



Disaster Recovery with CVO and AVS (guest-connected storage)

NetApp public and hybrid cloud solutions

NetApp
July 30, 2025

Table of Contents

Disaster Recovery with CVO and AVS (guest-connected storage)	1
Overview	1
Assumptions.....	1
Deploying the DR Solution.....	2
Solution Deployment Overview	2
Deployment Details	2
Benefits of this solution	26

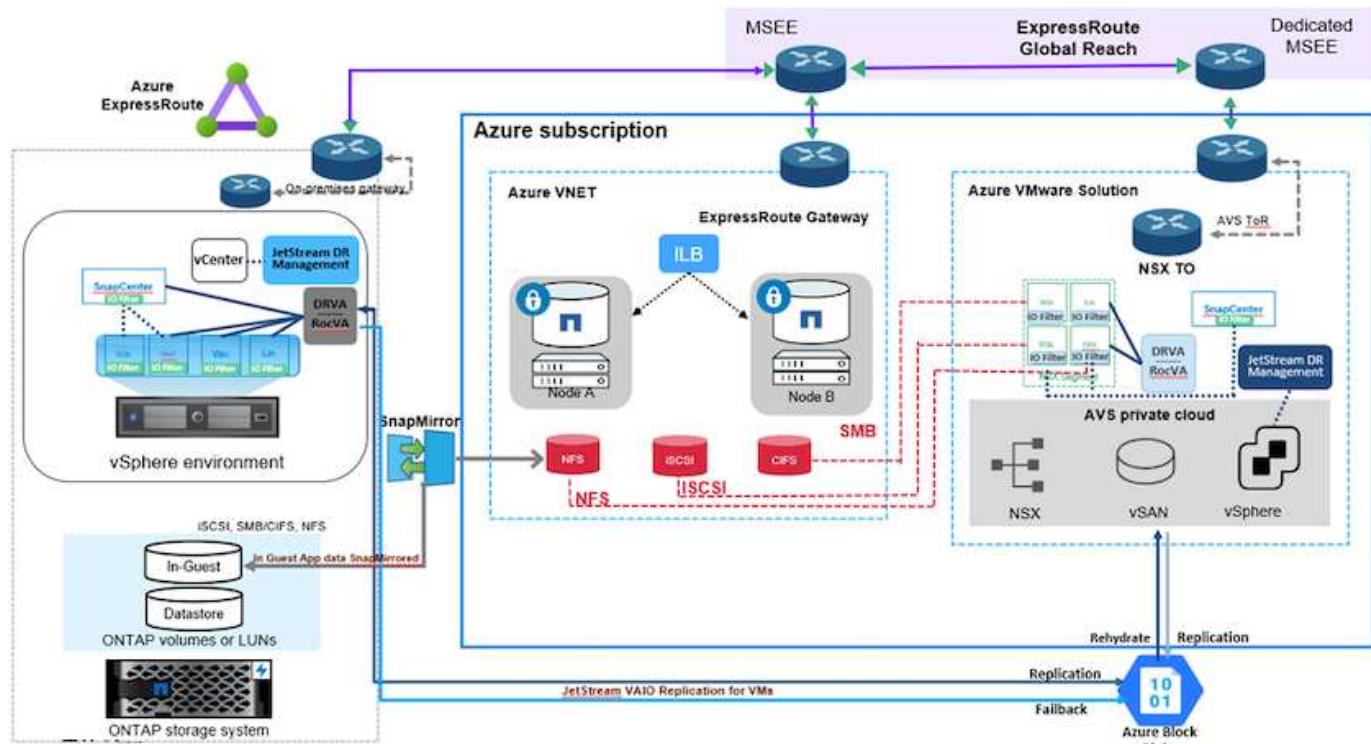
Disaster Recovery with CVO and AVS (guest-connected storage)

Disaster recovery to cloud is a resilient and cost-effective way of protecting workloads against site outages and data corruption events such as ransomware. With NetApp SnapMirror, on-premises VMware workloads that use guest-connected storage can be replicated to NetApp Cloud Volumes ONTAP running in Azure.

Overview

This covers application data; however, what about the actual VMs themselves. Disaster recovery should cover all dependent components, including virtual machines, VMDKs, application data, and more. To accomplish this, SnapMirror along with Jetstream can be used to seamlessly recover workloads replicated from on-premises to Cloud Volumes ONTAP while using vSAN storage for VM VMDKs.

This document provides a step-by-step approach for setting up and performing disaster recovery that uses NetApp SnapMirror, JetStream, and the Azure VMware Solution (AVS).



Assumptions

This document focuses on in-guest storage for application data (also known as guest connected), and we

assume that the on-premises environment is using SnapCenter for application-consistent backups.



This document applies to any third-party backup or recovery solution. Depending on the solution used in the environment, follow best practices to create backup policies that meet organizational SLAs.

For connectivity between the on-premises environment and the Azure virtual network, use the express route global reach or a virtual WAN with a VPN gateway. Segments should be created based on the on-premises vLAN design.



There are multiple options for connecting on-premises datacenters to Azure, which prevents us from outlining a specific workflow in this document. Refer to the Azure documentation for the appropriate on-premises-to-Azure connectivity method.

Deploying the DR Solution

Solution Deployment Overview

1. Make sure that application data is backed up using SnapCenter with the necessary RPO requirements.
2. Provision Cloud Volumes ONTAP with the correct instance size using Cloud manager within the appropriate subscription and virtual network.
 - a. Configure SnapMirror for the relevant application volumes.
 - b. Update the backup policies in SnapCenter to trigger SnapMirror updates after the scheduled jobs.
3. Install the JetStream DR software in the on-premises data center and start protection for virtual machines.
4. Install JetStream DR software in the Azure VMware Solution private cloud.
5. During a disaster event, break the SnapMirror relationship using Cloud Manager and trigger failover of virtual machines to Azure NetApp Files or to vSAN datastores in the designated AVS DR site.
 - a. Reconnect the iSCSI LUNs and NFS mounts for the application VMs.
6. Invoke failback to the protected site by reverse resyncing SnapMirror after the primary site has been recovered.

Deployment Details

Configure CVO on Azure and replicate volumes to CVO

The first step is to configure Cloud Volumes ONTAP on Azure ([Link](#)) and replicate the desired volumes to Cloud Volumes ONTAP with the desired frequencies and snapshot retentions.

Health Status	Source Volume	Target Volume	Total Transfer Time	Status	Mirror State	Last Successful Transfer	...
✓	gcsdrsql ldb_sc46 ntaphci-a300e9u25	gcsdrsql ldb_sc46_copy ANFCVODRDemo	17 seconds	idle	snapmirrored	May 6, 2022, 11:43:18 AM 105.06 KiB	...
✓	gcsdrsql hld_sc46_copy ANFCVODRDemo	gcsdrsql hld_sc46 ntaphci-a300e9u25	7 seconds	idle	snapmirrored	May 6, 2022, 11:42:20 AM 7.22 MiB	...
✓	gcsdrsql log_sc46 ntaphci-a300e9u25	gcsdrsql log_sc46_copy ANFCVODRDemo	16 seconds	idle	snapmirrored	May 6, 2022, 11:43:52 AM 130.69 KiB	...

Configure AVS hosts and CVO data access

Two important factors to consider when deploying the SDDC are the size of the SDDC cluster in the Azure VMware solution and how long to keep the SDDC in service. These two key considerations for a disaster recovery solution help reduce the overall operational costs. The SDDC can be as small as three hosts, all the way up to a multi-host cluster in a full-scale deployment.

The decision to deploy an AVS cluster is primarily based on the RPO/RTO requirements. With the Azure VMware solution, the SDDC can be provisioned just in time in preparation for either testing or an actual disaster event. An SDDC deployed just in time saves on ESXi host costs when you are not dealing with a disaster. However, this form of deployment affects the RTO by a few of hours while SDDC is being provisioned.

The most common deployed option is to have SDDC running in an always-on, pilot-light mode of operation. This option provides a small footprint of three hosts that are always available, and it also speeds up recovery operations by providing a running baseline for simulation activities and compliance checks, thus avoiding the risk of operational drift between the production and DR sites. The pilot-light cluster can be scaled up quickly to the desired level when needed to handle an actual DR event.

