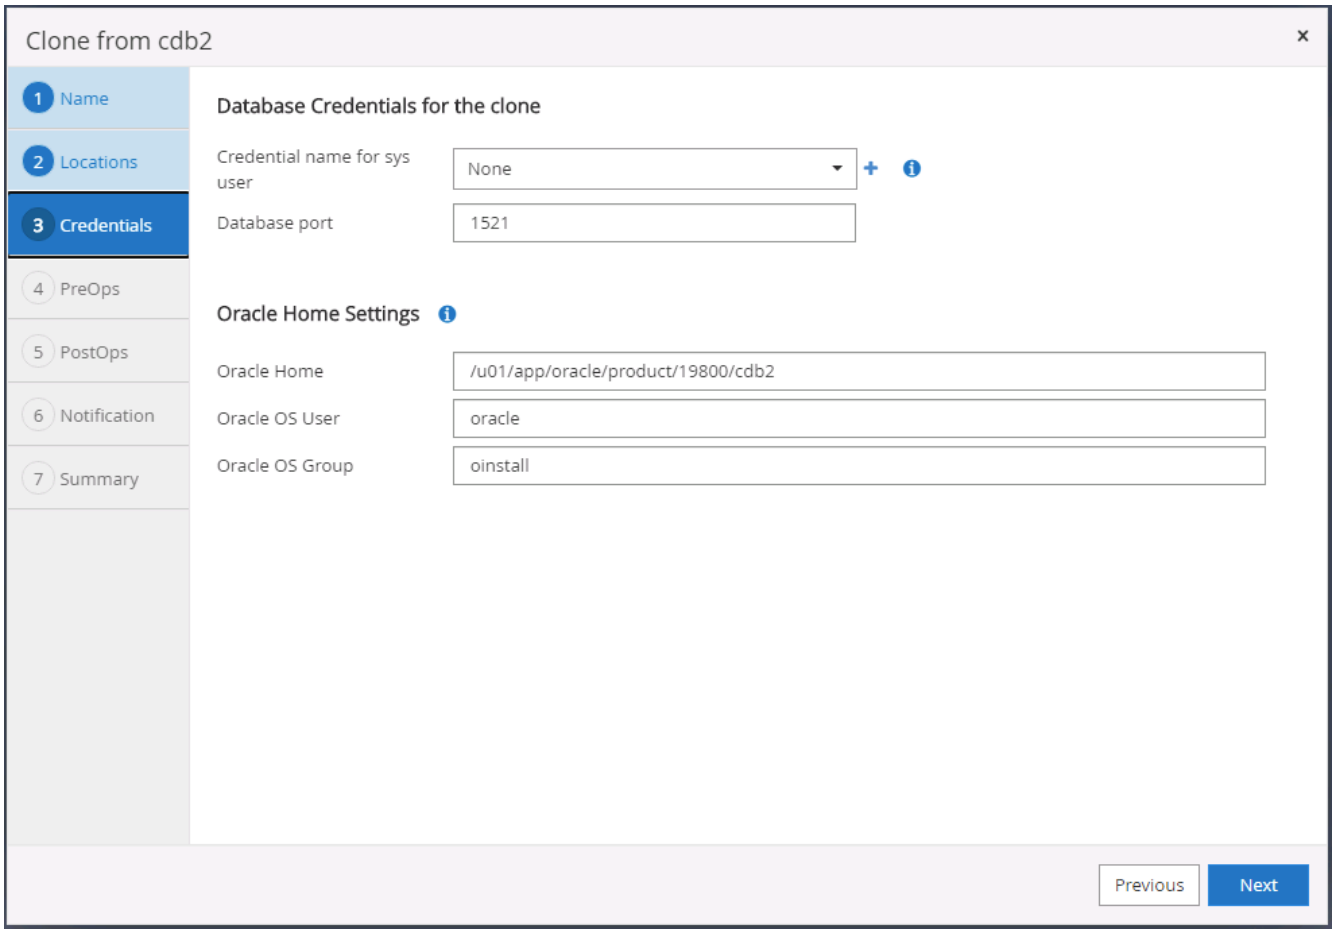

9. Specify the scripts to run before clone operation. More importantly, the database instance parameter can be adjusted or defined here.

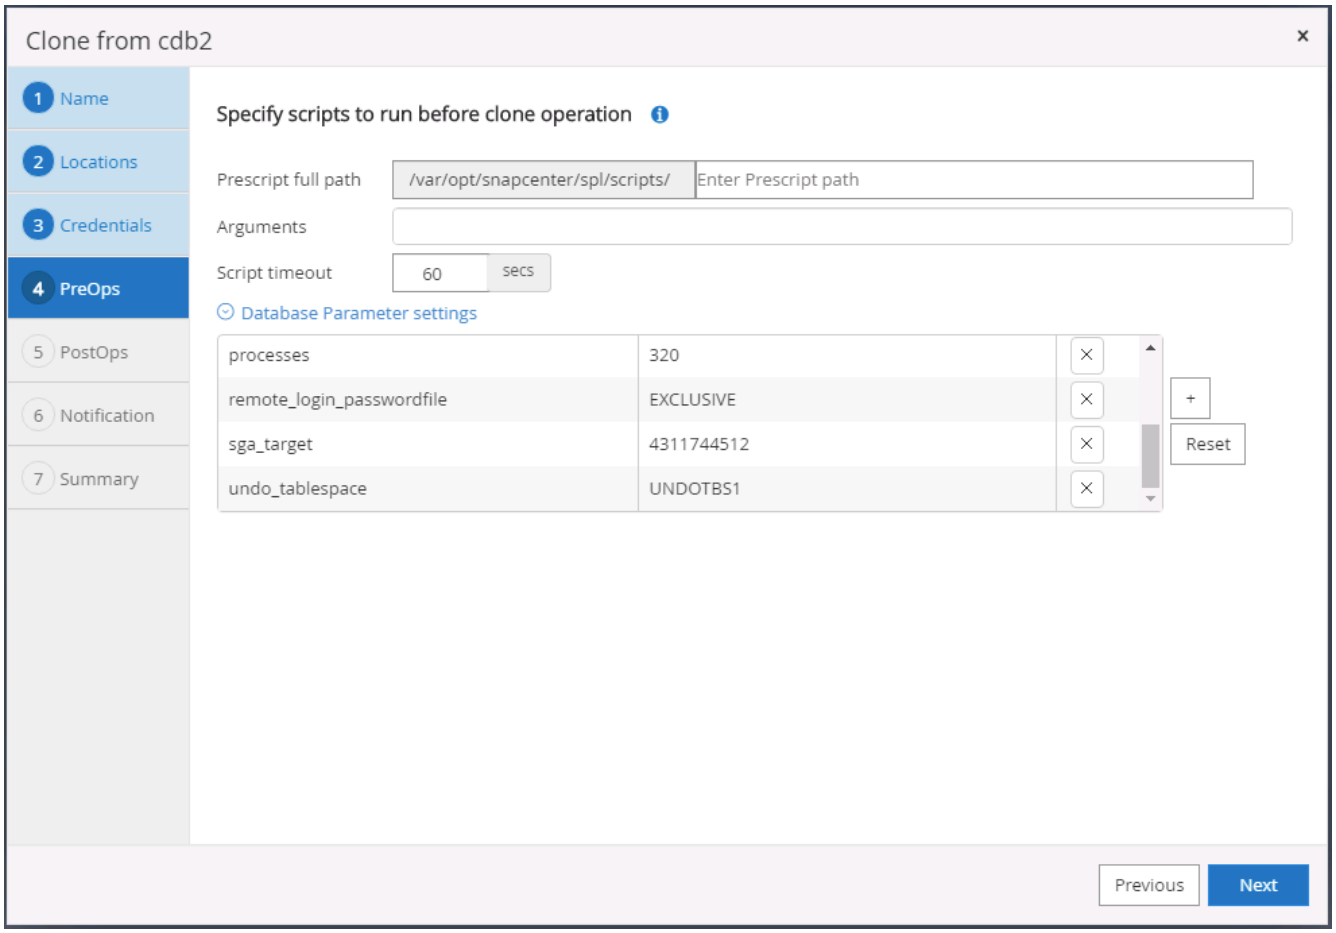

10. Specify the recovery point either by the date and time or SCN. Until Cancel recovers the database up to the available archive logs. Specify the external archive log location from the target host where the archive log volume is mounted. If target server Oracle owner is different from the on-premises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

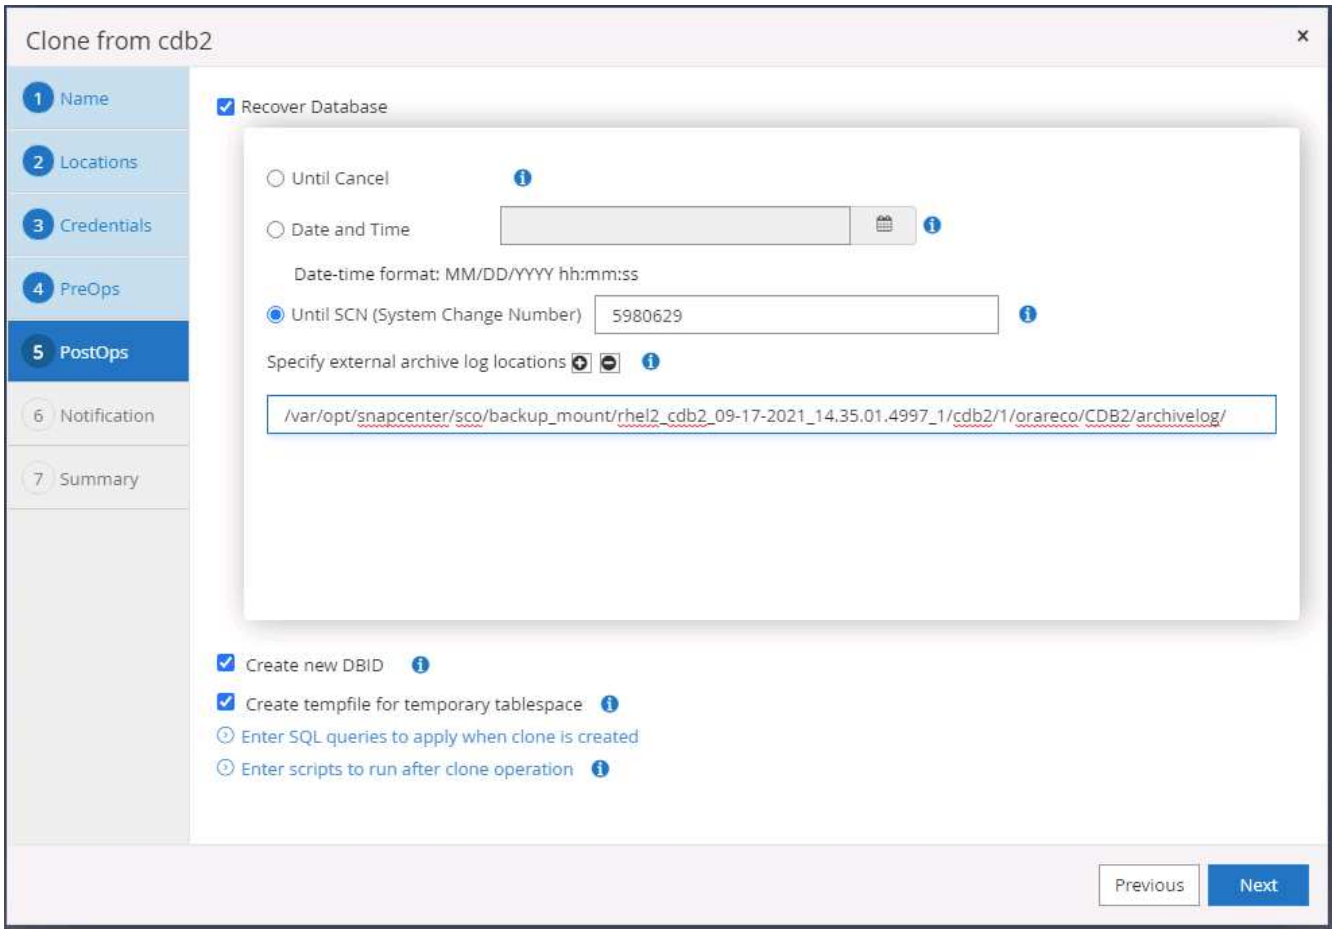

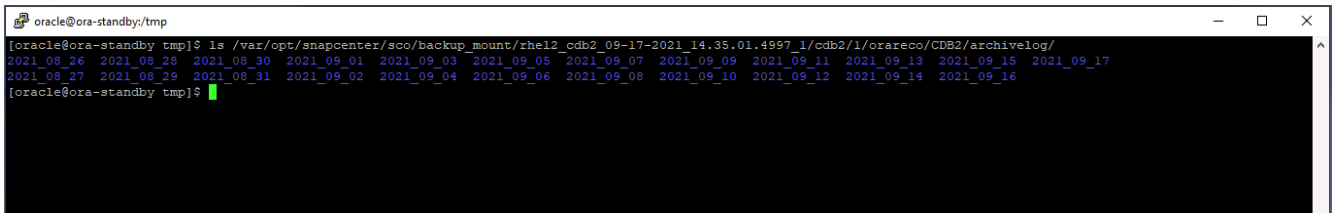

11. Configure the SMTP server for email notification if desired.

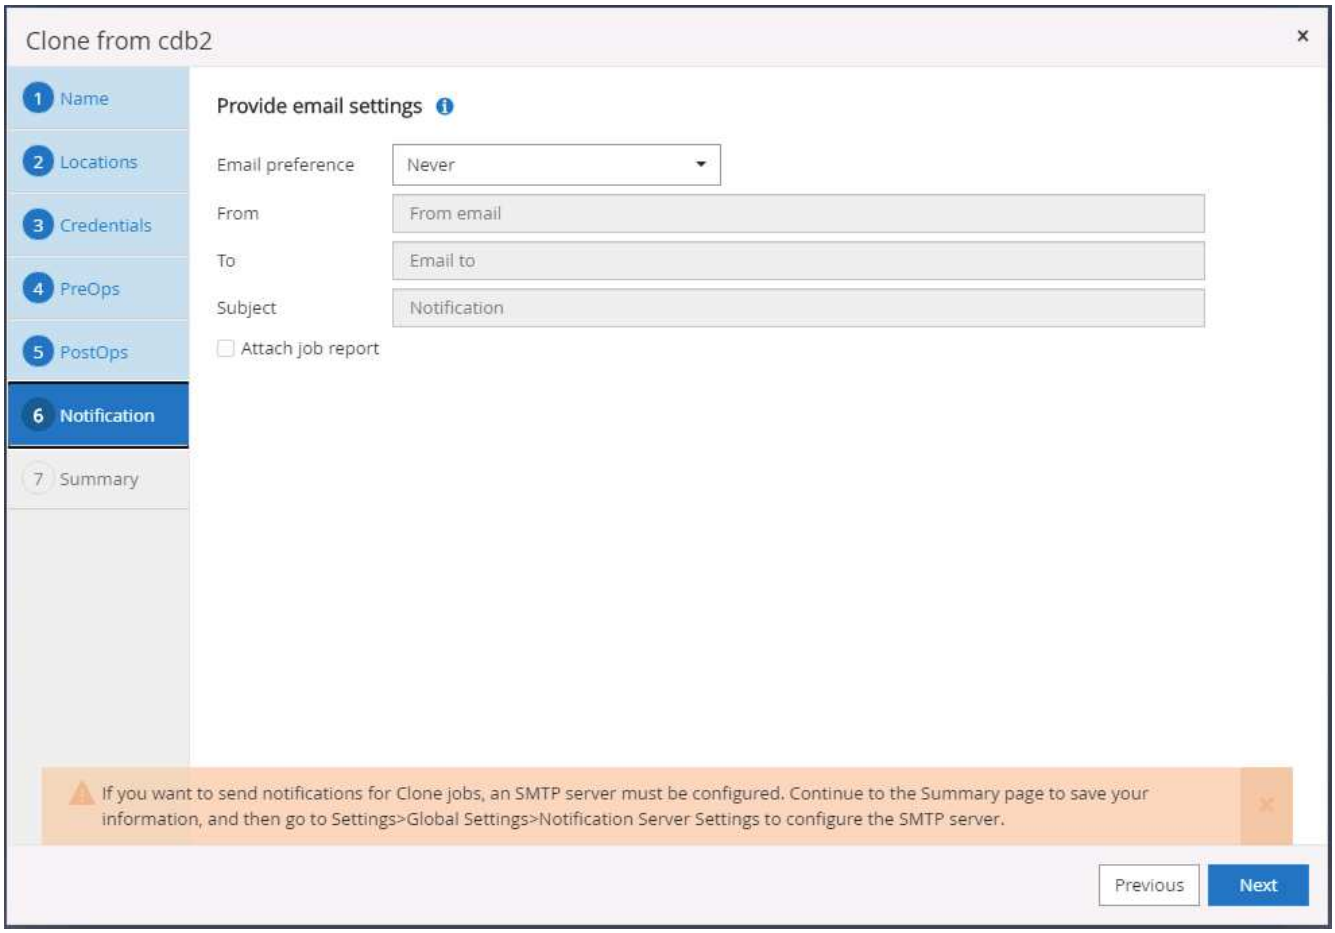

12. Clone summary.

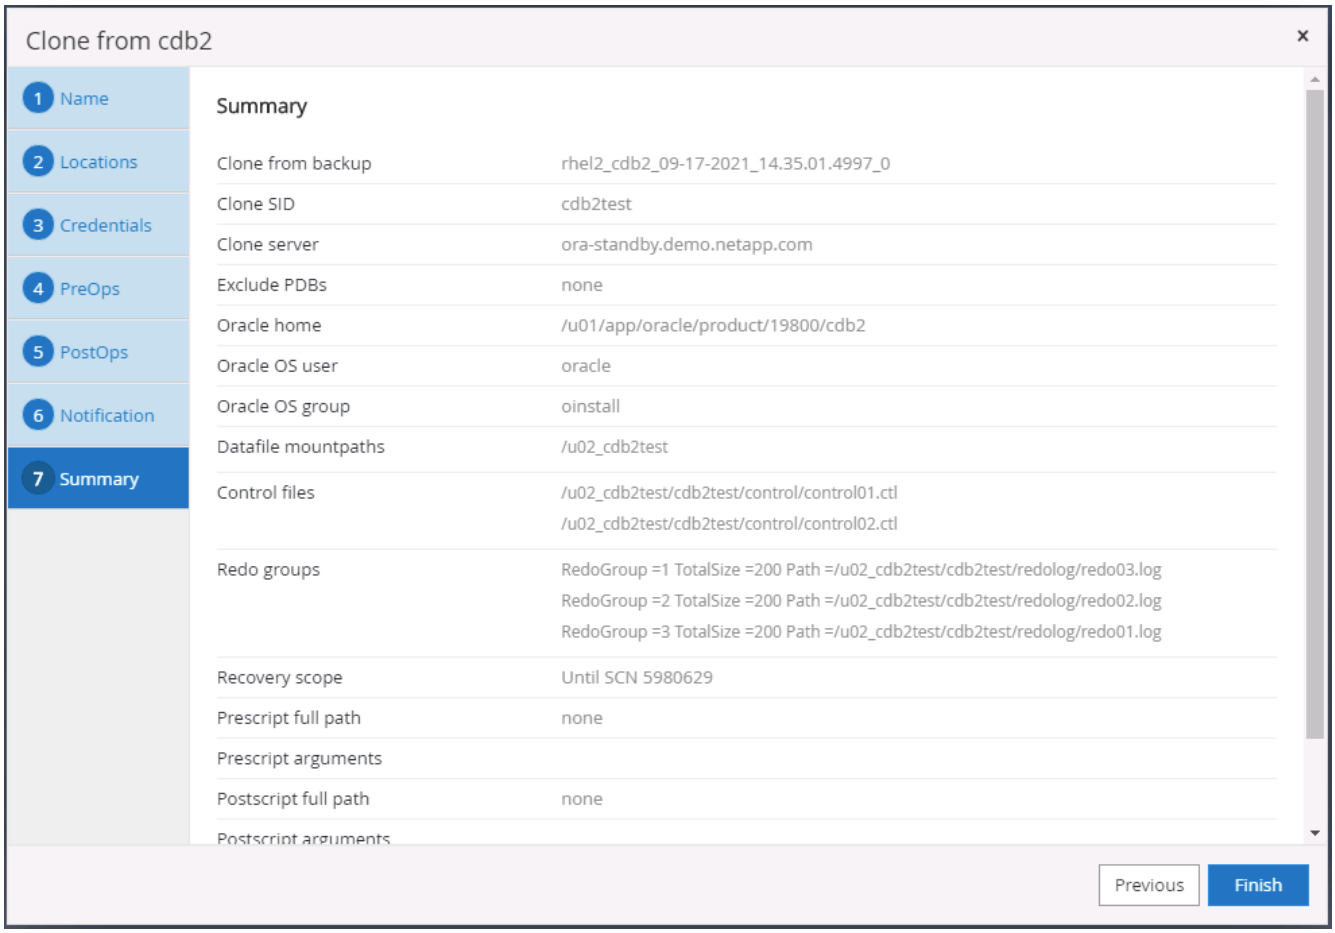

13. You should validate after cloning to make sure that the cloned database is operational. Some additional tasks, such as starting up the listener or turning off the DB log archive mode, can be performed on the dev/test database.

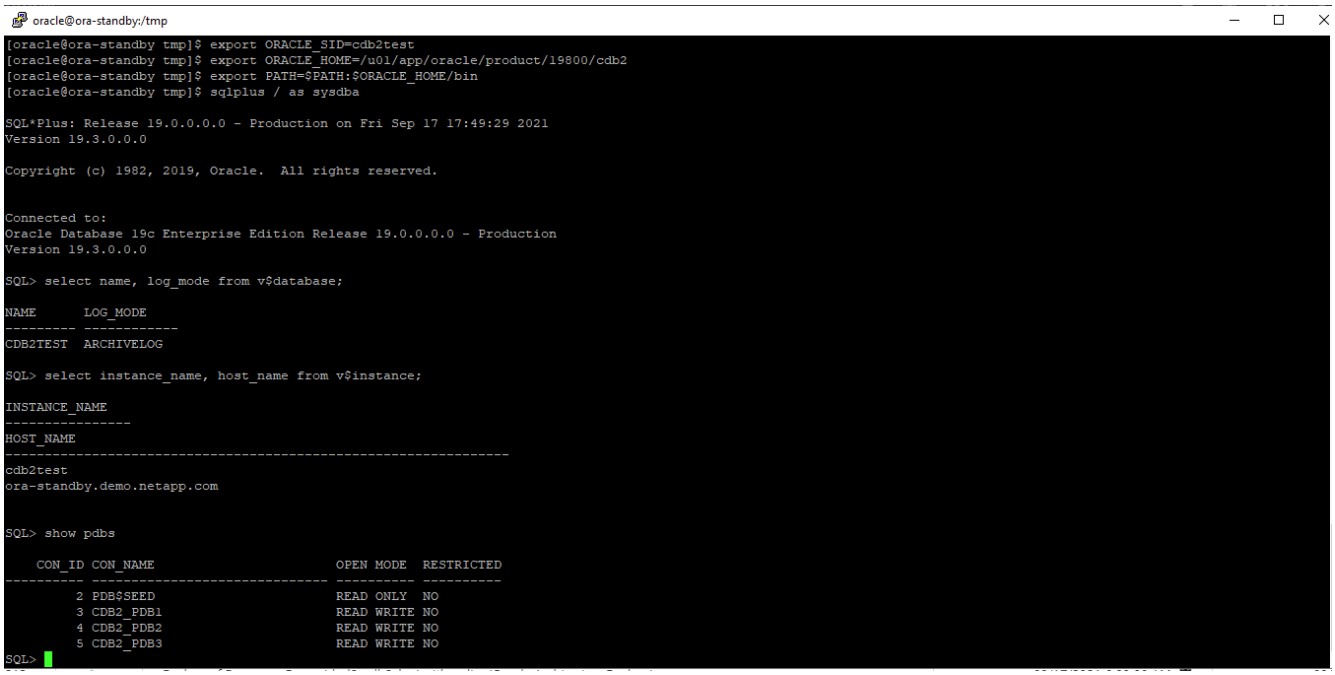

Clone a SQL database for dev/test from a replicated Snapshot backup

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.

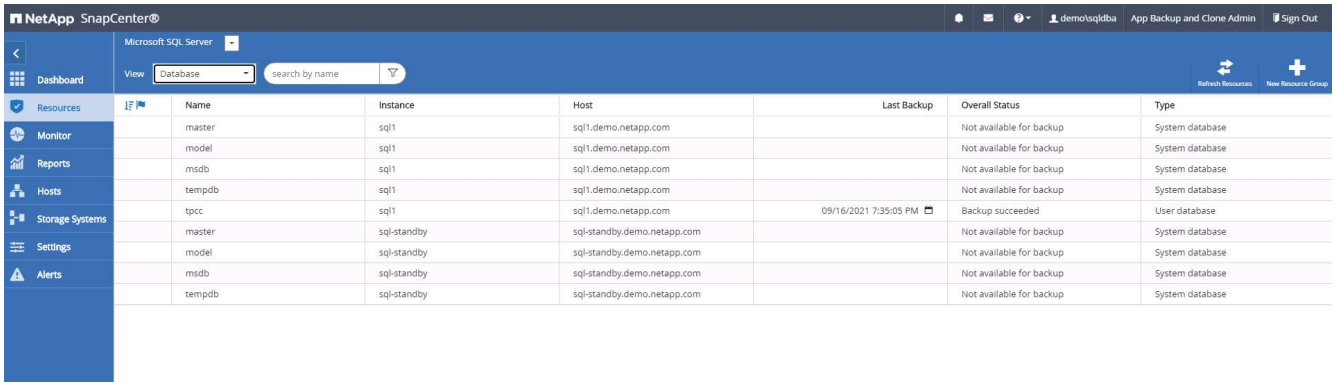

2. Click on the intended on-premises SQL Server user database name for the backups topology and detailed view. If a secondary replicated location is enabled, it shows linked mirror backups.

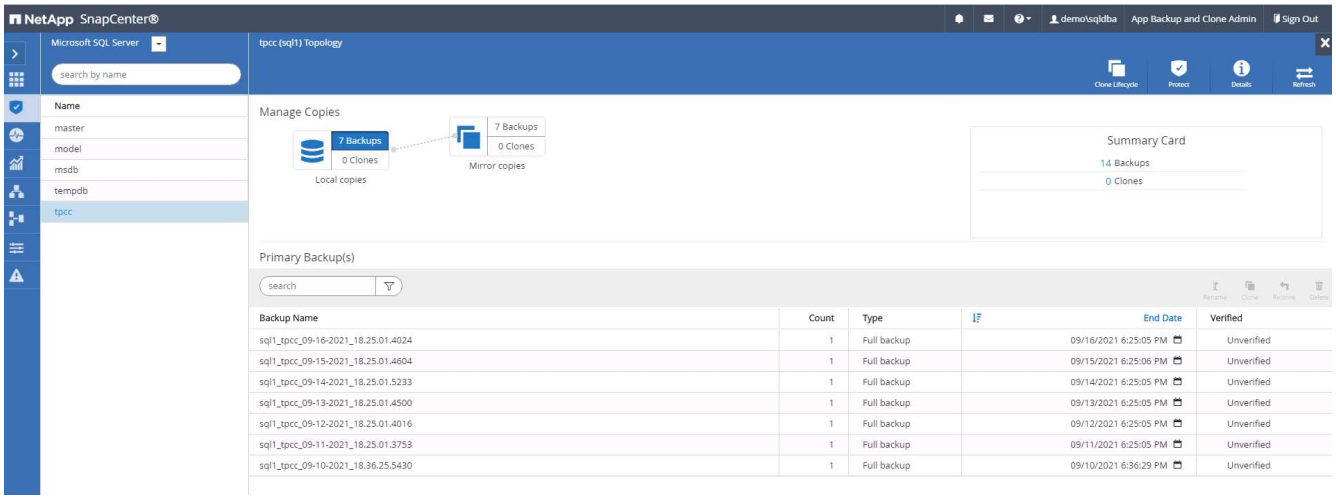

3. Toggle to the Mirrored Backups view by clicking Mirrored Backups. Secondary Mirror Backup(s) are then displayed. Because SnapCenter backs up the SQL Server transaction log to a dedicated drive for recovery, only full database backups are displayed here.