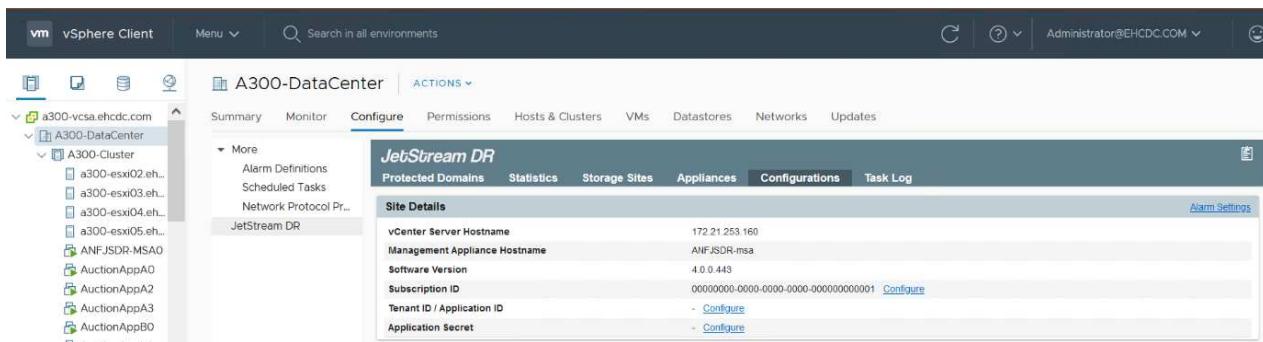
To configure AVS SDDC (be it on-demand or in pilot-light mode), see [Deploy and configure the Virtualization Environment on Azure](#). As a prerequisite, verify that the guest VMs residing on the AVS hosts are able to consume data from Cloud Volumes ONTAP after connectivity has been established.

After Cloud Volumes ONTAP and AVS have been configured properly, begin configuring Jetstream to automate the recovery of on-premises workloads to AVS (VMs with application VMDKs and VMs with in-guest storage) by using the VAIO mechanism and by leveraging SnapMirror for application volumes copies to Cloud Volumes ONTAP.

Install JetStream DR in on-premises datacenter

JetStream DR software consists of three major components: the JetStream DR Management Server Virtual Appliance (MSA), the DR Virtual Appliance (DRVA), and host components (I/O filter packages). The MSA is used to install and configure host components on the compute cluster and then to administer JetStream DR software. The installation process is as follows:

1. Check the prerequisites.
2. Run the Capacity Planning Tool for resource and configuration recommendations.
3. Deploy the JetStream DR MSA to each vSphere host in the designated cluster.
4. Launch the MSA using its DNS name in a browser.
5. Register the vCenter server with the MSA.
6. After JetStream DR MSA has been deployed and the vCenter Server has been registered, navigate to the JetStream DR plug-in with the vSphere Web Client. This can be done by navigating to Datacenter > Configure > JetStream DR.

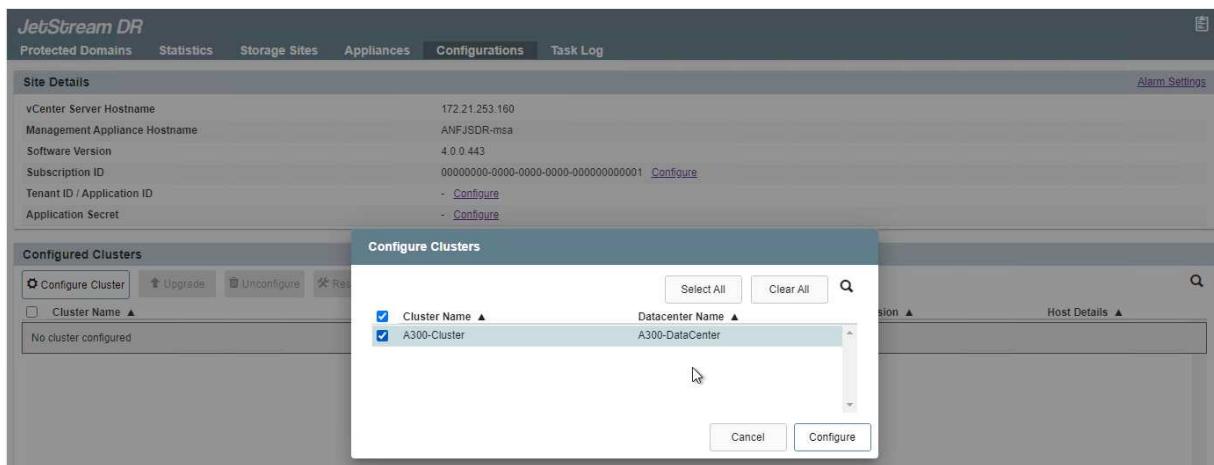


The screenshot shows the vSphere Client interface with the 'A300-DataCenter' cluster selected in the left navigation pane. The 'Configure' tab is active in the top navigation bar. The main content area is titled 'JetStream DR' and shows the 'Configurations' tab selected. The 'Site Details' section displays the following information:

Setting	Value
vCenter Server Hostname	172.21.253.160
Management Appliance Hostname	ANFJSDR-msa
Software Version	4.0.0.443
Subscription ID	00000000-0000-0000-0000-000000000001
Tenant ID / Application ID	- Configure
Application Secret	- Configure

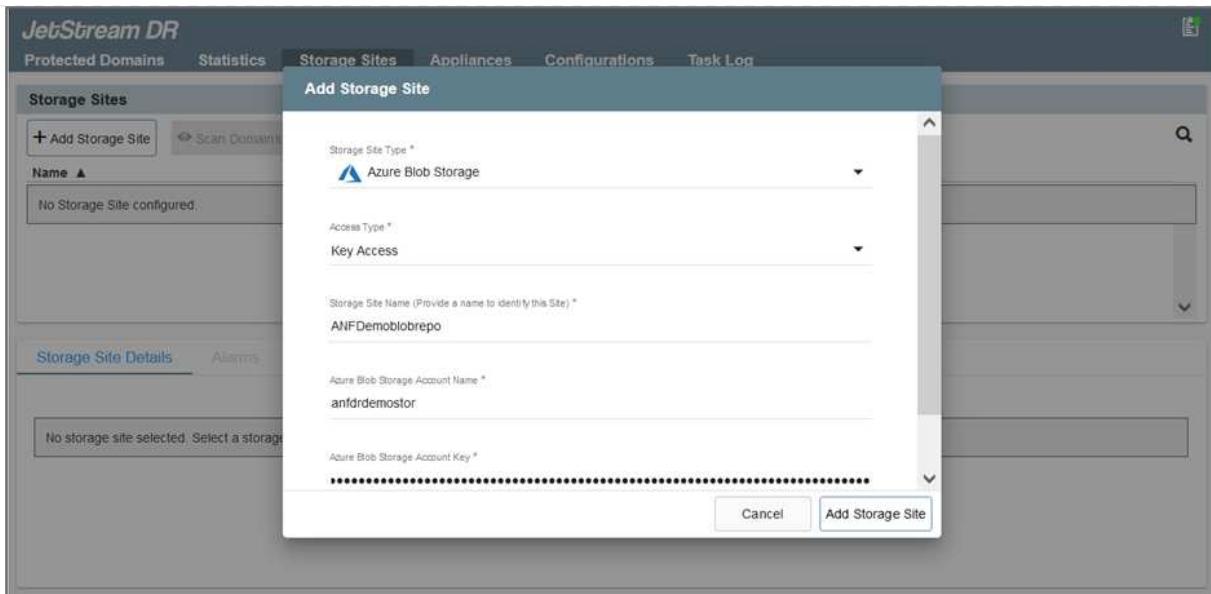
7. From the JetStream DR interface, complete the following tasks:

- a. Configure the cluster with the I/O filter package.



The screenshot shows the 'JetStream DR' configuration interface with the 'Configurations' tab selected. The 'Configured Clusters' section shows a list with one item: 'A300-Cluster'. A modal dialog box titled 'Configure Clusters' is open, showing a table with two columns: 'Cluster Name' and 'Datacenter Name'. The table contains one row with the values 'A300-Cluster' and 'A300-DataCenter'. The 'Configure' button is visible at the bottom right of the dialog.