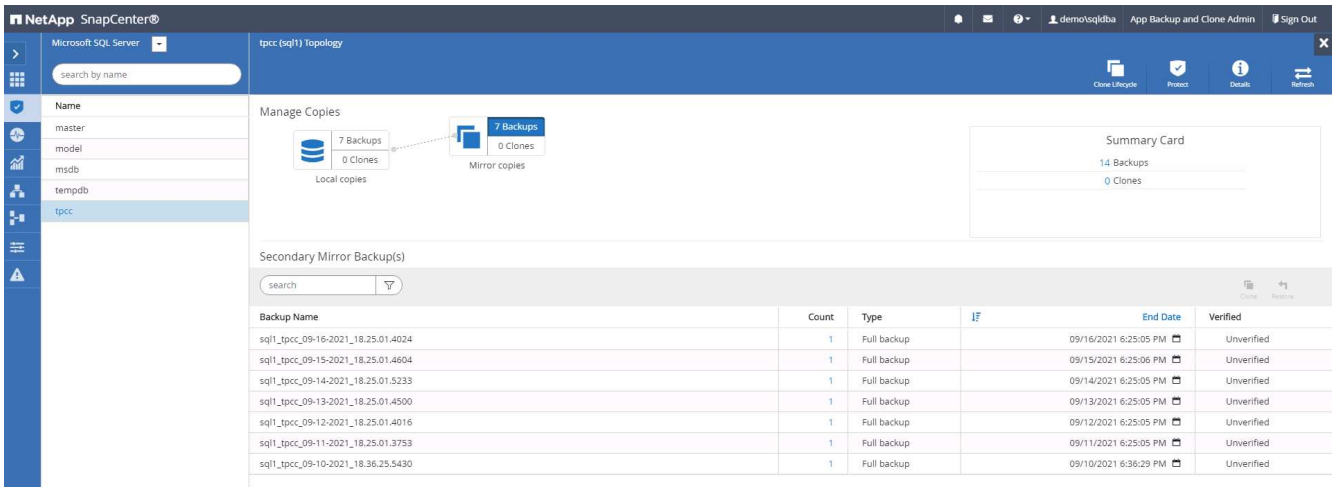

4. Choose a backup copy, and then click the Clone button to launch the Clone from Backup workflow.

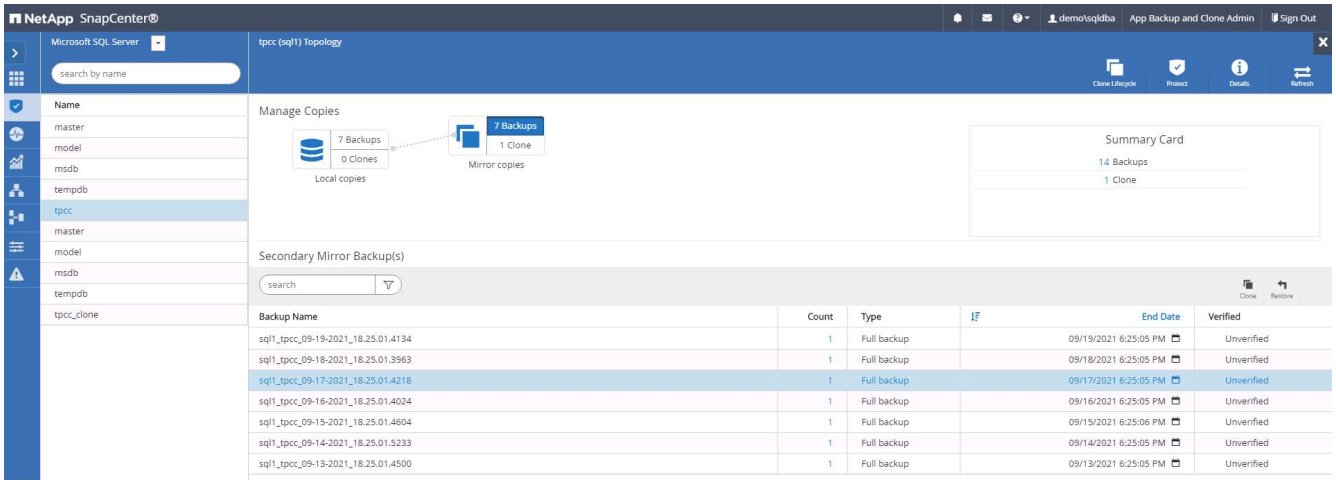

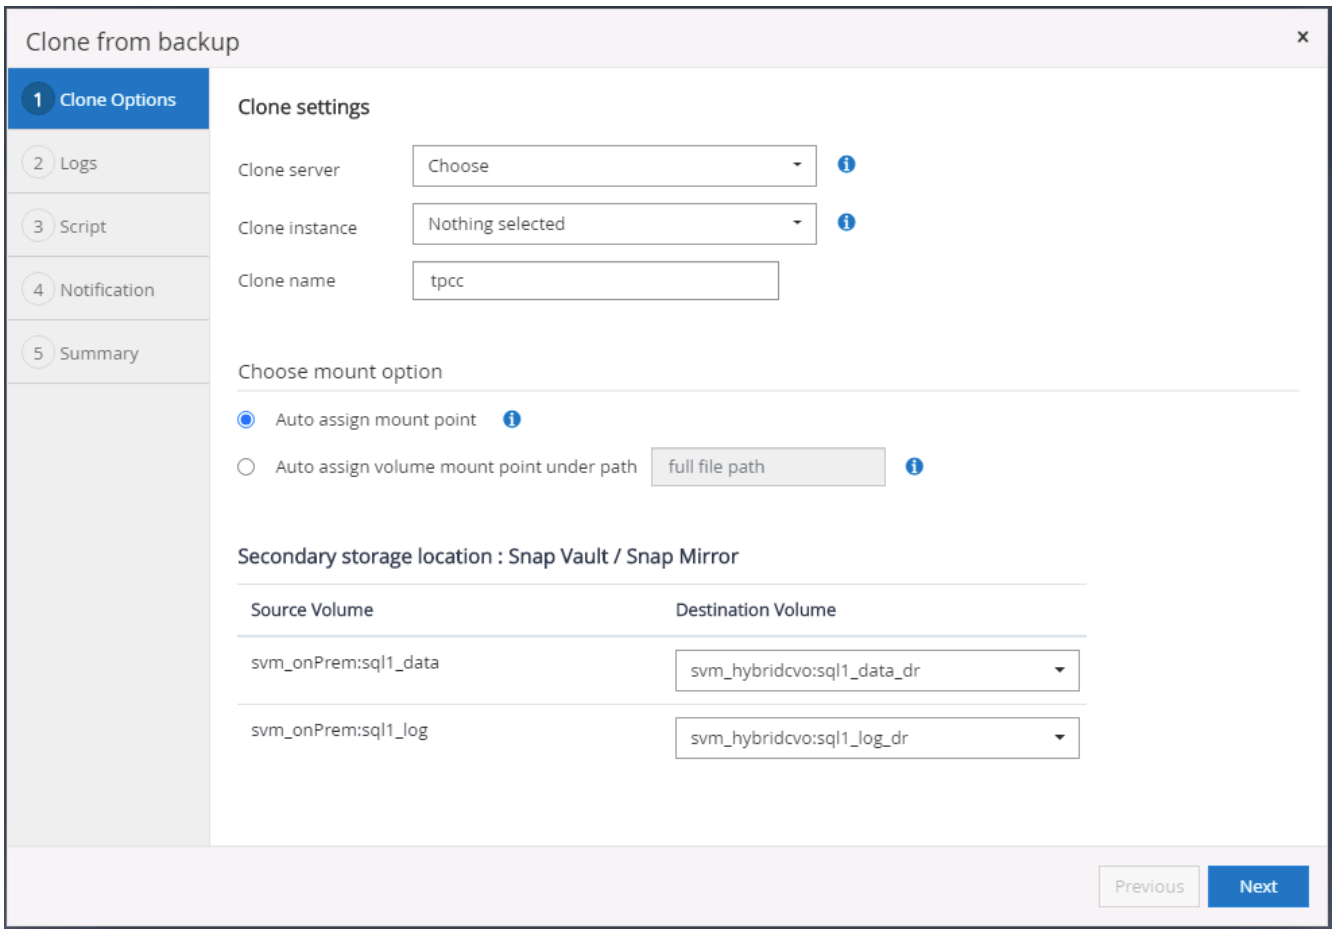

5. Select a cloud server as the target clone server, clone instance name, and clone database name. Choose either an auto-assign mount point or a user-defined mount point path.