- b. Add the Azure Blob storage located at the recovery site.



8. Deploy the required number of DR Virtual Appliances (DRVAs) from the Appliances tab.



Use the capacity planning tool to estimate the number of DRVAs required.

9. Create replication log volumes for each DRVA using the VMDK from the datastores available or the independent shared iSCSI storage pool.

The screenshot shows the JetStream DR interface with the following sections:

- DRVAs (DR Virtual Appliances):** A table with columns: Name, Status, Child Alarm, Software Version, and Details. One entry is shown: GCSDRPD001, Status: Running, Child Alarm: 0, Software Version: 4.0.0.134, Details: [link].
- Replication Log Volume:** A table with columns: Disk Path Name, Status, Child Alarm, Size (available/total), and Details. One entry is shown: /dev/sdb, Status: Ok, Child Alarm: 0, Size: 179.88 GB / 200 GB, Details: [link].
- Replication Log Volume Details:** A detailed view of the replication log volume, showing its configuration and status.

- From the Protected Domains tab, create the required number of protected domains using information about the Azure Blob Storage site, the DRVA instance, and the replication log. A protected domain defines a specific VM or set of application VMs within the cluster that are protected together and assigned a priority order for failover/failback operations.

The screenshot shows the 'Create Protected Domain' dialog box, Step 1: General. The fields are as follows:

Protected Domain Name	GCSDRPD_Demo01
Priority Level (Optional)	-
Description	Protection domain ANF
Total estimated data size to be protected	1000GB
DR Virtual Appliance	GCSDRPD001
Compression	Yes
Compression Level	Default
Normal GC Storage Overhead	50%
Maximum GC Storage Overhead	300%
Replication Log Storage	/dev/sdb
Replication Log Size	400GB

The screenshot shows the 'Create Protected Domain' dialog box, Step 2: Primary Site. The fields are as follows:

Compression	Yes
Compression Level	Default
Normal GC Storage Overhead	50%
Maximum GC Storage Overhead	300%
Replication Log Storage	/dev/sdb
Replication Log Size	50GB
Metadata Size	31.56GB
Primary Site Datacenter	A300-DataCenter
Primary Site Cluster	A300-Cluster
Storage Site	ANFDRDemoFailoverSite
Enable PITR	No

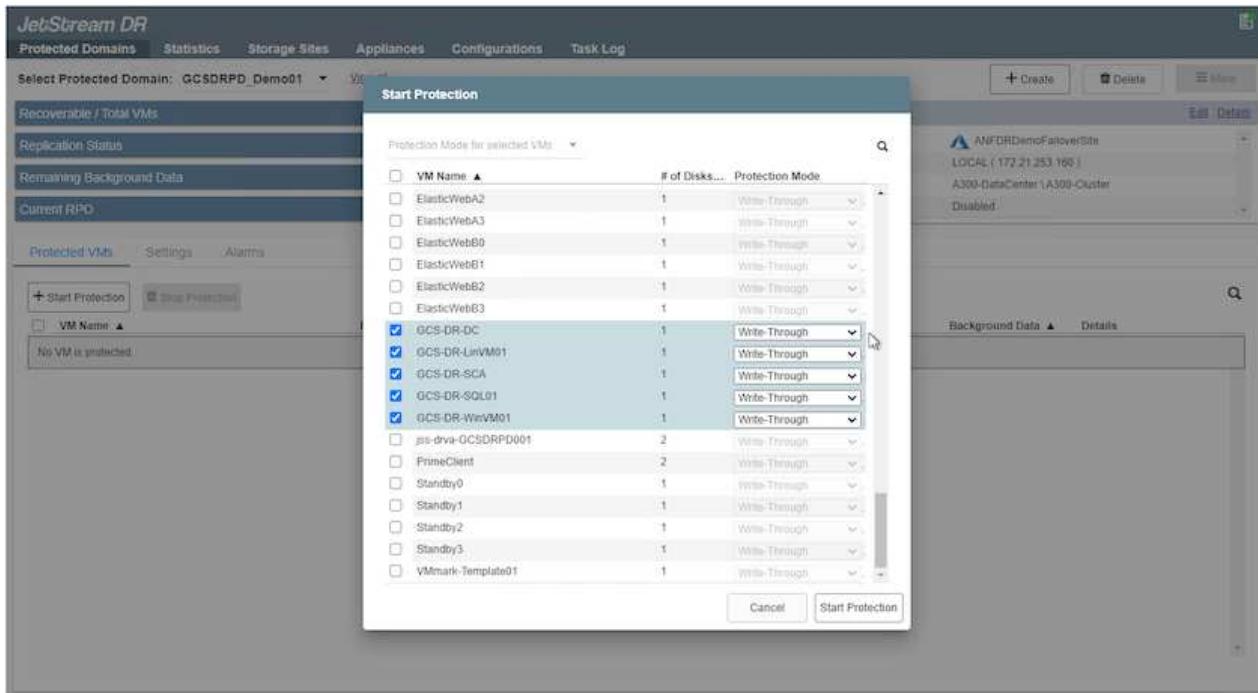
- Select the VMs to be protected and group the VMs into applications groups based on dependency. Application definitions allow you to group sets of VMs into logical groups that contain their boot orders, boot delays, and optional application validations that can be executed upon recovery.



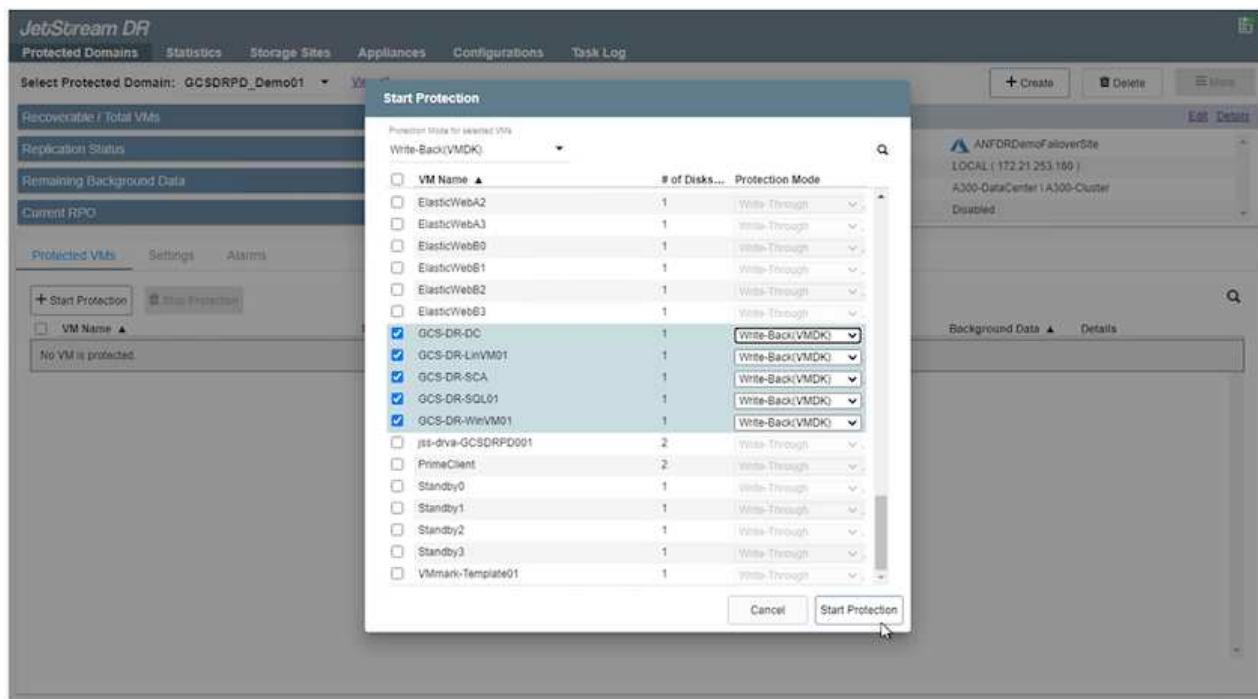
Make sure that the same protection mode is used for all VMs in a protected domain.



Write-Back(VMDK) mode offers higher performance.



12. Make sure that replication log volumes are placed on high- performance storage.



13. After you are done, click Start Protection for the protected domain. This starts data replication for the selected VMs to the designated Blob store.

Protected VMs

VM Name	Protection Status	Replication Status	Protection Mode	Background Data	Details
GCS-DR-DC	Initializing	-	Write-Back(VMDK)	-	Details
GCS-DR-LinVM01	Initializing	-	Write-Back(VMDK)	-	Details
GCS-DR-SCA	Initializing	-	Write-Back(VMDK)	-	Details
GCS-DR-SQL01	Initializing	-	Write-Back(VMDK)	-	Details
GCS-DR-WinVM01	Initializing	-	Write-Back(VMDK)	-	Details

14. After replication is completed, the VM protection status is marked as Recoverable.

Protected VMs

VM Name	Protection Status	Replication Status	Protection Mode	Background Data	Details
GCS-DR-DC	Recoverable	OK	Write-Back(VMDK)	0 B	Details
GCS-DR-LinVM01	Recoverable	OK	Write-Back(VMDK)	0 B	Details
GCS-DR-SCA	Recoverable	OK	Write-Back(VMDK)	0 B	Details
GCS-DR-SQL01	Recoverable	OK	Write-Back(VMDK)	0 B	Details
GCS-DR-WinVM01	Recoverable	OK	Write-Back(VMDK)	0 B	Details



Failover runbooks can be configured to group the VMs (called a recovery group), set the boot order sequence, and modify the CPU/memory settings along with the IP configurations.

15. Click Settings and then click the runbook Configure link to configure the runbook group.

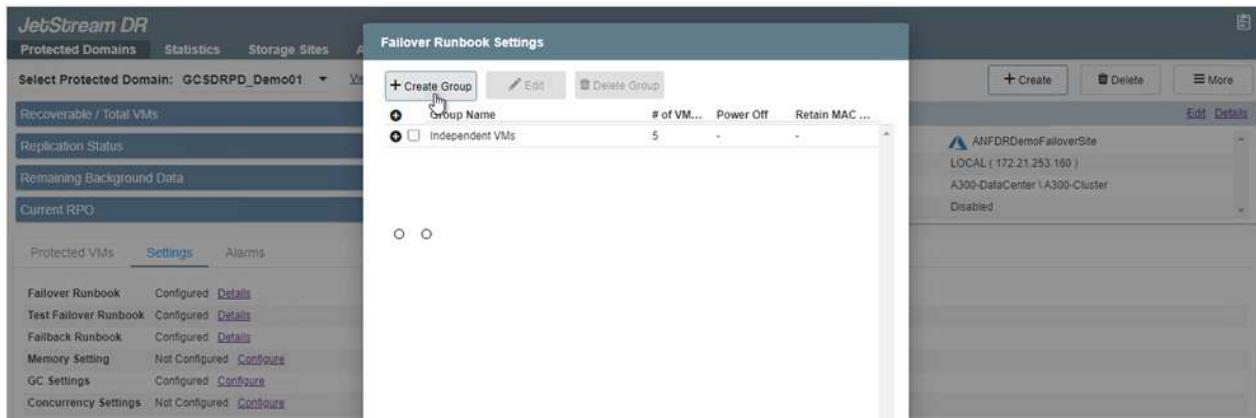
Settings

Failover Runbook	Not Configured	Configure
Test Failover Runbook	Not Configured	Configure
Fallback Runbook	Not Configured	Configure
Memory Setting	Not Configured	Configure
GC Settings	Configured	Configure
Concurrency Settings	Not Configured	Configure

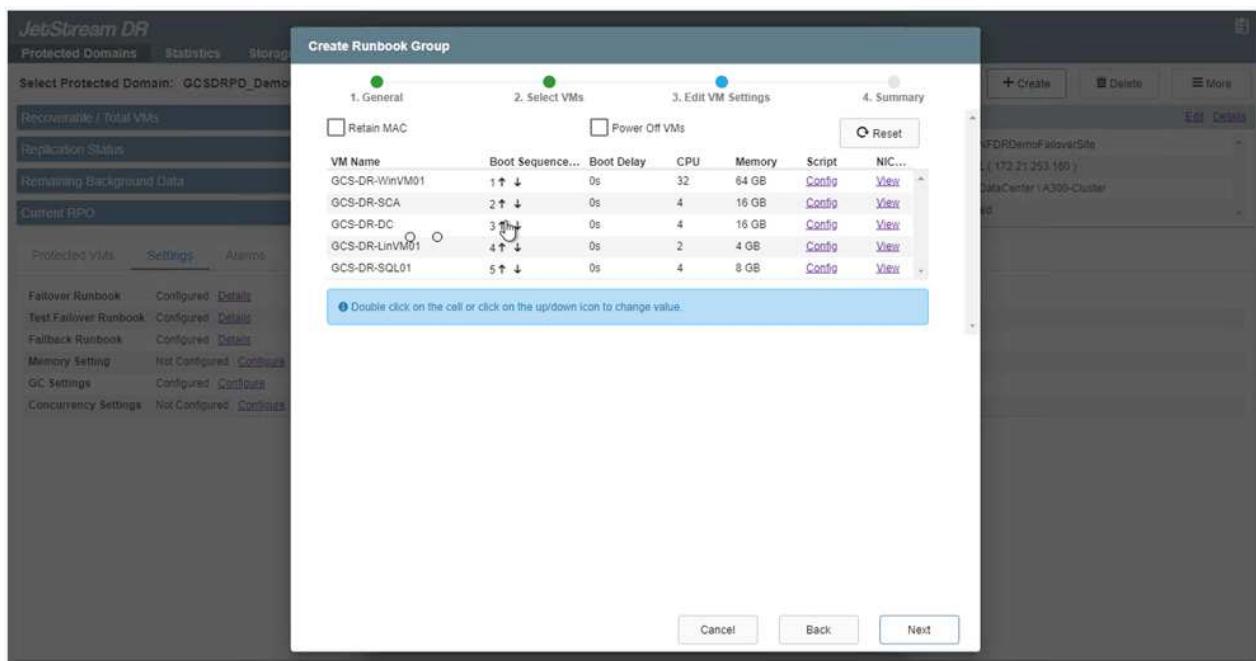
16. Click the Create Group button to begin creating a new runbook group.



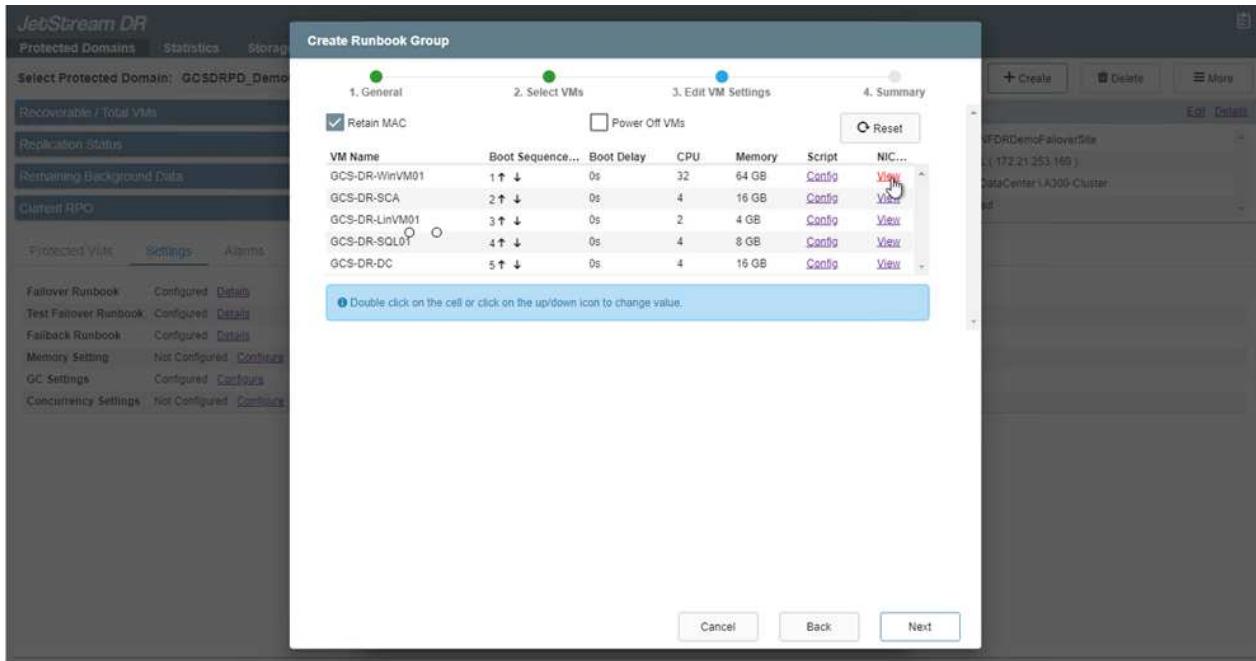
If needed, in the lower portion of the screen, apply custom pre-scripts and post-scripts to automatically run prior to and following operation of the runbook group. Make sure that the Runbook scripts are residing on the management server.