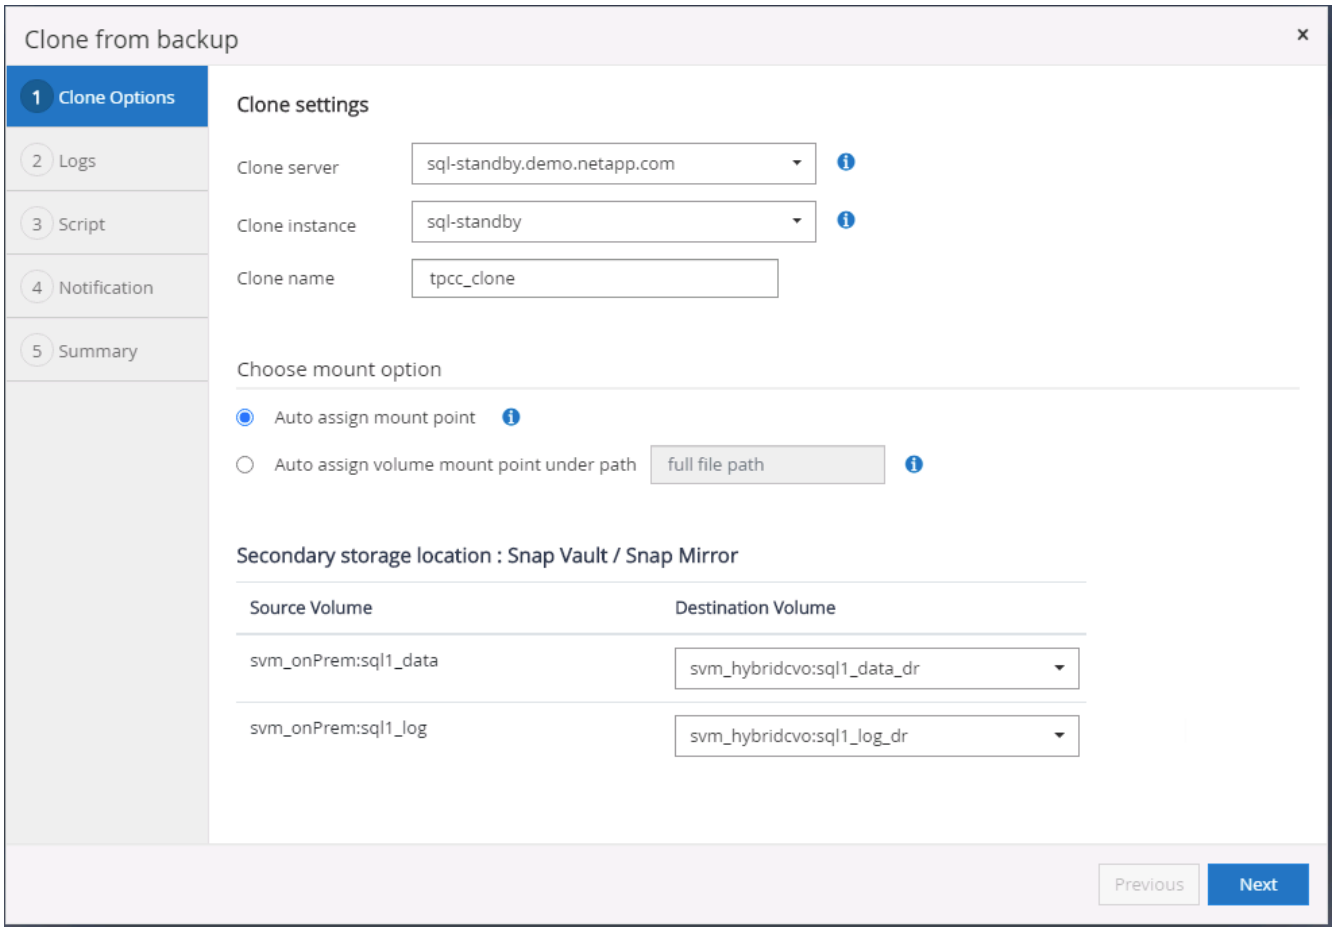

6. Determine a recovery point either by a log backup time or by a specific date and time.

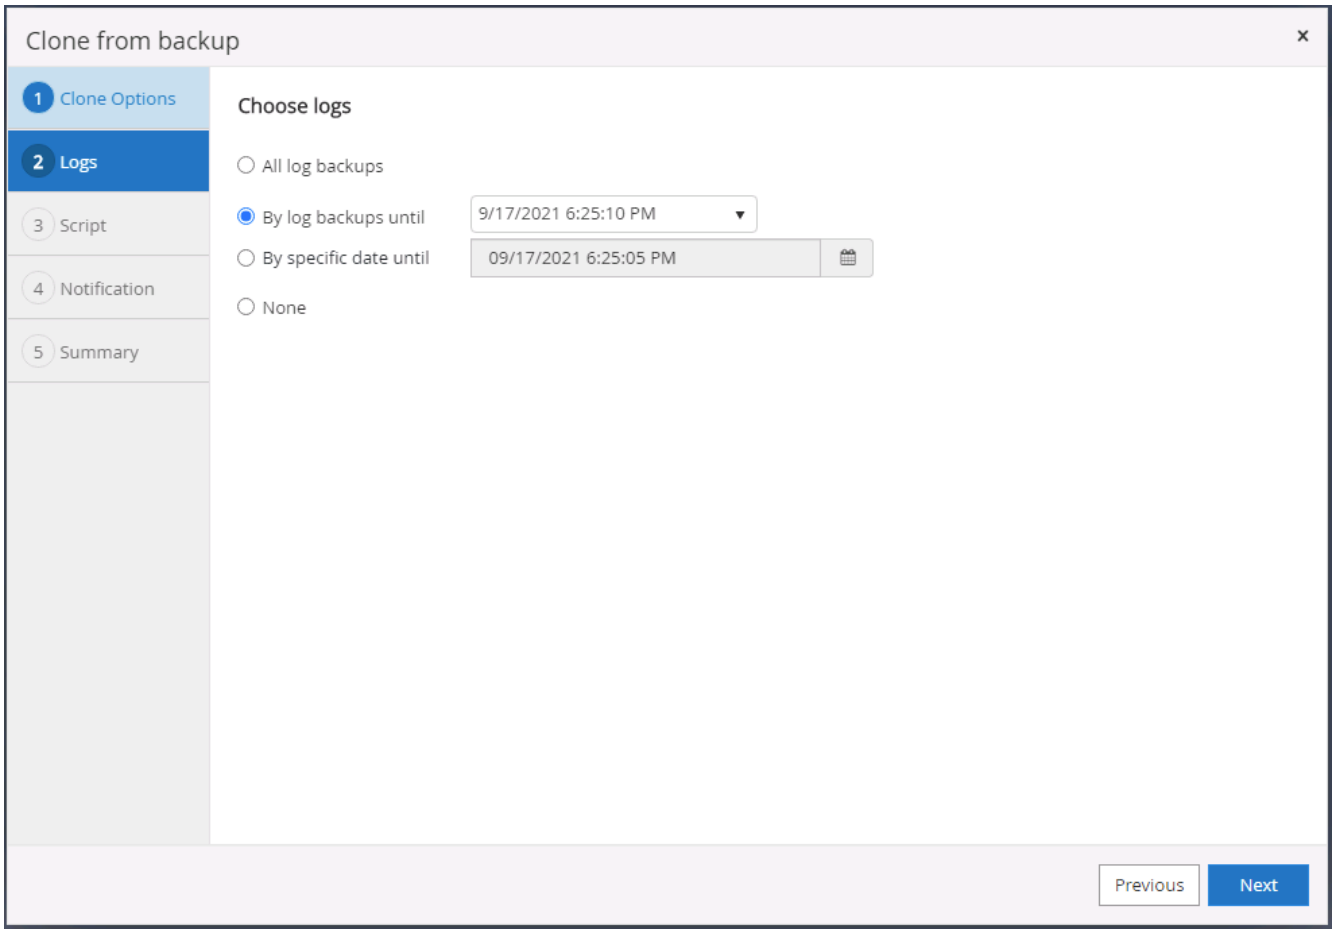

7. Specify optional scripts to run before and after the cloning operation.

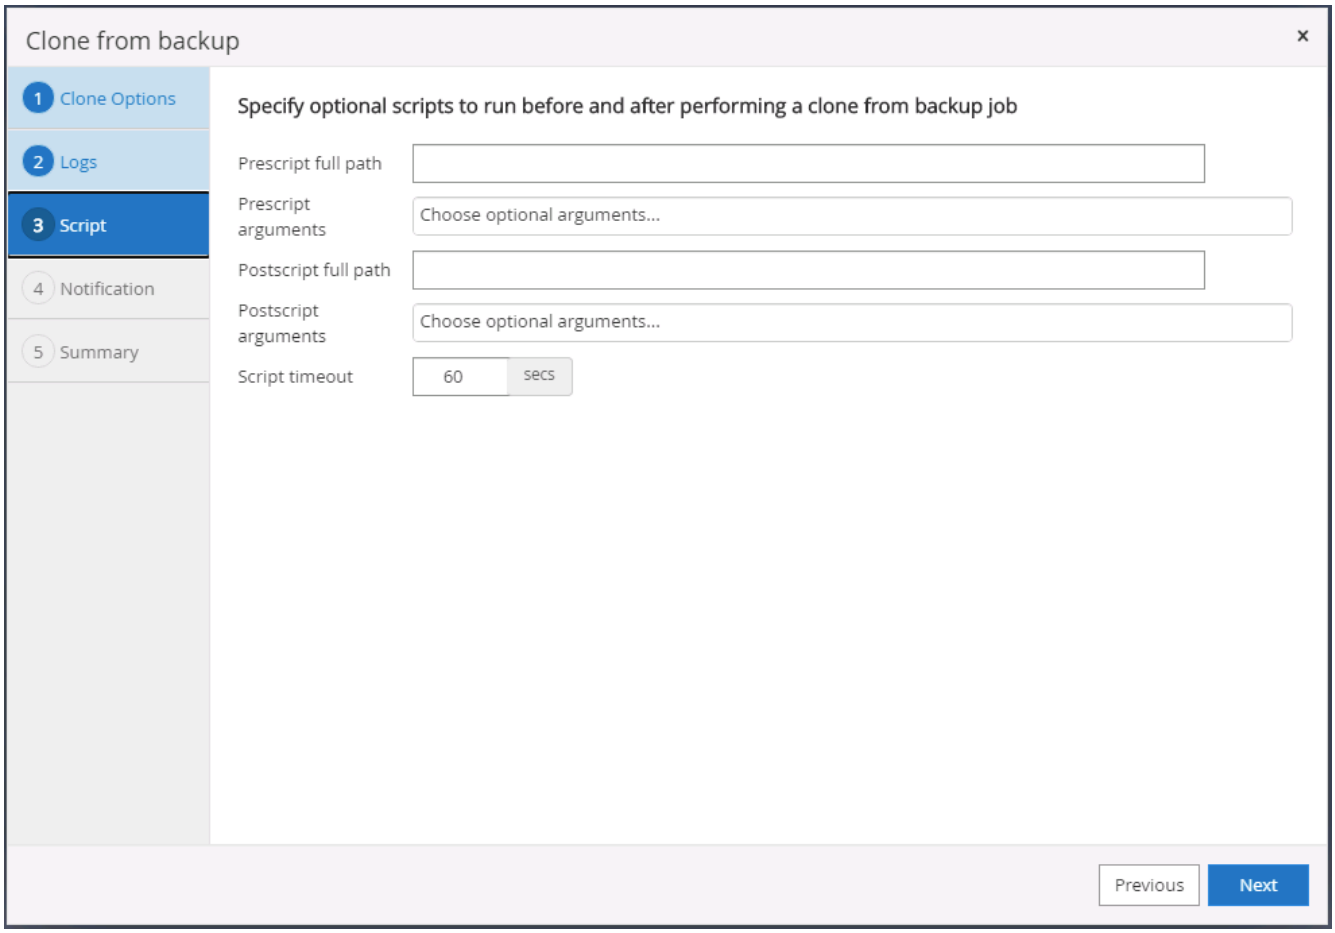

8. Configure an SMTP server if email notification is desired.

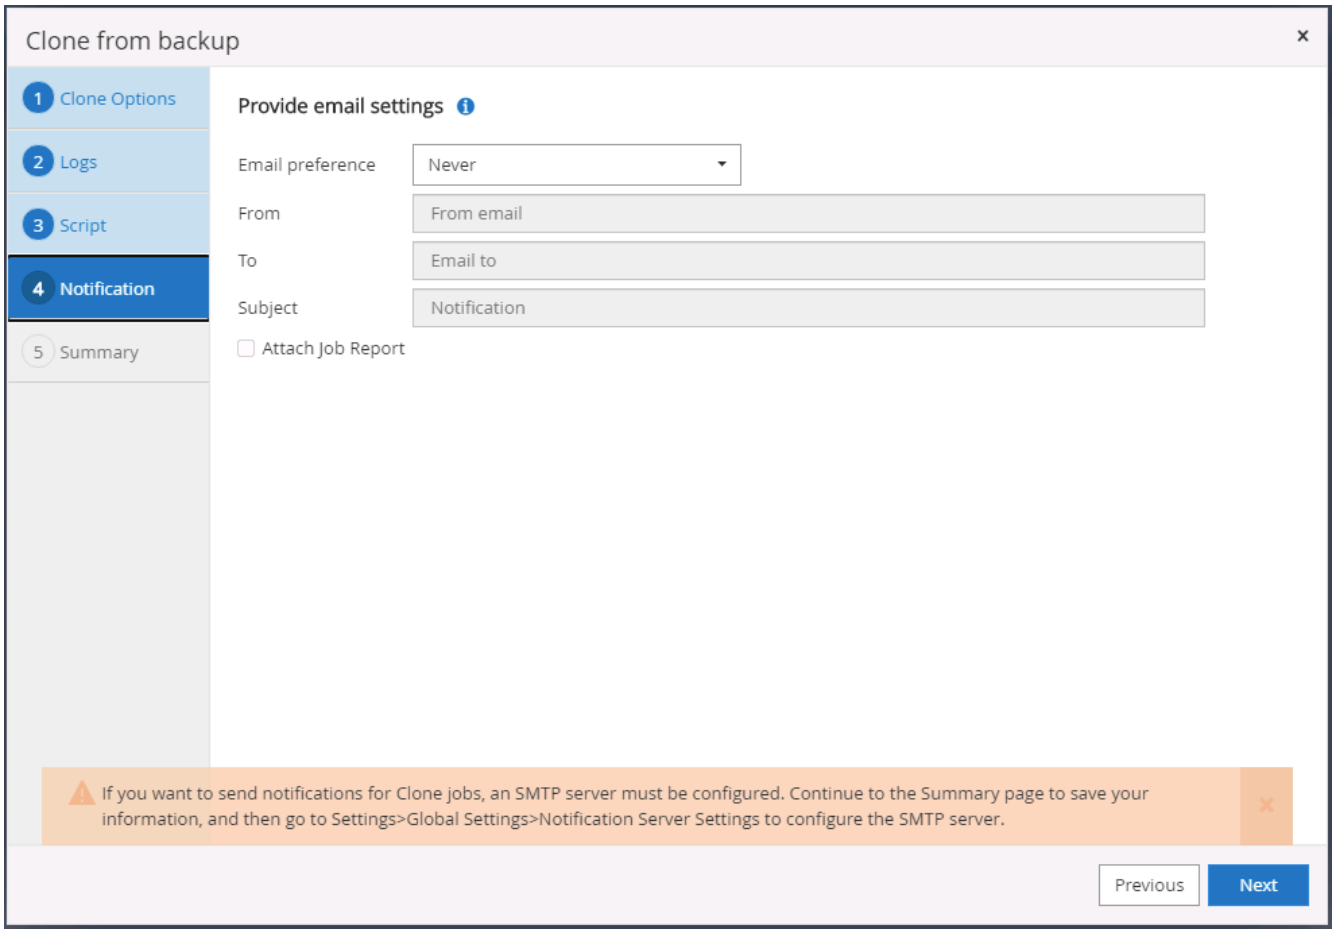

9. Clone Summary.

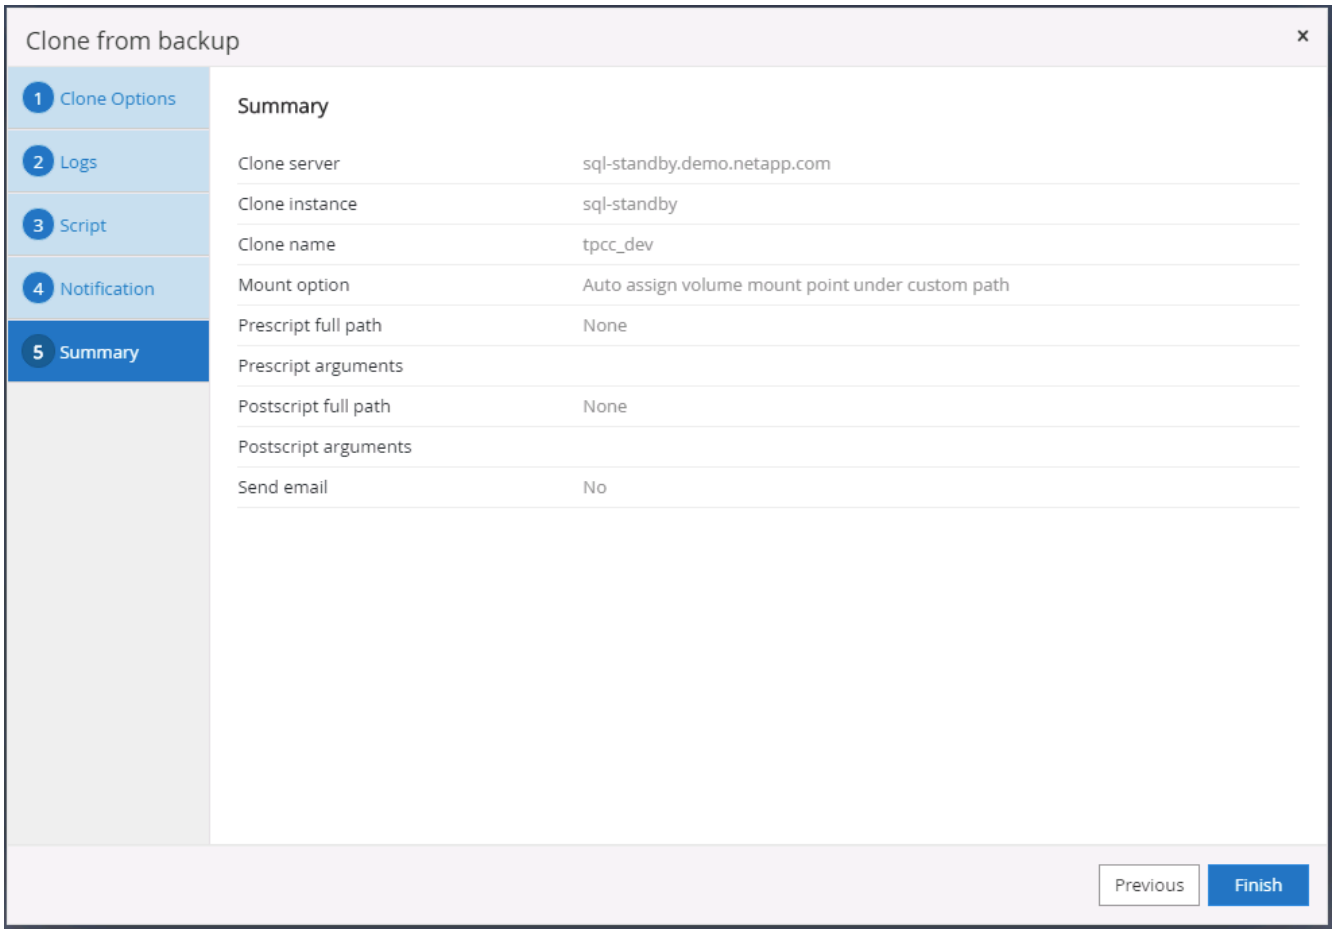

10. Monitor the job status and validate that the intended user database has been attached to a target SQL instance in the cloud clone server.

Post-clone configuration

1. An Oracle production database on-premises is usually running in log archive mode. This mode is not necessary for a development or test database. To turn off log archive mode, log into the Oracle DB as sysdba, execute a log mode change command, and start the database for access.
2. Configure an Oracle listener, or register the newly cloned DB with an existing listener for user access.
3. For SQL Server, change the log mode from Full to Easy so that the SQL Server dev/test log file can be readily shrunk when it is filling up the log volume.

Refresh clone database

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

Where to go for help?

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

Disaster recovery workflow

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

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

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

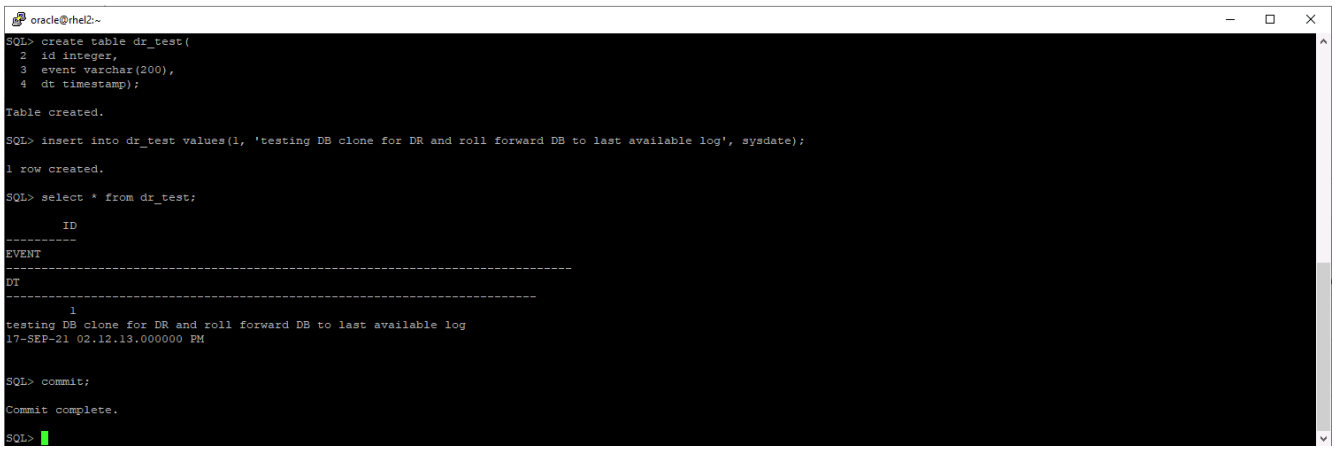

2. Log into SnapCenter as a database management user ID for Oracle. Navigate to the Resources tab, which shows the Oracle databases being protected by SnapCenter.

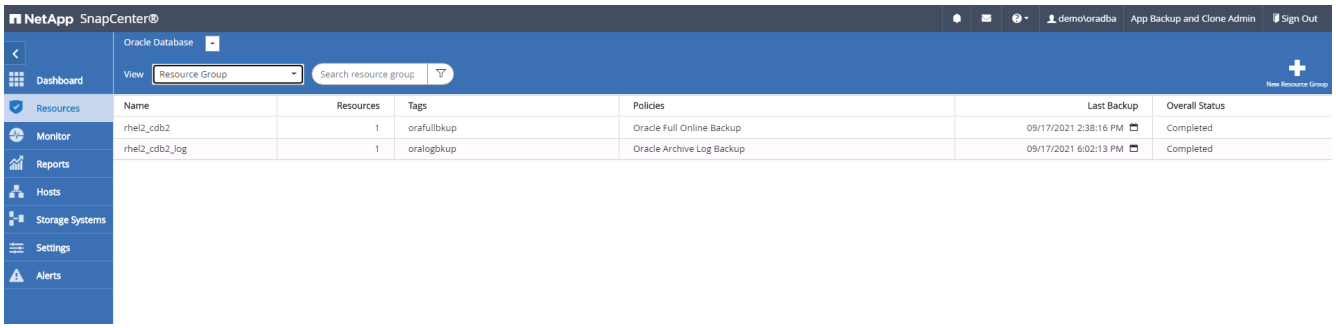

3. Select the Oracle log resource group and click Backup Now to manually run an Oracle log backup to flush the latest transaction to the destination in the cloud. In a real DR scenario, the last transaction recoverable depends on the database log volume replication frequency to the cloud, which in turn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

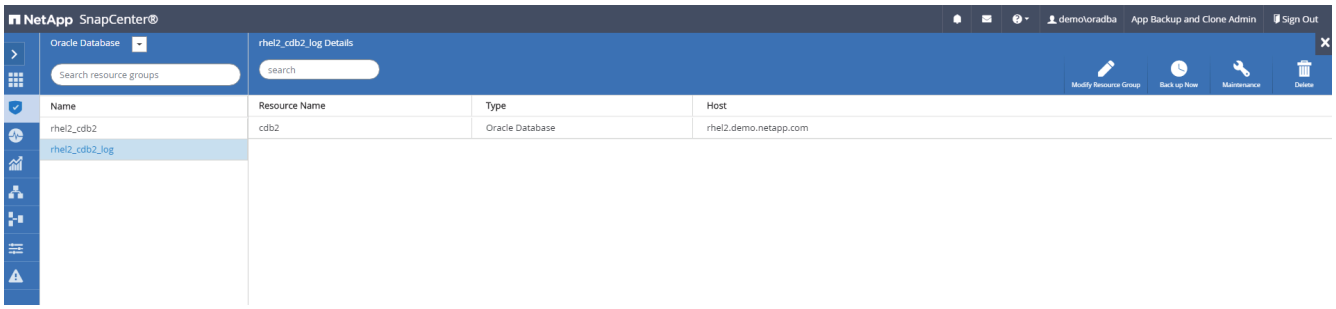

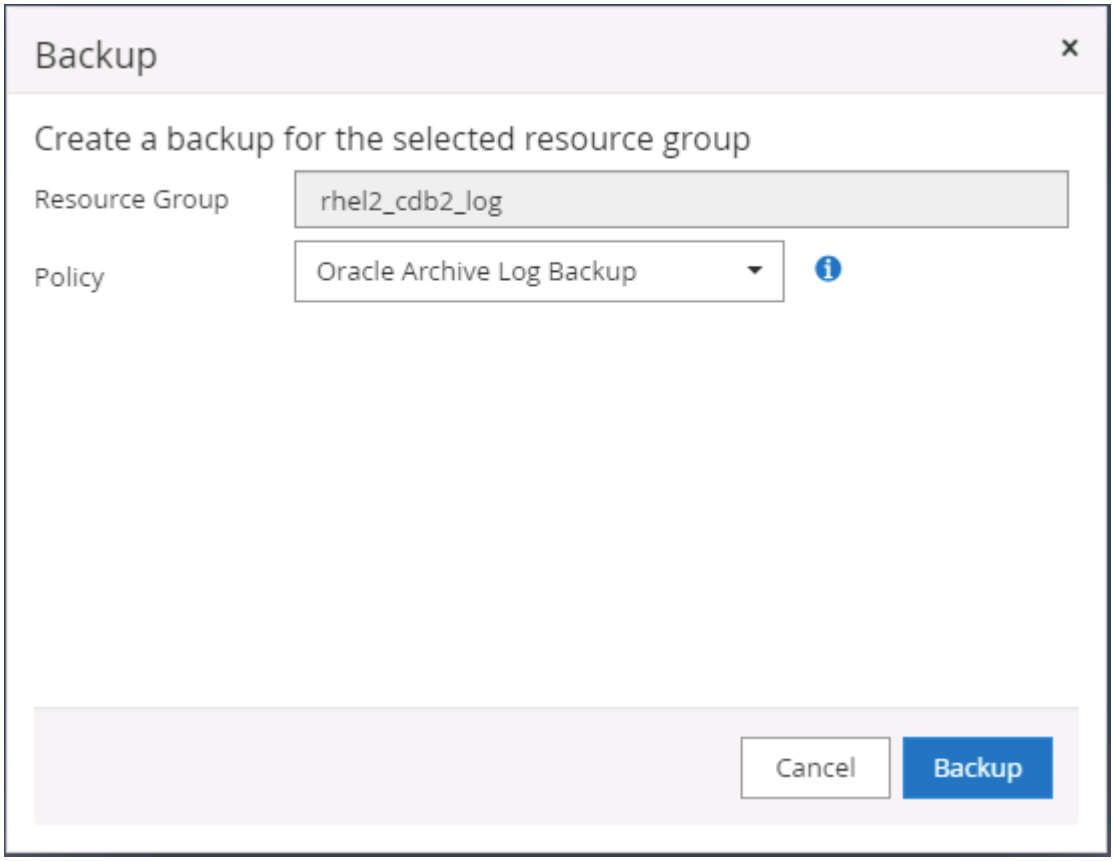

> Asynchronous SnapMirror loses data that has not made it to the cloud destination in the database log backup interval in a disaster recovery scenario. To minimize data loss, more frequent log backup can be scheduled. However there is a limit to the log backup frequency that is technically achievable.