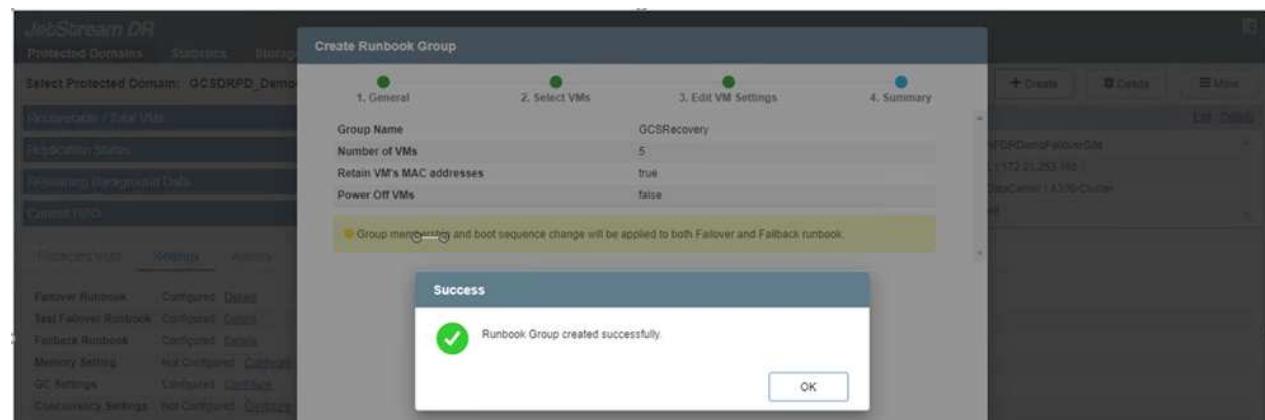
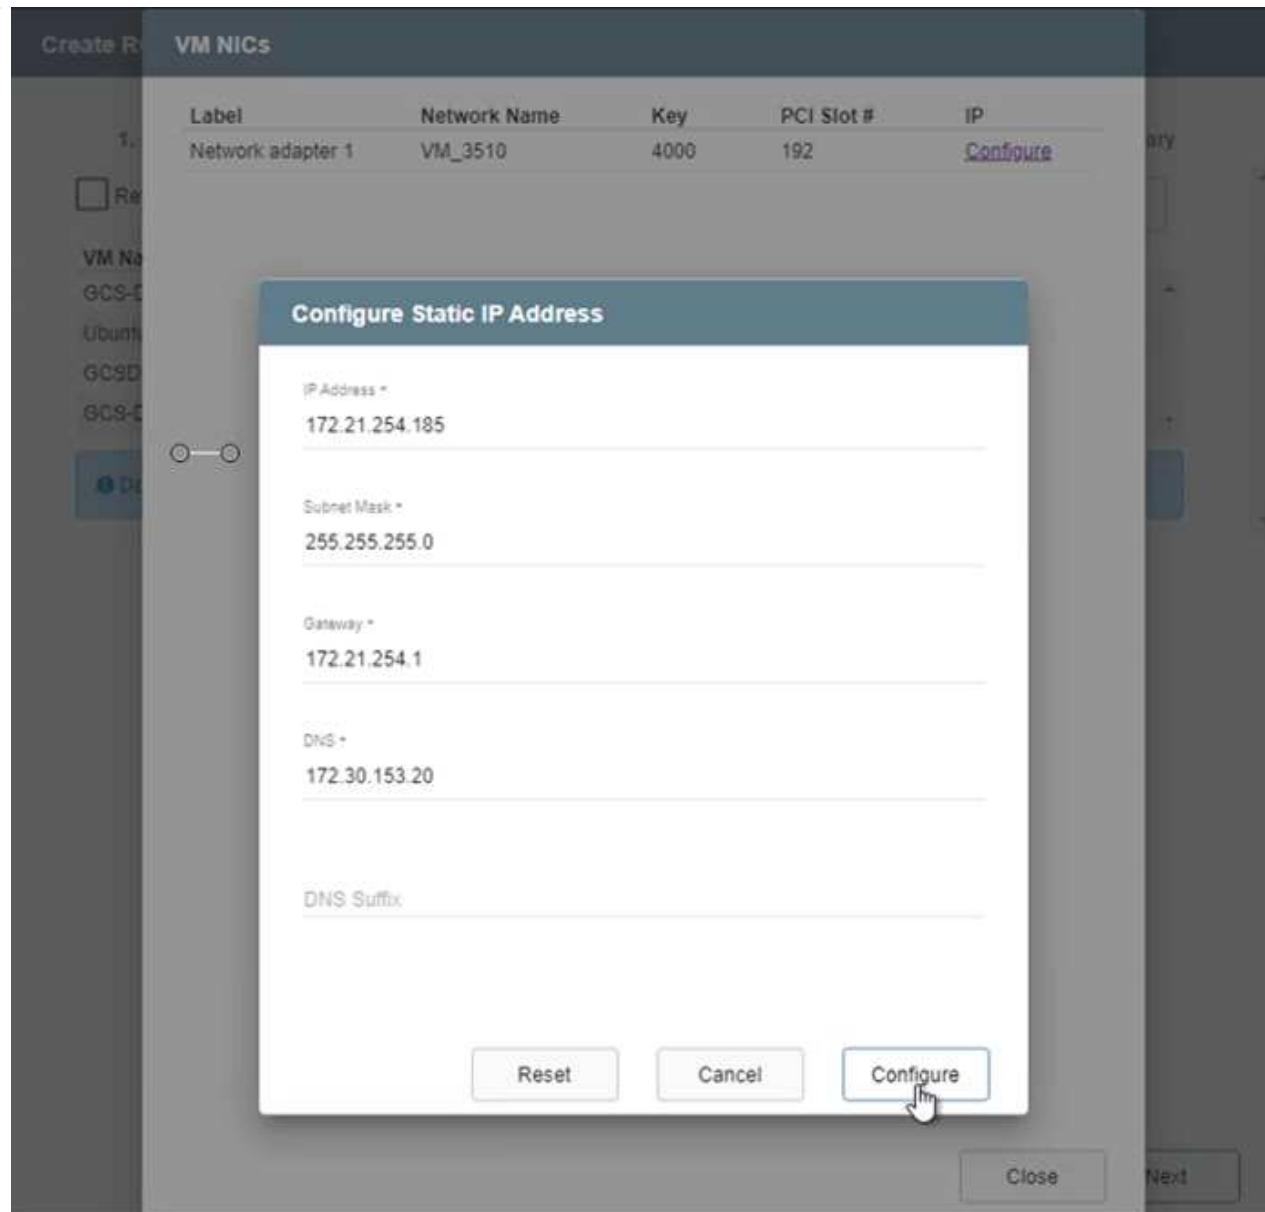
17. Edit the VM settings as required. Specify the parameters for recovering the VMs, including the boot sequence, the boot delay (specified in seconds), the number of CPUs, and the amount of memory to allocate. Change the boot sequence of the VMs by clicking the up or down arrows. Options are also provided to Retain MAC.