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

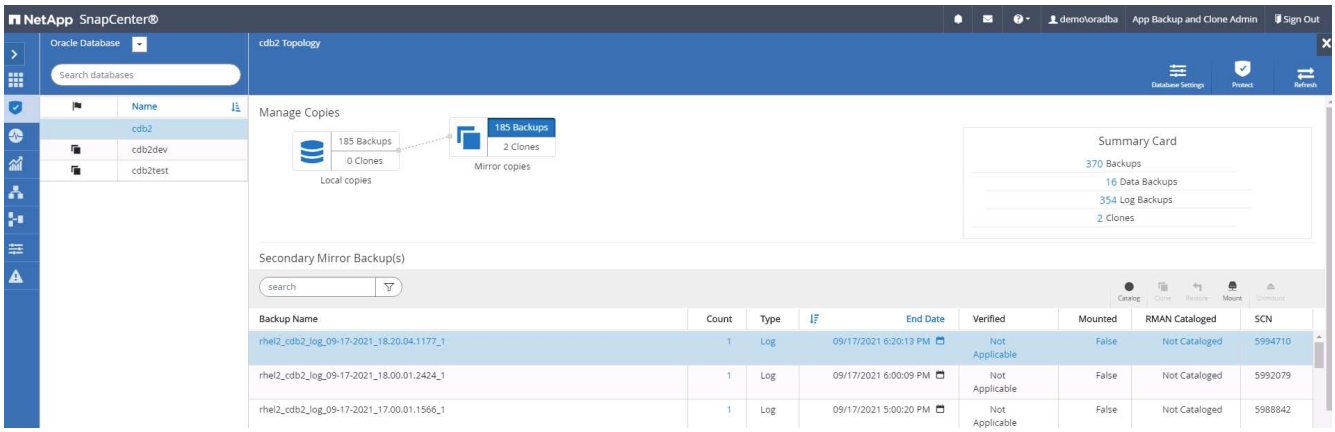

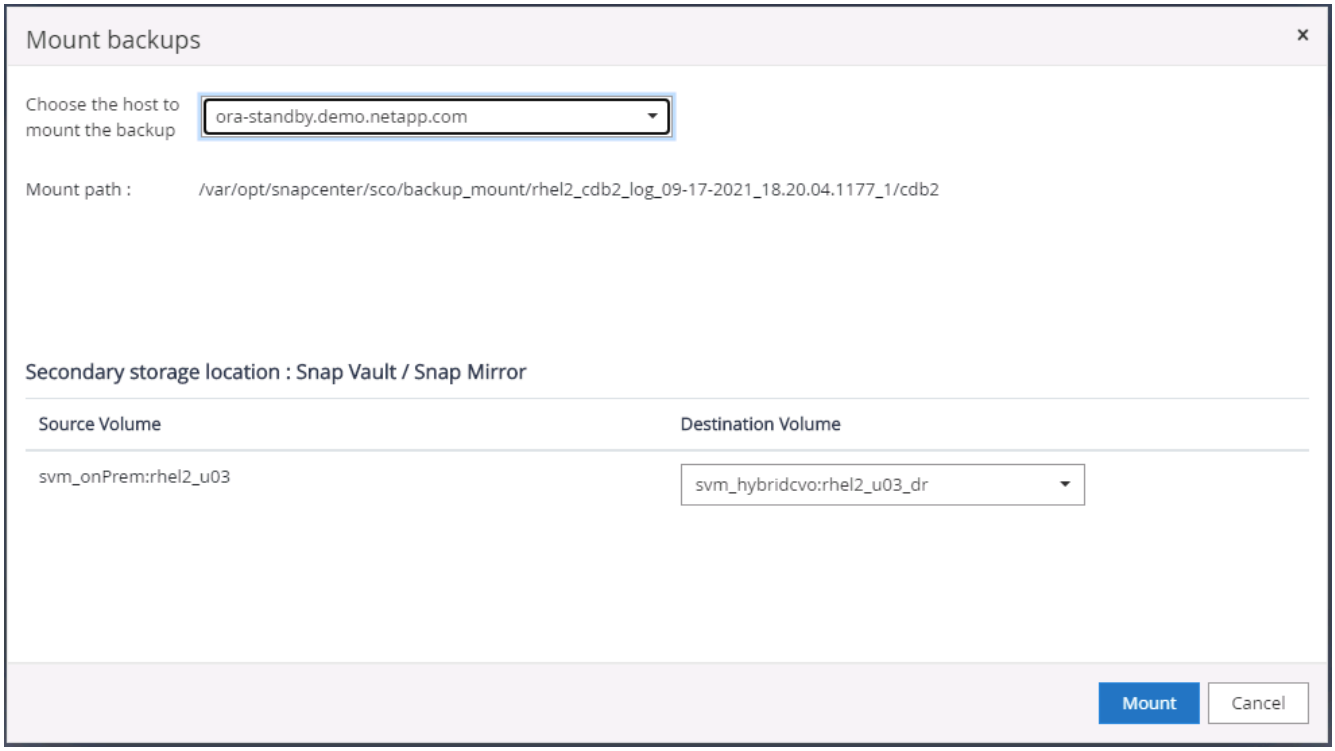

5. Select the last full database backup and click Clone to initiate the clone workflow.

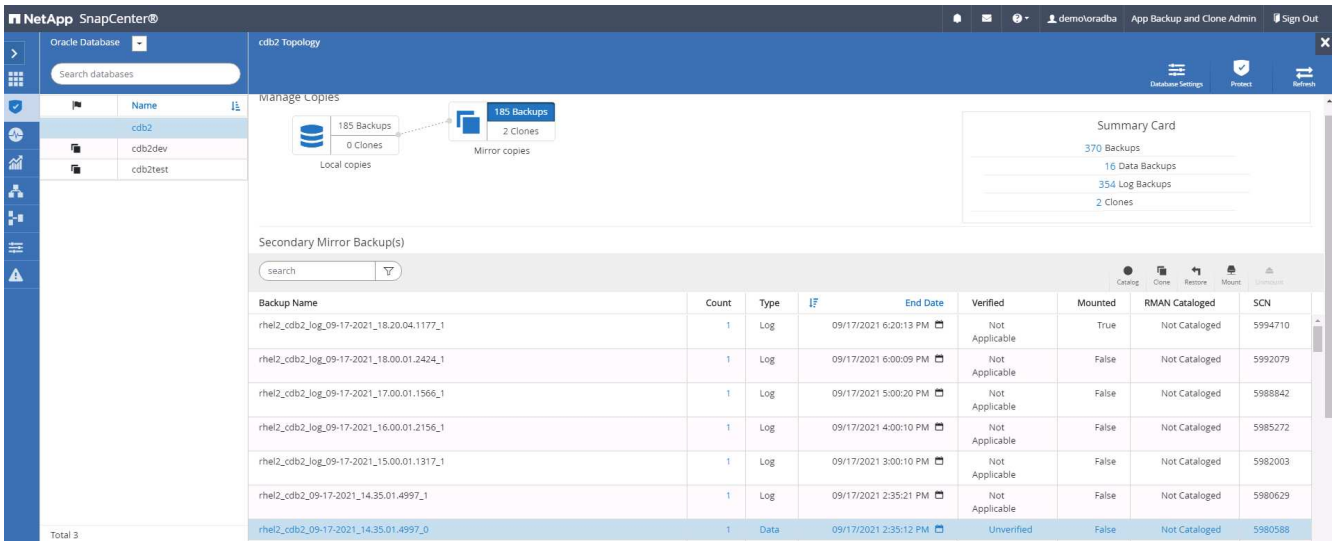

6. Select a unique clone DB ID on the host.

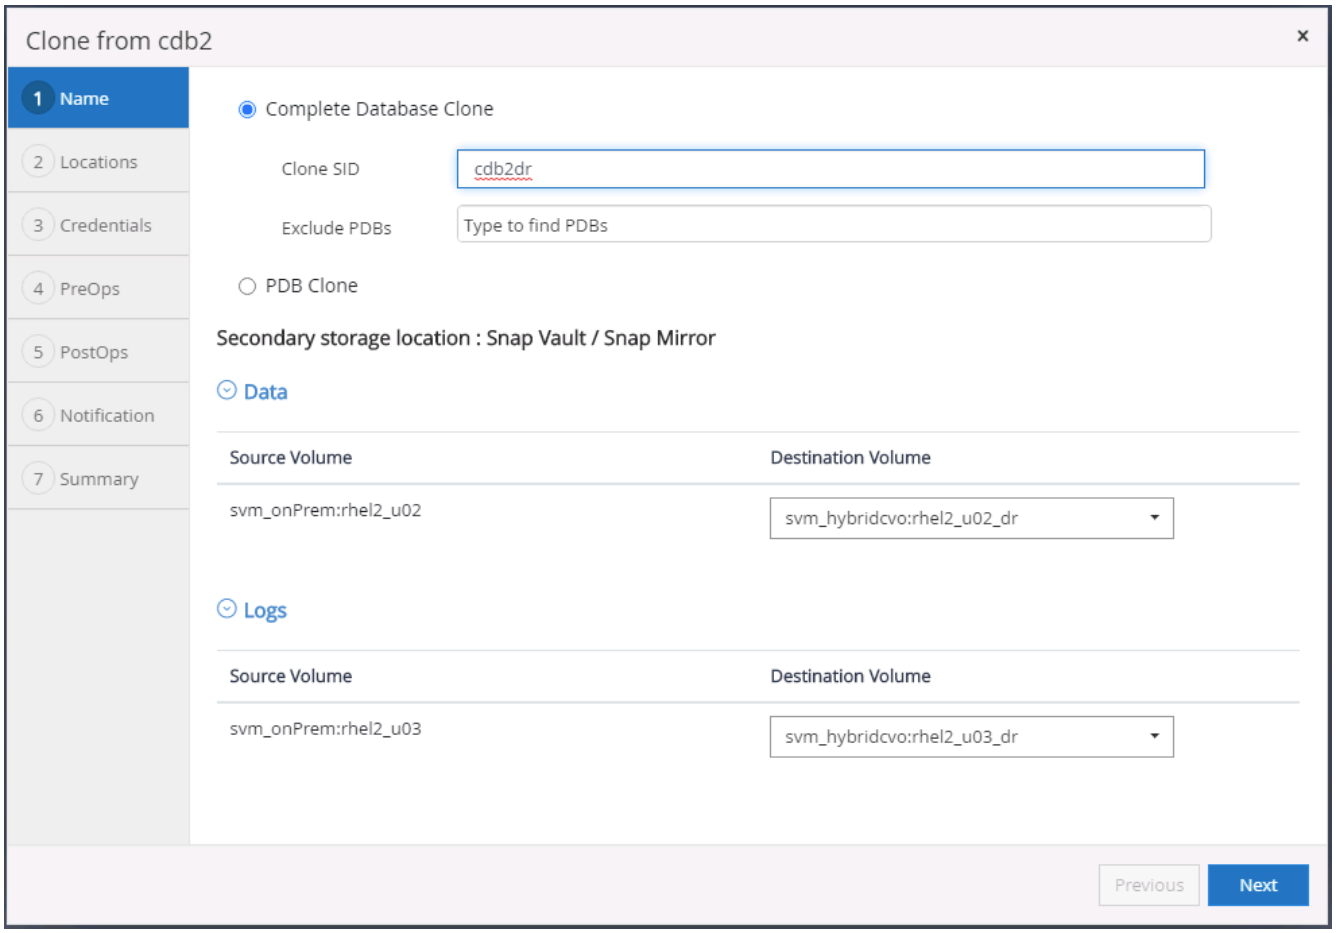

7. Provision a log volume and mount it to the target DR server for the Oracle flash recovery area and online logs.

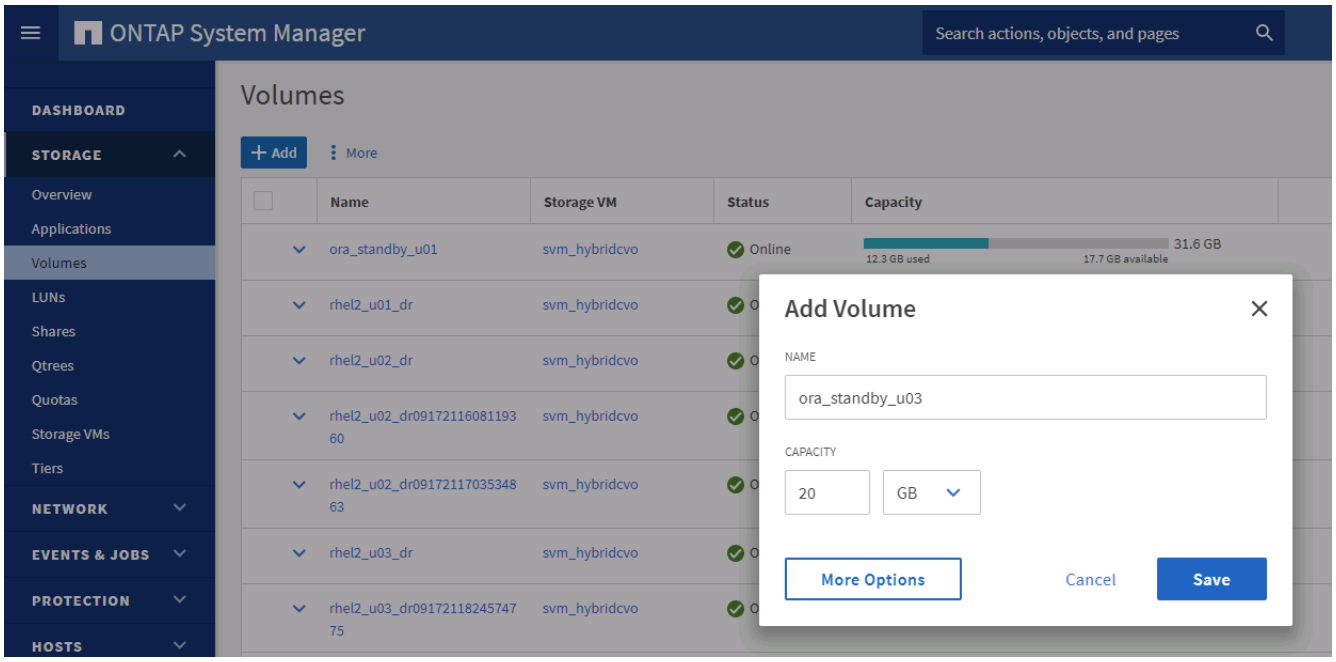

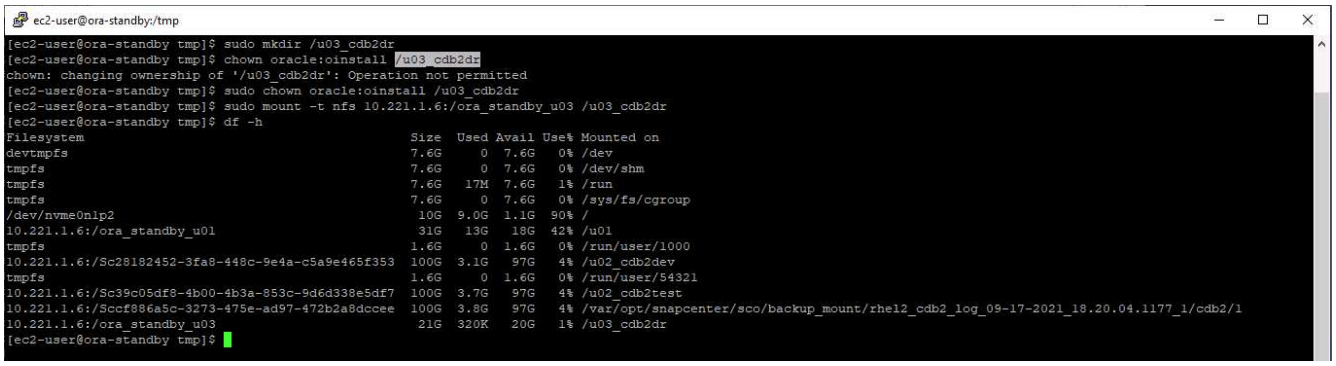

The Oracle clone procedure does not create a log volume, which needs to be provisioned on the DR server before cloning.