18. Static IP addresses can be manually configured for the individual VMs of the group. Click the NIC View link of a VM to manually configure its IP address settings.



19. Click the Configure button to save NIC settings for the respective VMs.



The status of both the failover and fallback runbooks is now listed as Configured. Failover and fallback runbook groups are created in pairs using the same initial group of VMs and settings. If necessary, the settings of any runbook group can be individually customized by clicking its respective Details link and making changes.

Install JetStream DR for AVS in private cloud

A best practice for a recovery site (AVS) is to create a three-node pilot-light cluster in advance. This allows the recovery site infrastructure to be preconfigured, including the following:

- Destination networking segments, firewalls, services like DHCP and DNS, and so on
- Installation of JetStream DR for AVS
- Configuration of ANF volumes as datastores and more

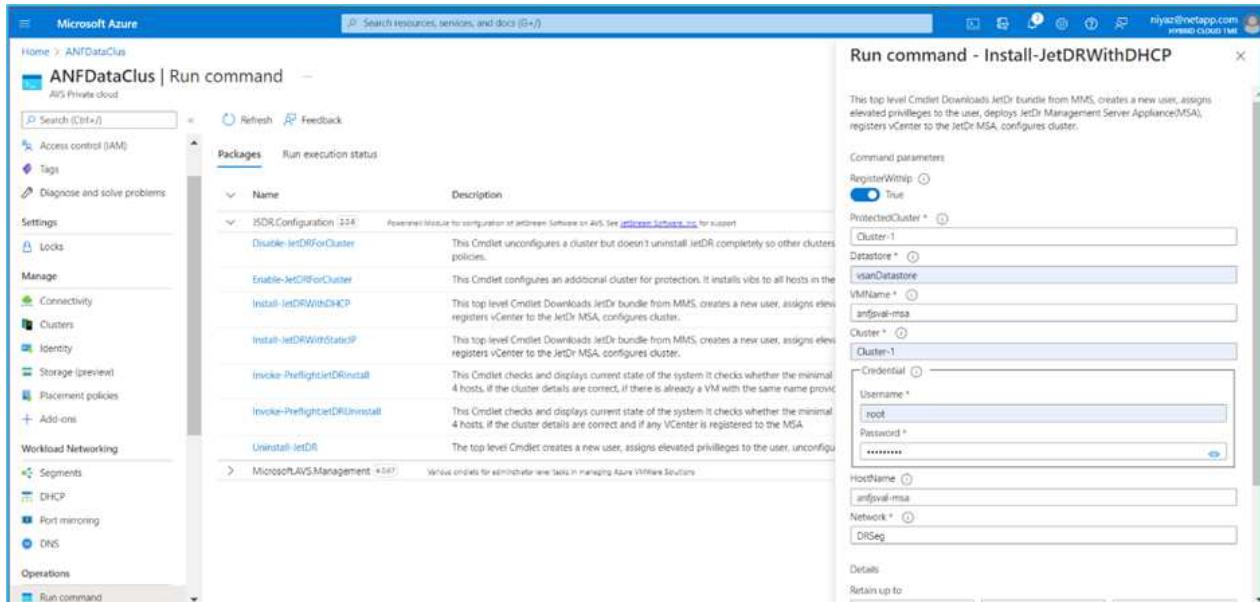
JetStream DR supports a near-zero RTO mode for mission-critical domains. For these domains, destination storage should be preinstalled. ANF is a recommended storage type in this case.

- i Network configuration including segment creation should be configured on the AVS cluster to match on-premises requirements.
- i Depending on the SLA and RTO requirements, you can use continuous failover or regular (standard) failover mode. For near-zero RTO, you should start continuous rehydration at the recovery site.

1. To install JetStream DR for AVS on an Azure VMware Solution private cloud, use the Run command. From the Azure portal, go to Azure VMware solution, select the private cloud, and select Run command > Packages > JSDR.Configuration.

- i The default CloudAdmin user of the Azure VMware Solution doesn't have sufficient privileges to install JetStream DR for AVS. The Azure VMware Solution enables simplified and automated installation of JetStream DR by invoking the Azure VMware Solution Run command for JetStream DR.

The following screenshot shows installation using a DHCP-based IP address.



2. After JetStream DR for AVS installation is complete, refresh the browser. To access the JetStream DR UI, go to SDDC Datacenter > Configure > JetStream DR.

JetStream DR

Protected Domains Statistics Storage Sites Appliances Configurations Task Log

Site Details

[Alarm Settings](#)

vCenter Server Hostname 172.30.156.2

Management Appliance Hostname anfjsval-msa

Software Version 4.0.2.450

Subscription ID - [Configure](#)

Tenant ID / Application ID - [Configure](#)

Application Secret - [Configure](#)

[Configure Cluster](#)

[Upgrade](#)

[Unconfigure](#)

[Resolve Configure Issue](#)



Cluster Name ▲

Datacenter Name ▲

Status ▲

Software Version ▲

Host Details ▲

Cluster-1

SDDC-Datacenter

Ok

4.0.2.132

[Details](#)

3. From the JetStream DR interface, complete the following tasks:

- Add the Azure Blob Storage account that was used to protect the on-premises cluster as a storage site and then run the Scan Domains option.
- In the pop-up dialog window that appears, select the protected domain to import and then click its Import link.

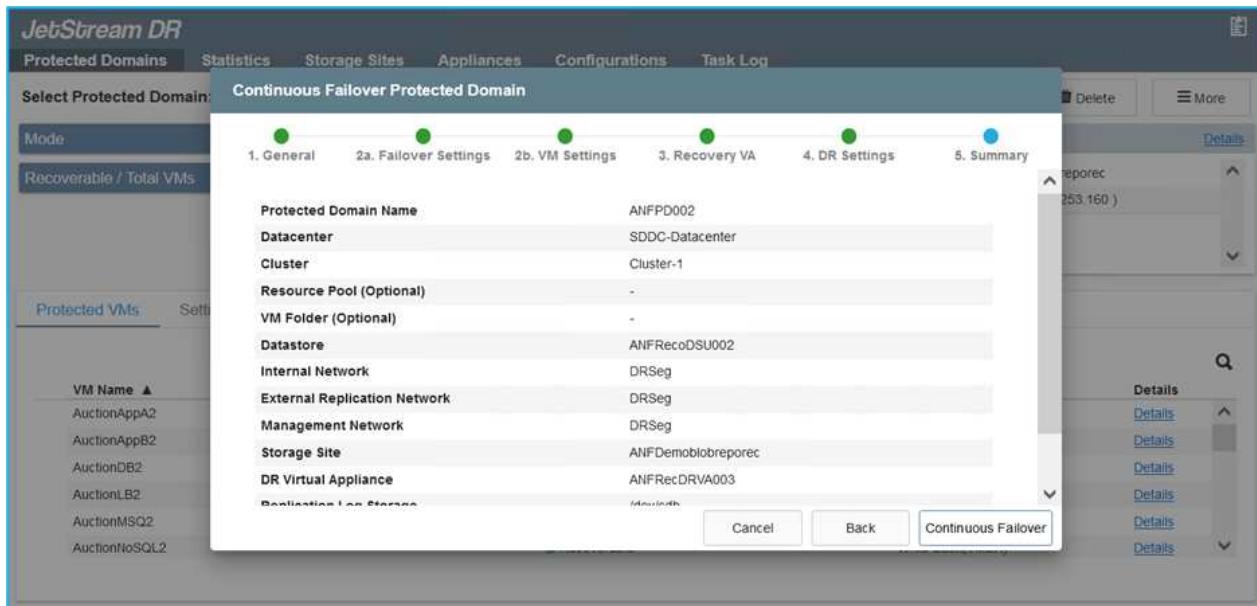
4. The domain is imported for recovery. Go to the Protected Domains tab and verify that the intended domain has been selected or choose the desired one from the Select Protected Domain menu. A list of the recoverable VMs in the protected domain is displayed.

5. After the protected domains are imported, deploy DRVA appliances.



These steps can also be automated using CPT- created plans.

6. Create replication log volumes using available vSAN or ANF datastores.
7. Import the protected domains and configure the recovery VA to use an ANF datastore for VM placements.



Make sure that DHCP is enabled on the selected segment and that enough IPs are available. Dynamic IPs are temporarily used while domains are recovering. Each recovering VM (including continuous rehydration) requires an individual dynamic IP. After recovery is complete, the IP is released and can be reused.

8. Select the appropriate failover option (continuous failover or failover). In this example, continuous rehydration (continuous failover) is selected.



Although Continuous Failover and Failover modes differ on when configuration is performed, both failover modes are configured using the same steps. Failover steps are configured and performed together in response to a disaster event. Continuous failover can be configured at any time and then allowed to run in the background during normal system operation. After a disaster event has occurred, continuous failover is completed to immediately transfer ownership of the protected VMs to the recovery site (near-zero RTO).

JetStream DR

Protected Domains Statistics Storage Sites Appliances Configurations Task Log

Select Protected Domain: GCSDRPD_Demo01 [View all](#)

Mode Imported

Recoverable / Total VMs 5 / 5

Configurations

Storage Site ANFDemoblobrepor
Owner Site REMOTE (172.21.253.10)

+ Create Delete More

Restore → Failover → Continuous Failover → Test Failover

Protected VMs Settings Alarms

VM Name Protection Status Protection Mode Details

VM Name	Protection Status	Protection Mode	Details
GCS-DR-DC	Recoverable	Write-Back(VMDDK)	Details
GCS-DR-LinVM01	Recoverable	Write-Back(VMDDK)	Details
GCS-DR-SCA	Recoverable	Write-Back(VMDDK)	Details
GCS-DR-SQL01	Recoverable	Write-Back(VMDDK)	Details
GCS-DR-WinVM01	Recoverable	Write-Back(VMDDK)	Details

The continuous failover process begins, and its progress can be monitored from the UI. Clicking the blue icon in the Current Step section exposes a pop-up window showing details of the current step of the failover process.

Failover and Failback

1. After a disaster occurs in the protected cluster of the on-premises environment (partial or complete failure), you can trigger the failover for VMs using Jetstream after breaking the SnapMirror relationship for the respective application volumes.

Health Status	Source Volume	Target Volume	Total Transfer Time	Status	Mirror State	Last Successful Transfer
Green	gcsdrsqldb_sc46_ntaphci-a300e9u25	gcsdrsqldb_sc46_copy_ANFCVODRDemo	6 minutes 41 seconds	idle	snapmirrored	May 5, 2022, 12:08:34 PM 33.66 KB
Green	gcsdrsqlhld_sc46_ntaphci-a300e9u25	gcsdrsqlhld_sc46_copy_ANFCVODRDemo	4 minutes 56 seconds	idle	snapmirrored	Information
Green	gcsdrsqllog_sc46_ntaphci-a300e9u25	gcsdrsqllog_sc46_copy_ANFCVODRDemo	10 minutes 18 seconds	idle	snapmirrored	Break Reverse Resync Edit Schedule Edit Max Transfer Rate Update Delete

Health Status	Source Volume	Target Volume	Total Transfer Time	Status	Mirror State	Last Successful Transfer
Green	gcsdrsqldb_sc46_ntaphci-a300e9u25	gcsdrsqldb_sc46_copy_ANFCVODRDemo	6 minutes 41 seconds	idle	snapmirrored	May 5, 2022, 12:08:34 PM 33.66 KB
Green	gcsdrsqlhld_sc46_ntaphci-a300e9u25	gcsdrsqlhld_sc46_copy_ANFCVODRDemo	4 minutes 56 seconds	idle	snapmirrored	May 5, 2022, 12:09:15 PM 69.84 KB
Green	gcsdrsqllog_sc46_ntaphci-a300e9u25	gcsdrsqllog_sc46_copy_ANFCVODRDemo	10 minutes 18 seconds	idle	snapmirrored	May 5, 2022, 12:08:34 PM 104.34 KB



This step can easily be automated to facilitate the recovery process.

2. Access the Jetstream UI on AVS SDDC (destination side) and trigger the failover option to complete failover. The task bar shows progress for failover activities.

In the dialog window that appears when completing failover, the failover task can be specified as planned or assumed to be forced.

JetStream DR

Protected Domains Statistics Storage Sites Appliances Configurations Task Log

Select Protected Domain: GCSDRPD_Demo01 [View all](#)

Mode: Continuous Rehydration in Progress 4 / 4

Recoverable / Total VMs: 4 / 4

Data (Processed/Known Remaining): 329.01 GB / 6.19 GB

Current Step: Recover VMs' data from Storage Site

[+ Create](#) [Failover](#) [More](#)

Configurations

Storage Site	ANFDemobilobrerec
Owner Site	REMOTE (172.21.253.160)
Datacenter / Cluster	SDDC-Datacenter / Cluster-1
Point-in-time Recovery	Disabled

Protected VMs

VM Name	Protection Status	Protection Mode	Details
GCS-DR-DC	Recoverable	Write-Back(VMDK)	Details
GCS-DR-LinVM01	Recoverable	Write-Back(VMDK)	Details
GCS-DR-SCA	Recoverable	Write-Back(VMDK)	Details
GCS-DR-SQL01	Recoverable	Write-Back(VMDK)	Details
GCS-DR-WinVM01	Recoverable	Write-Back(VMDK)	Details

Complete Continuous Failover for Protected Domain

VM Network Mapping

Protected VM Network	Recovery VM Network
VM_3510	DRStretchSeg

Other Settings

Planned Failover

Force Failover

Some VMs' guest credential are required because of network configuration: [Configure](#)

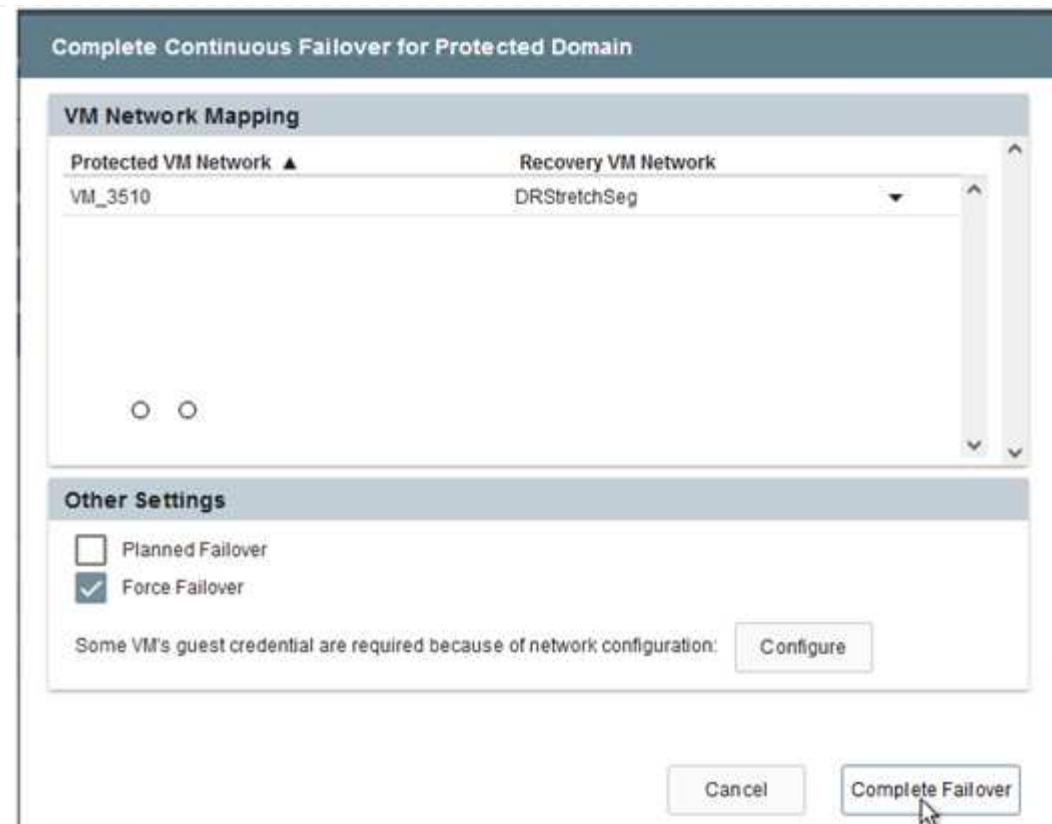
[Cancel](#) [Complete Failover](#)

Forced failover assumes the primary site is no longer accessible and ownership of the protected domain should be directly assumed by the recovery site.