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

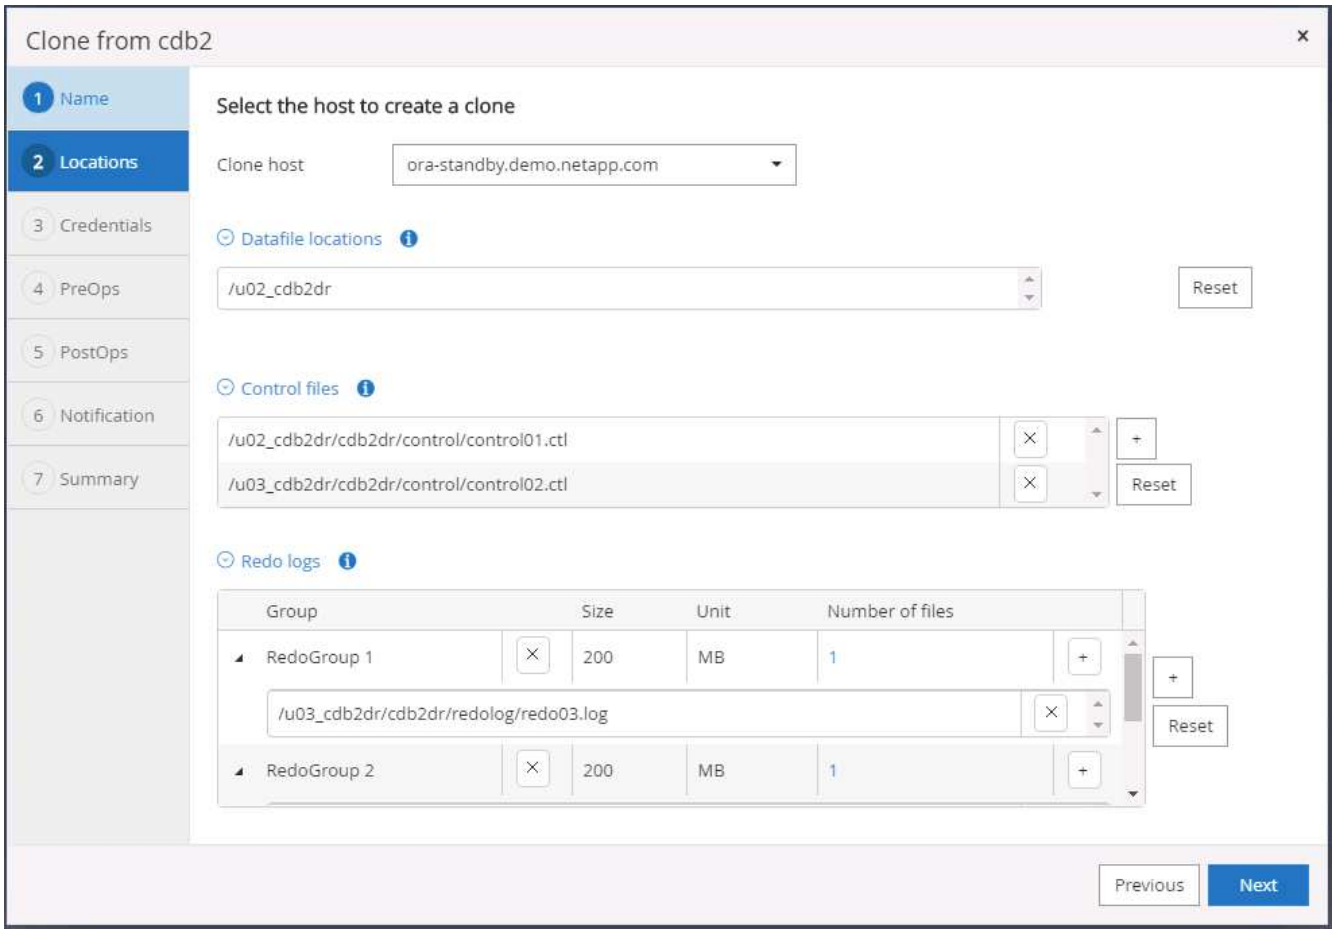

9. Select the credentials for the clone. Fill in the details of the Oracle home configuration on the target server.

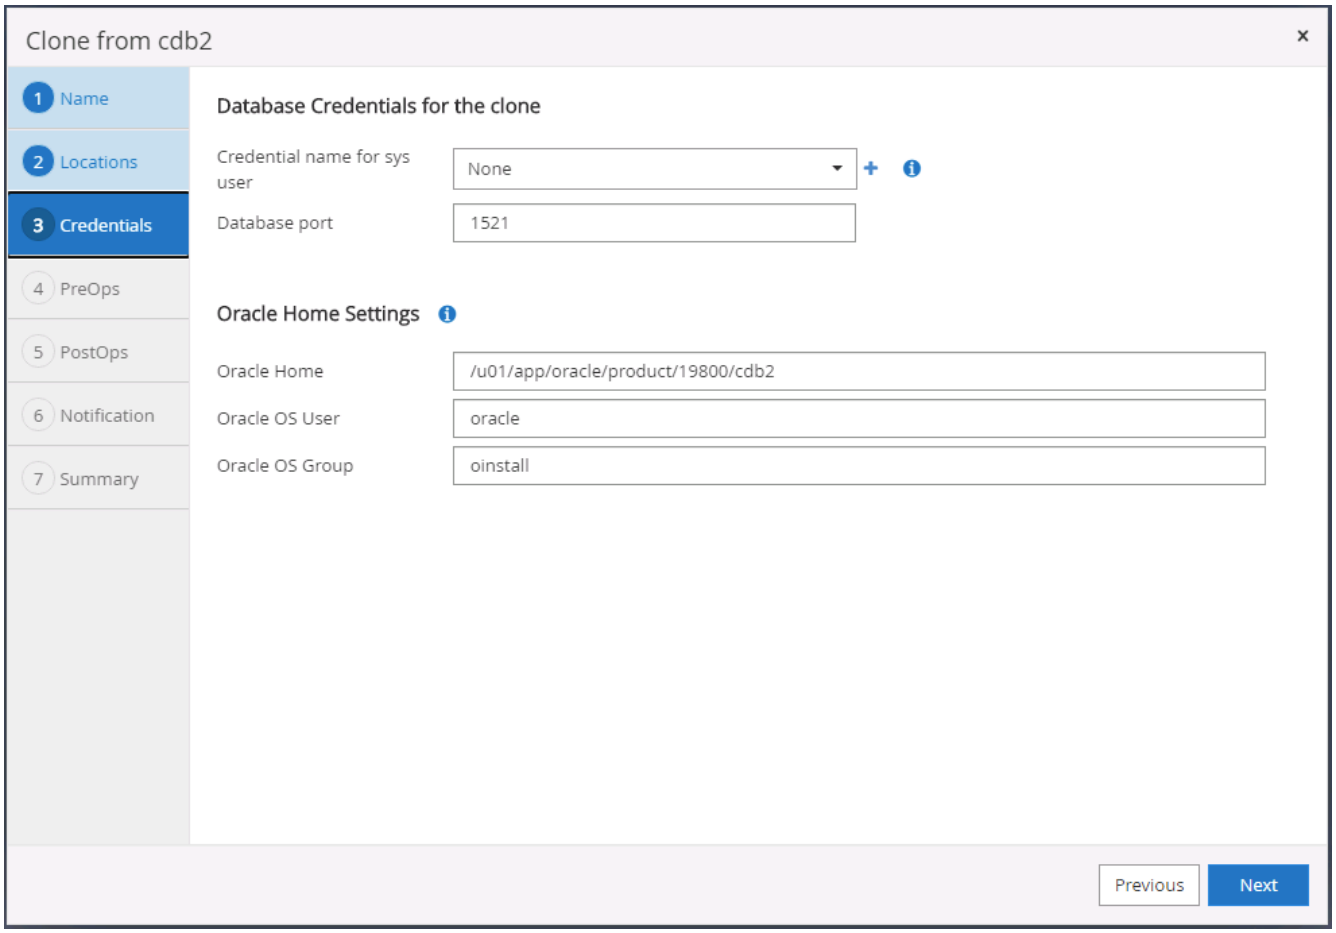

10. Specify the scripts to run before cloning. Database parameters can be adjusted if needed.

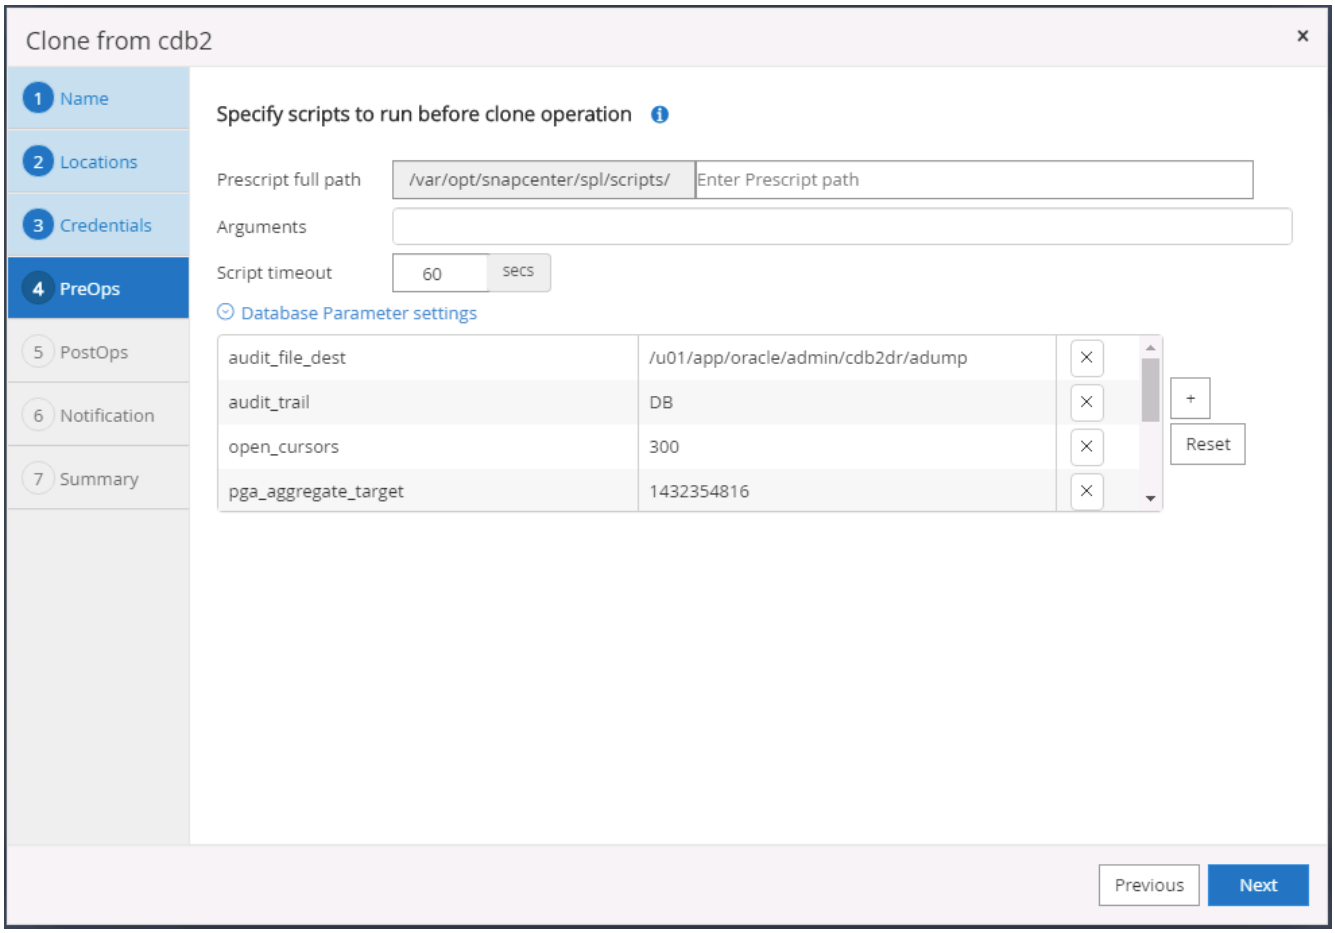

11. Select Until Cancel as the recovery option so that the recovery runs through all available archive logs to recoup the last transaction replicated to the secondary cloud location.

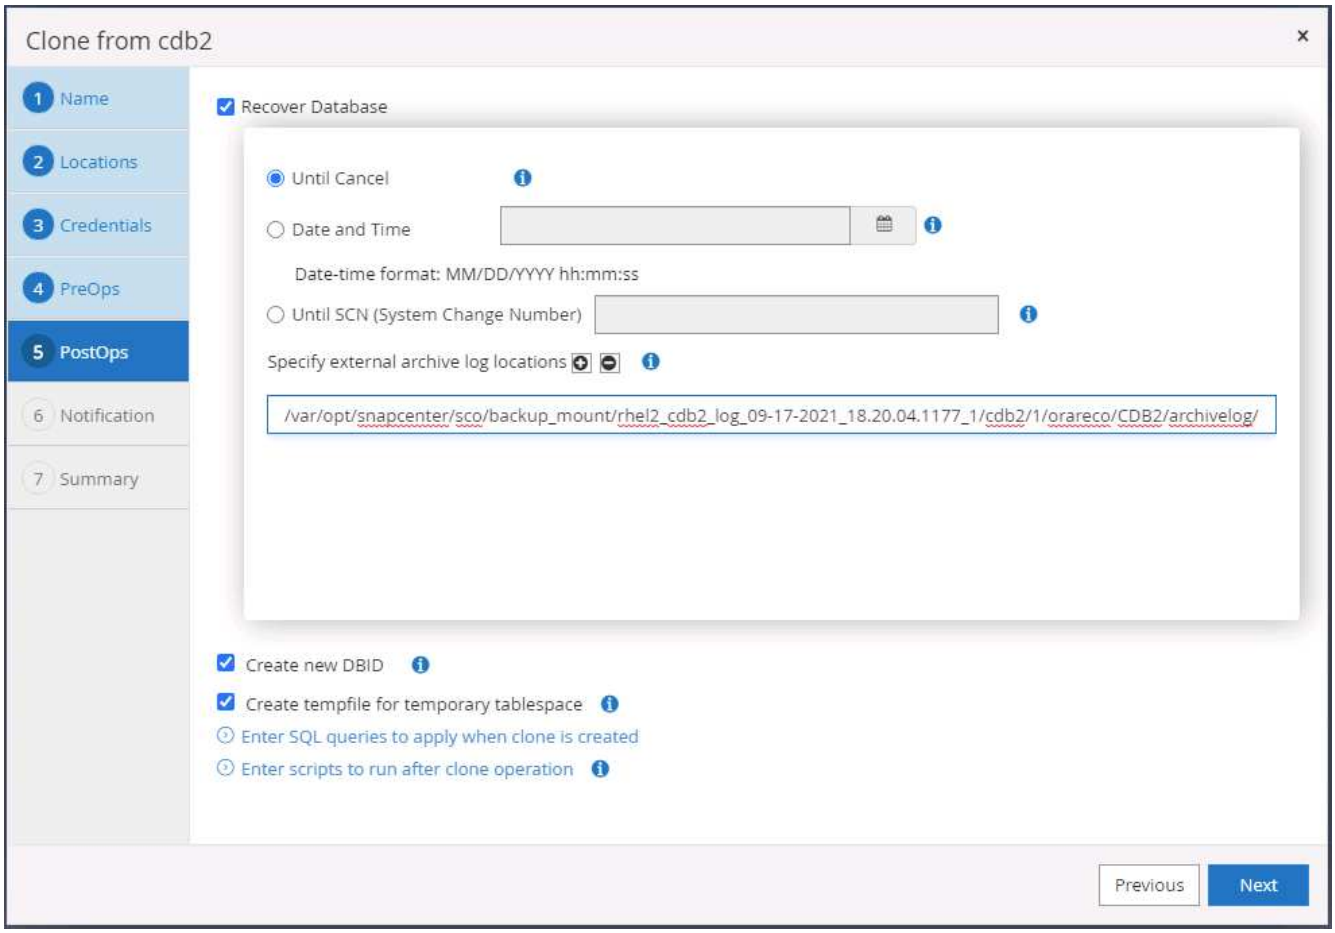

12. Configure the SMTP server for email notification if needed.

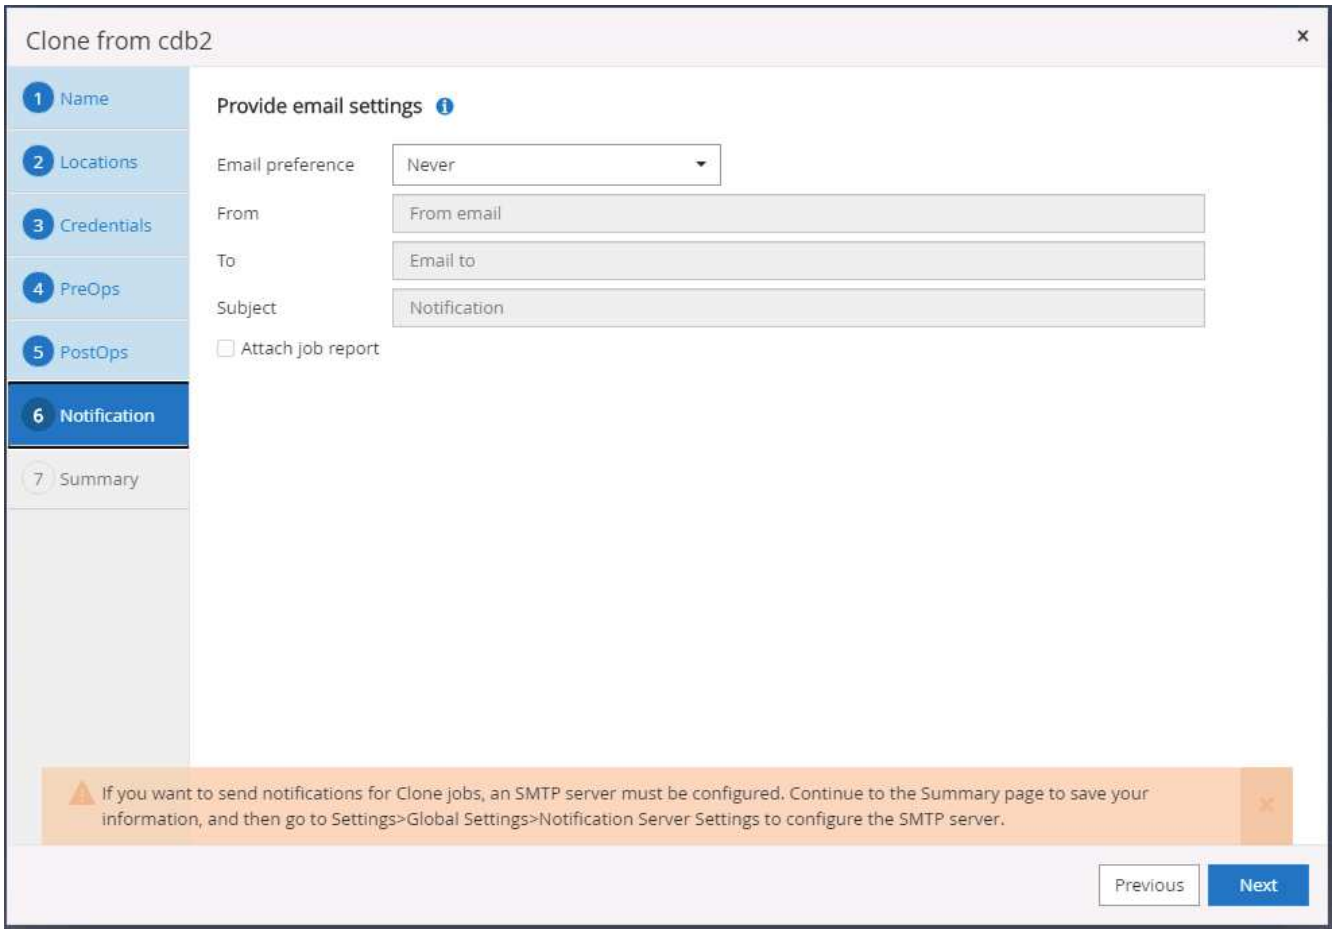

13. DR clone summary.

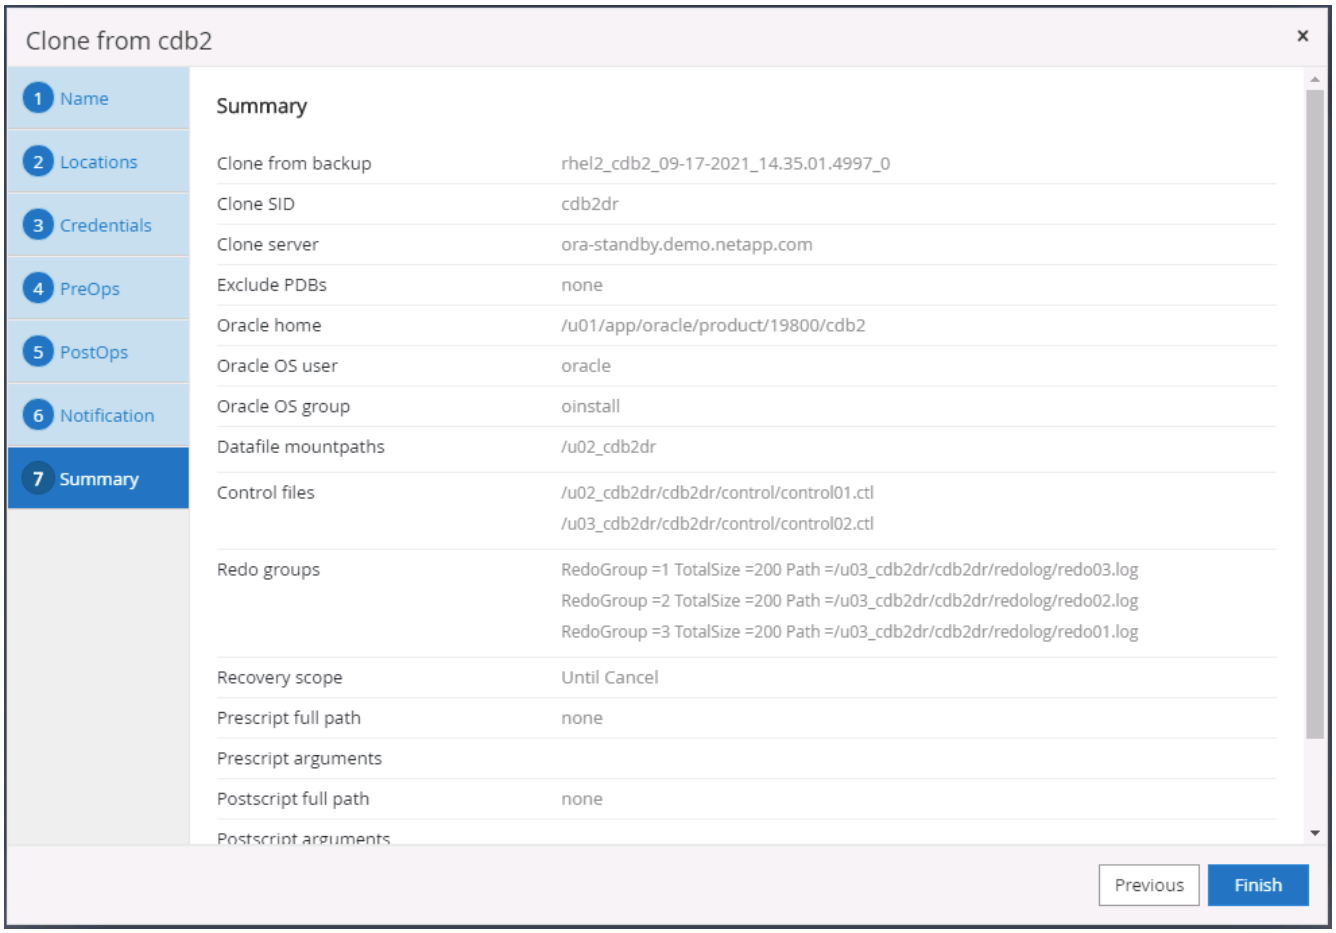

14. Cloned DBs are registered with SnapCenter immediately after clone completion and are then available for backup protection.

Post DR clone validation and configuration for Oracle

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

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

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

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

SQL> alter pluggable database cdb2_pdb1 open;

Pluggable database altered.

SQL> alter session set container=cdb2_pdb1;

Session altered.

SQL> select * from pdbadmin.dr_test;

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

SQL>
```

2. Configure the flash recovery area.

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

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

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


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

SQL> show parameter db_recovery_file_dest

NAME                                 TYPE        VALUE
------------------------------------ ----------- ------------------------------
db_recovery_file_dest                string
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.

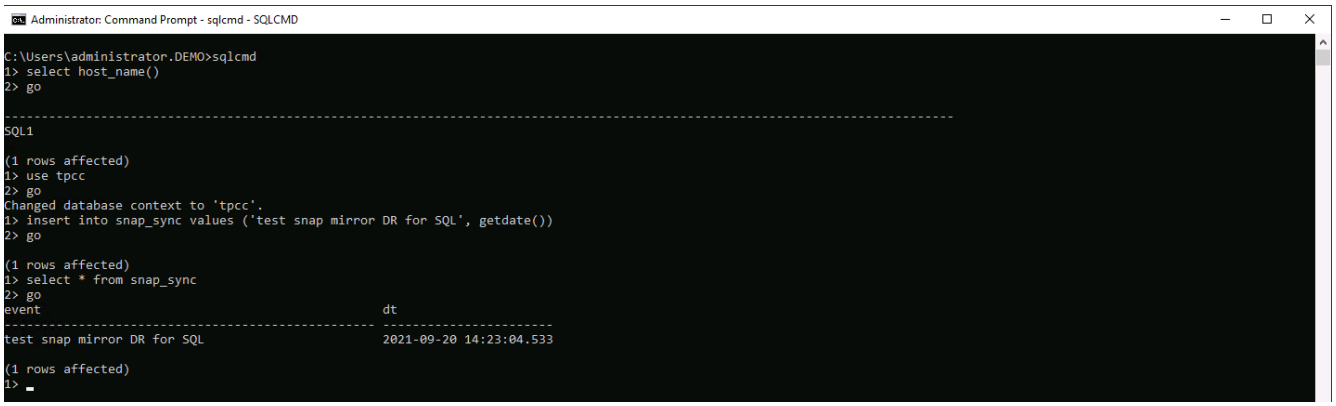

2. Log into SnapCenter with a database management user ID for SQL Server. Navigate to the Resources tab, which shows the SQL Server protection resources group.

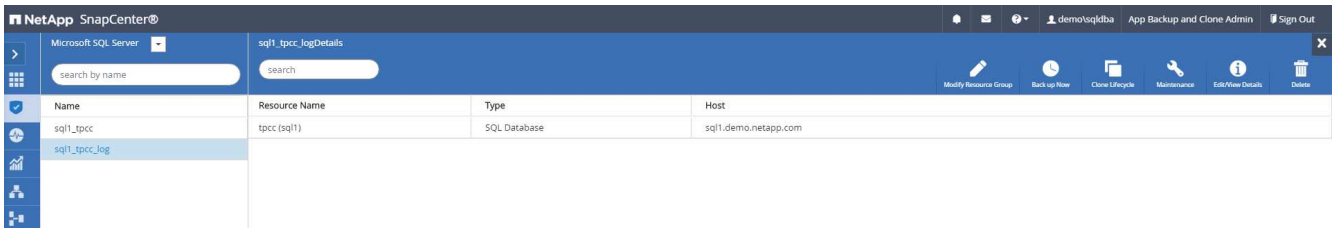

3. Manually run a log backup to flush the last transaction to be replicated to secondary storage in the public cloud.

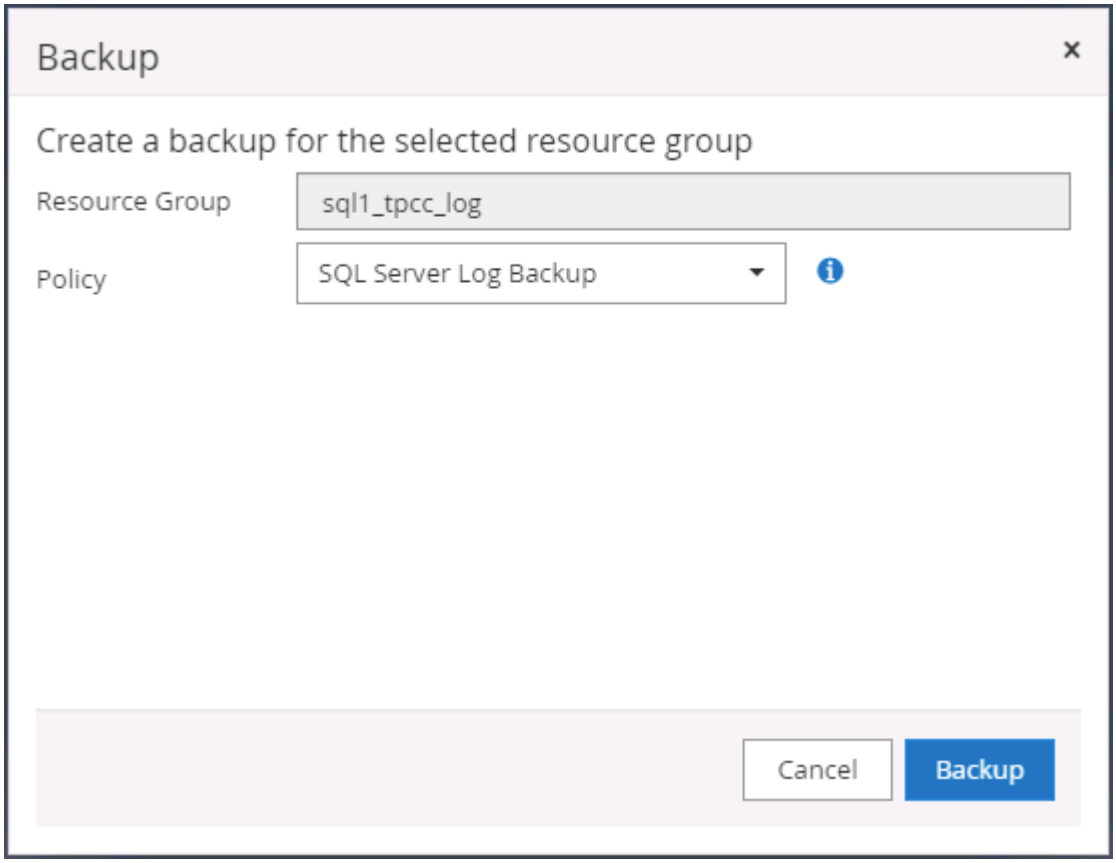

4. Select the last full SQL Server backup for the clone.

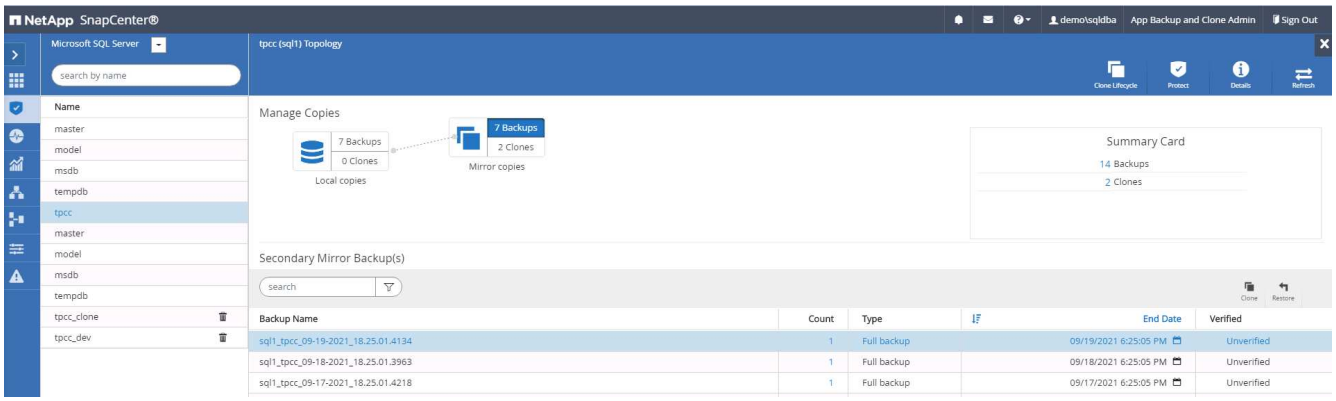

5. Set the clone setting such as the Clone Server, Clone Instance, Clone Name, and mount option. The secondary storage location where cloning is performed is auto-populated.

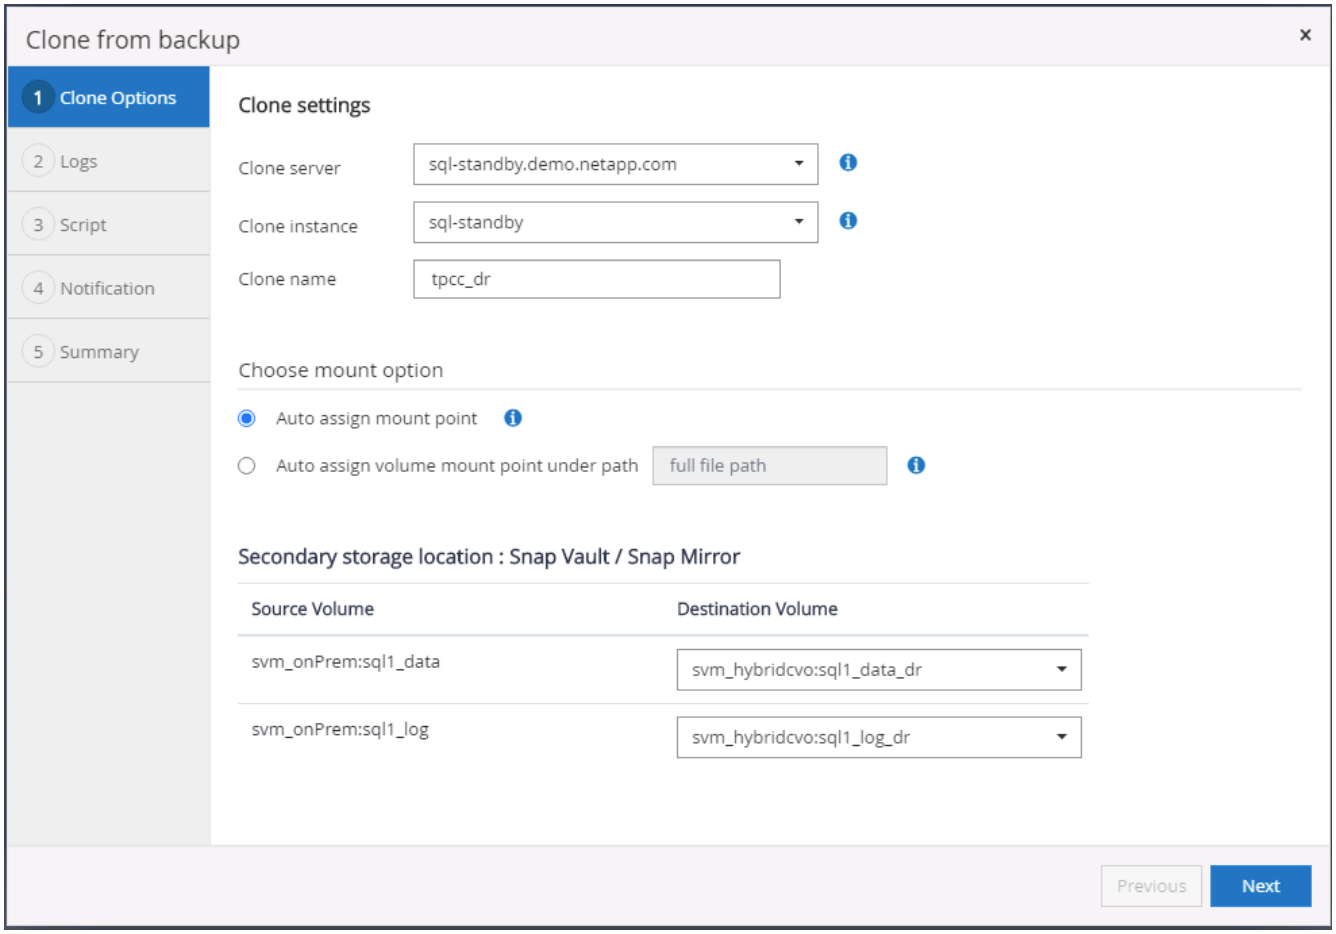

6. Select all log backups to be applied.

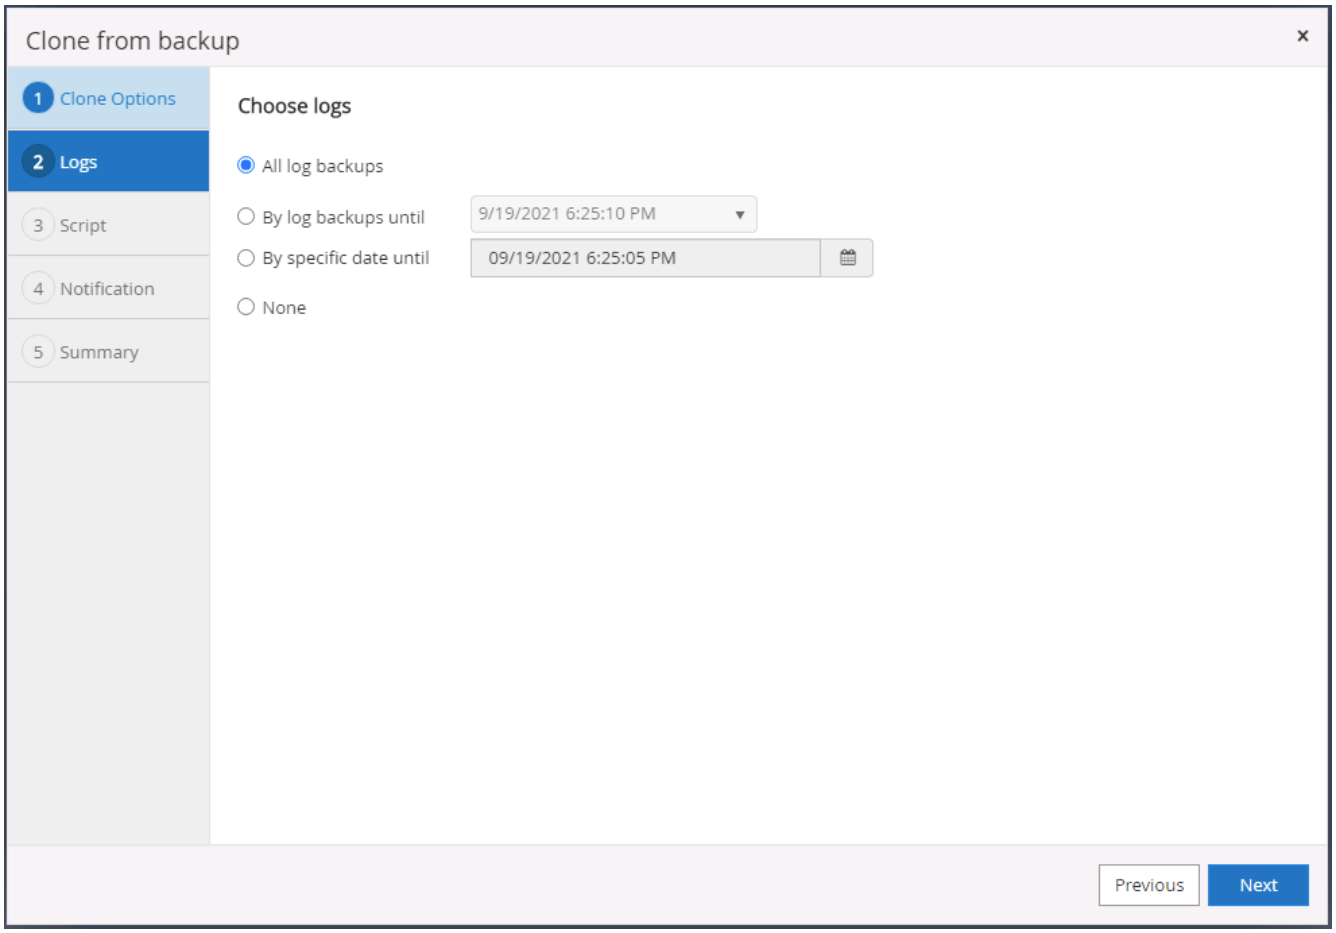

7. Specify any optional scripts to run before or after cloning.

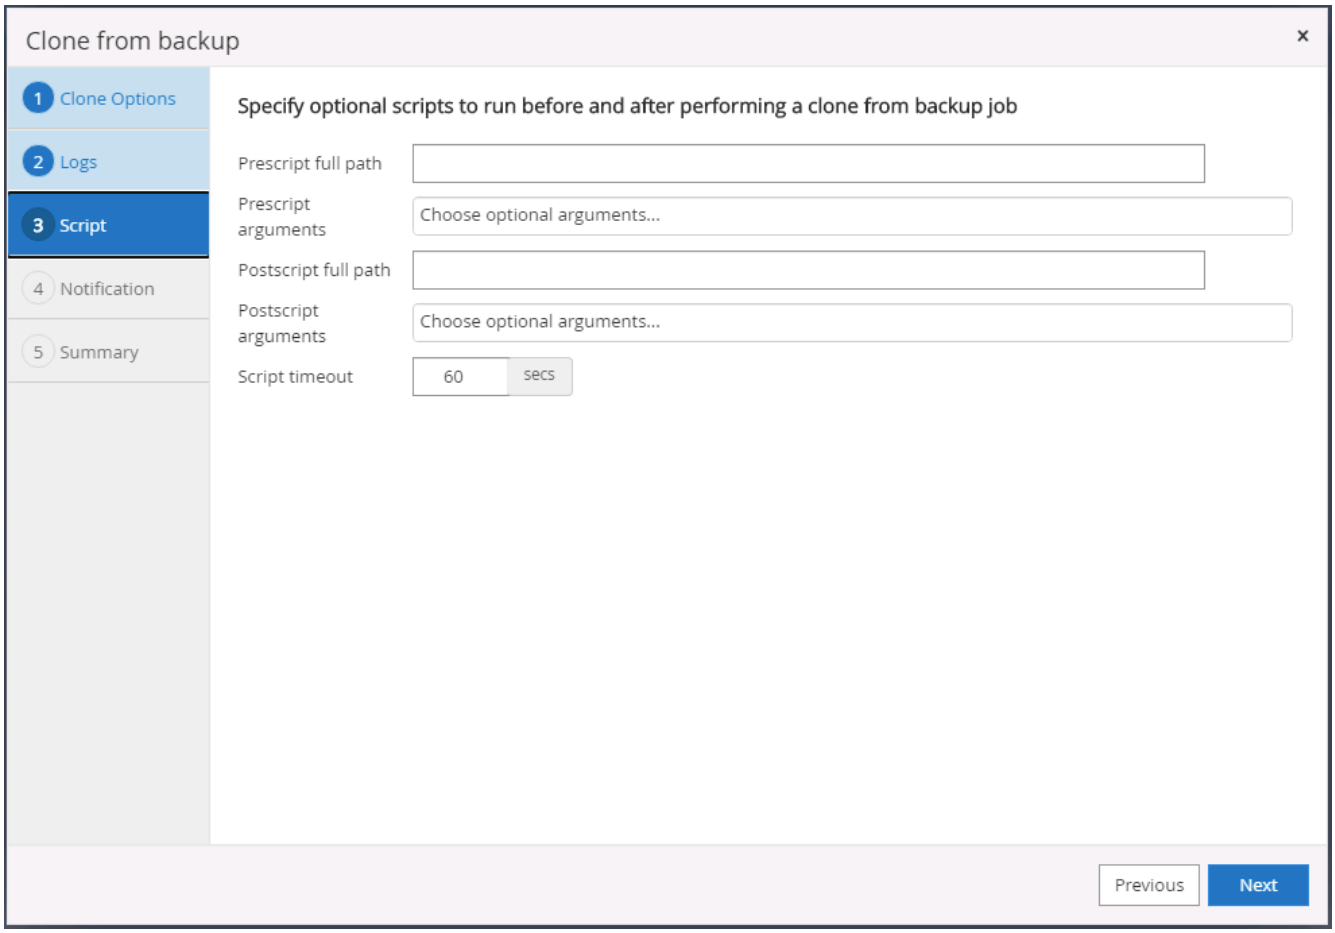

8. Specify an SMTP server if email notification is desired.

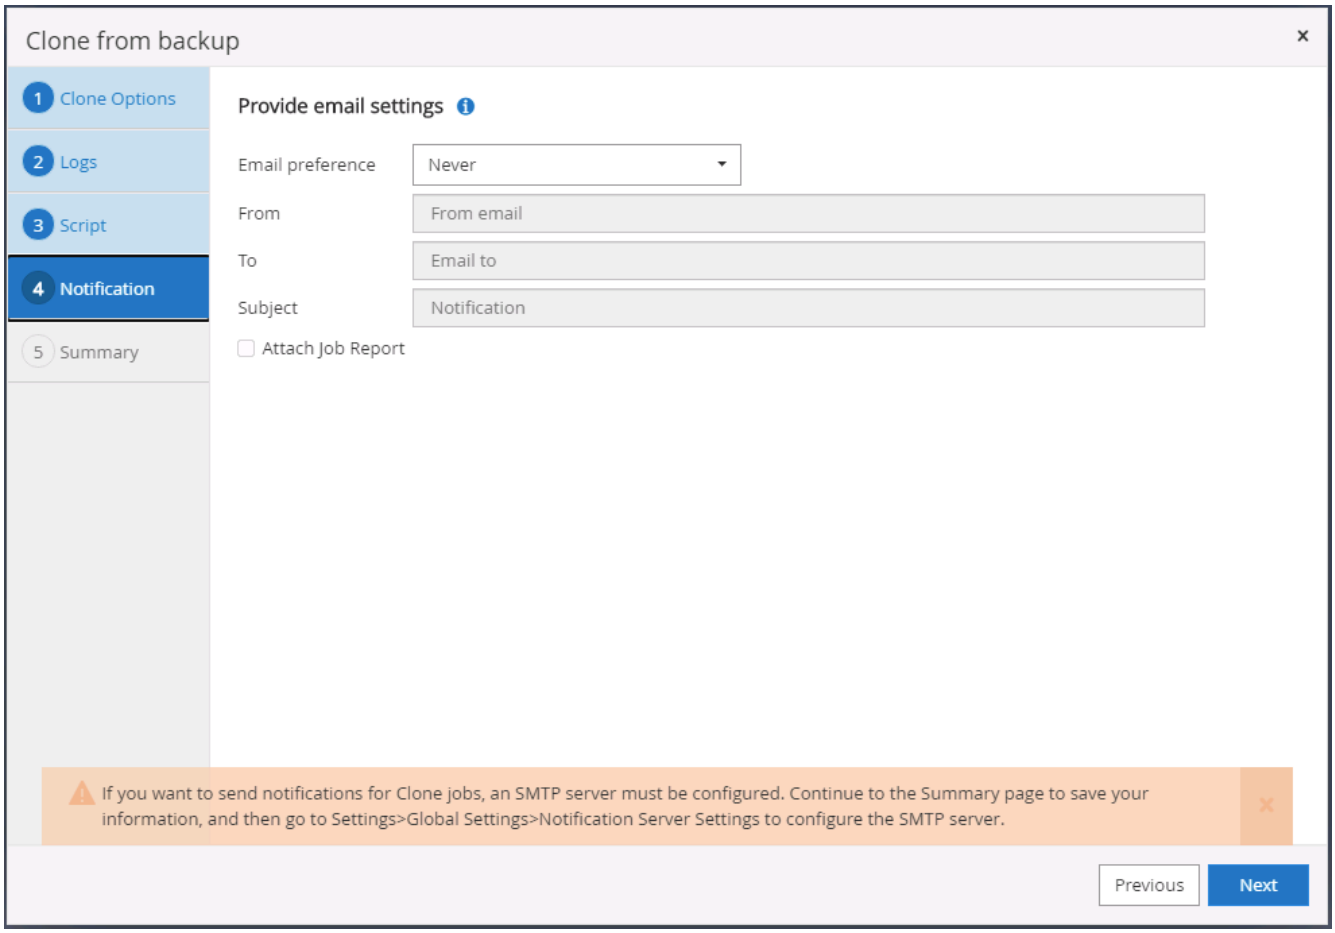

9. DR clone summary. Cloned databases are immediately registered with SnapCenter and available for backup protection.

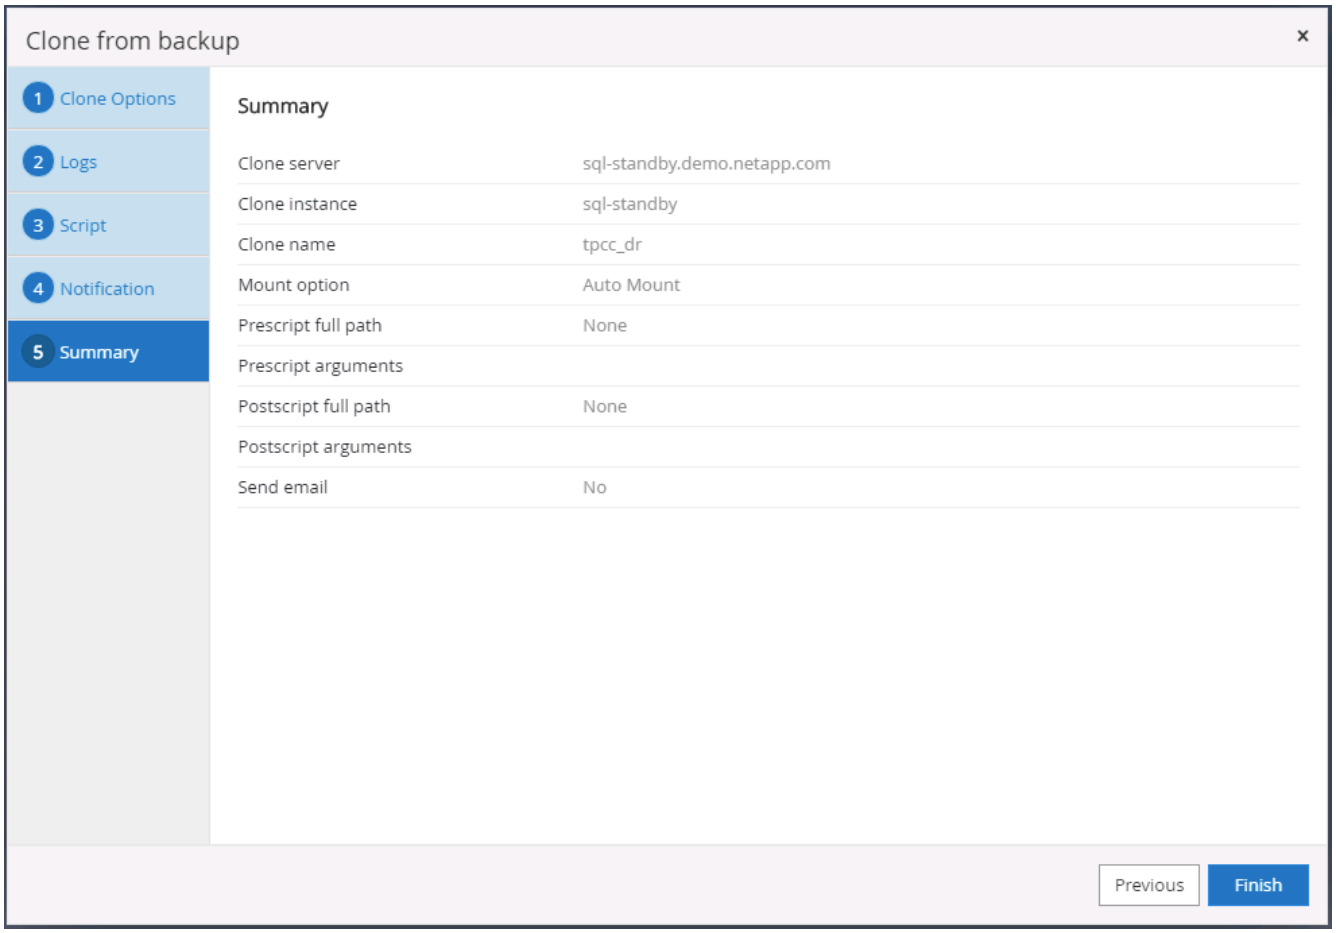

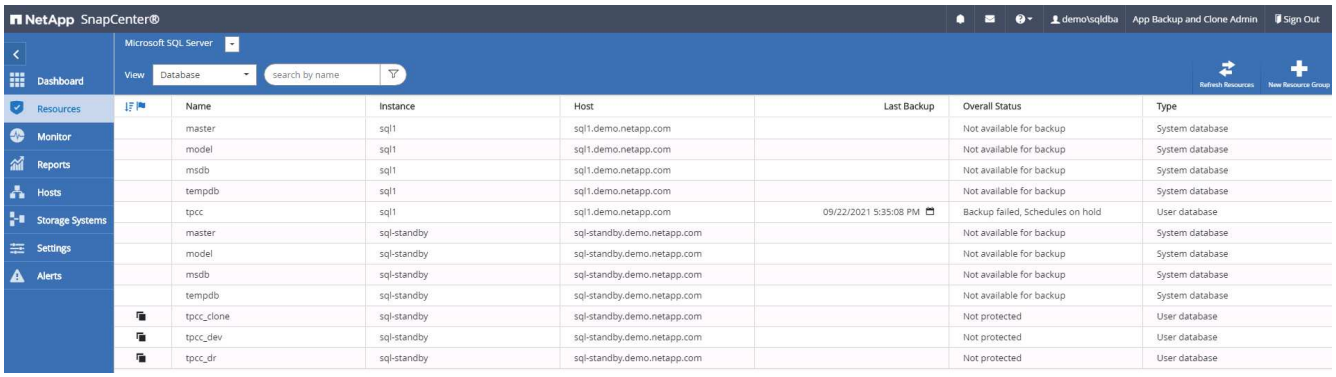

Post DR clone validation and configuration for SQL

1. Monitor clone job status.

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 inquires.