Force Failover

! Force Failover of Protected Domain requested. Administrator consent is required!
Complete ownership of this Protected Domain will be taken over by this Site.
Are you sure you want to continue?

[Cancel](#) [Confirm](#)



3. After continuous failover is complete, a message appears confirming completion of the task. When the task is complete, access the recovered VMs to configure iSCSI or NFS sessions.



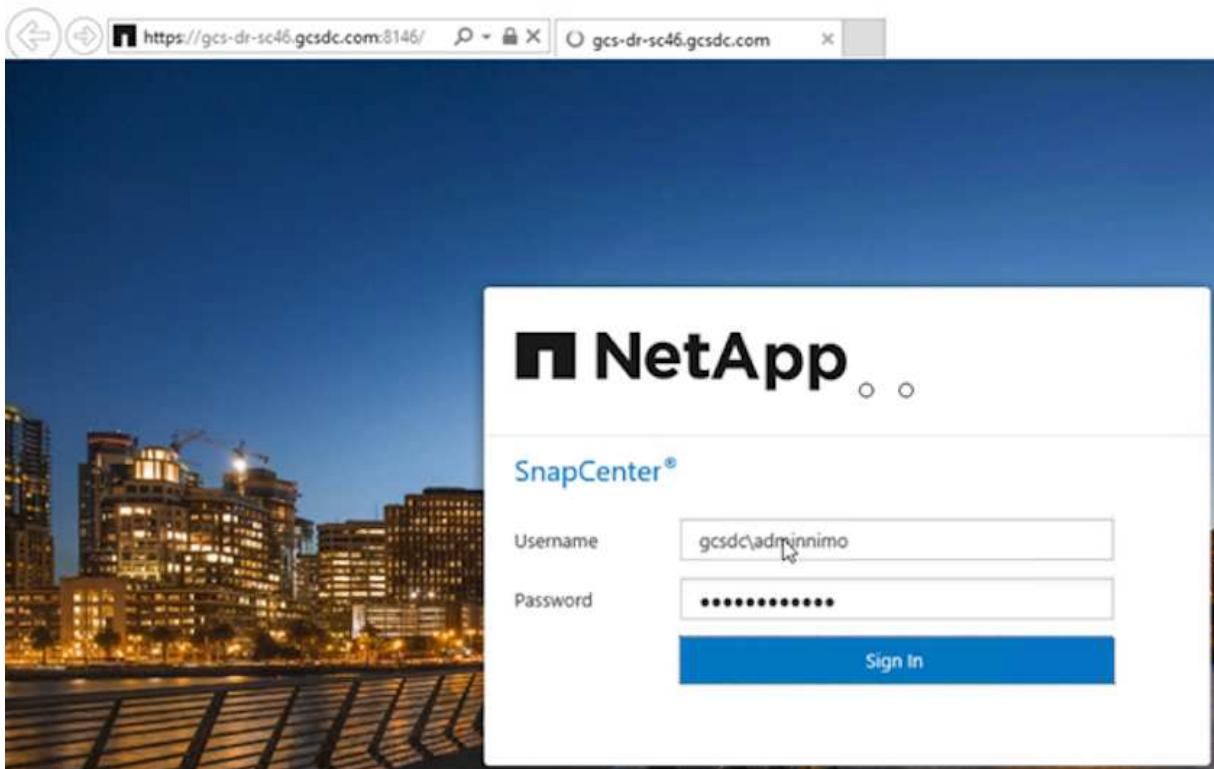
The failover mode changes to Running in Failover and the VM status is Recoverable. All the VMs of the protected domain are now running at the recovery site in the state specified by the failover runbook settings.



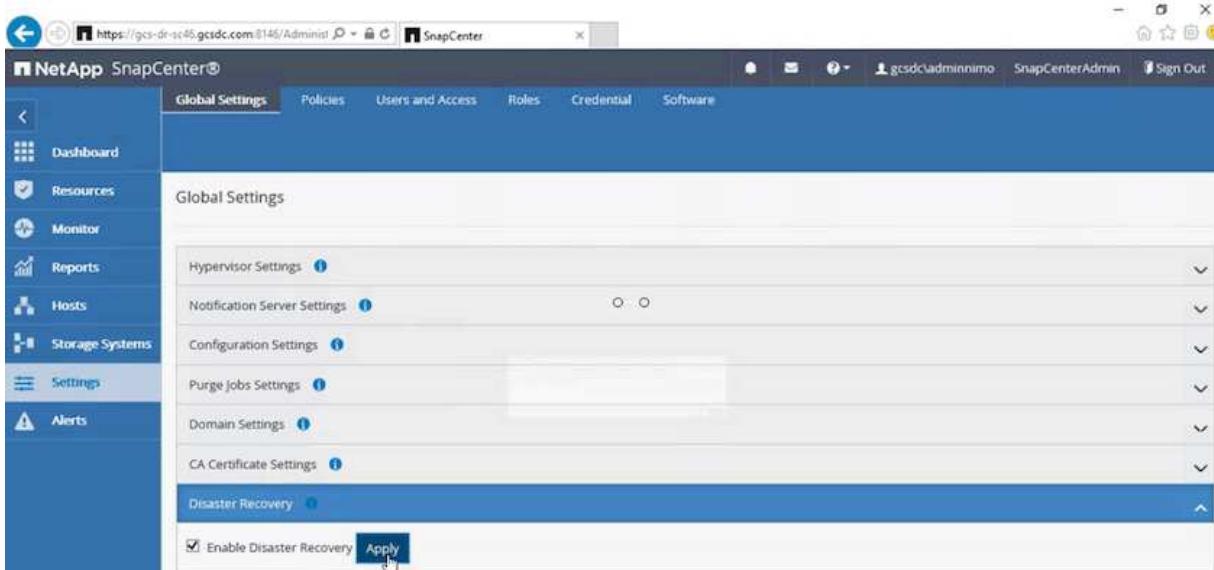
To verify the failover configuration and infrastructure, JetStream DR can be operated in test mode (Test Failover option) to observe the recovery of virtual machines and their data from the object store into a test recovery environment. When a failover procedure is executed in test mode, its operation resembles an actual failover process.

Protected Domain	VMs Recovery Status
GCSDRPD002	Success with warnings
	Total VMs Recovered: 4
	VM(s) with warning: 2 View
GCSRRecovery03 Status:	
Pre-script Execution Status	Not defined
Runbook Execution Status	Success
Post-script Execution Status	Not defined

4. After the virtual machines are recovered, use storage disaster recovery for in-guest storage. To demonstrate this process, SQL server is used in this example.
5. Log into the recovered SnapCenter VM on AVS SDDC and enable DR mode.
 - a. Access the SnapCenter UI using the browser.



- b. In the Settings page, navigate to Settings > Global Settings > Disaster Recovery.
- c. Select Enable Disaster Recovery.
- d. Click Apply.



- e. Verify whether the DR job is enabled by clicking Monitor > Jobs.



NetApp SnapCenter 4.6 or later should be used for storage disaster recovery. For previous versions, application-consistent snapshots (replicated using SnapMirror) should be used and manual recovery should be executed in case previous backups must be recovered in the disaster recovery site.

6. Make sure that the SnapMirror relationship is broken.

7. Attach the LUN from Cloud Volumes ONTAP to the recovered SQL guest VM with same drive letters.

8. Open iSCSI Initiator, clear the previous disconnected session and add the new target along with multipath for the replicated Cloud Volumes ONTAP volumes.

iSCSI Initiator Properties

X

Targets Discovery Favorite Targets Volumes and Devices RADIUS Configuration

Quick Connect

To discover and log on to a target using a basic connection, type the IP address or DNS name of the target and then click Quick Connect.

Target:

Quick Connect...

Discovered targets

Refresh

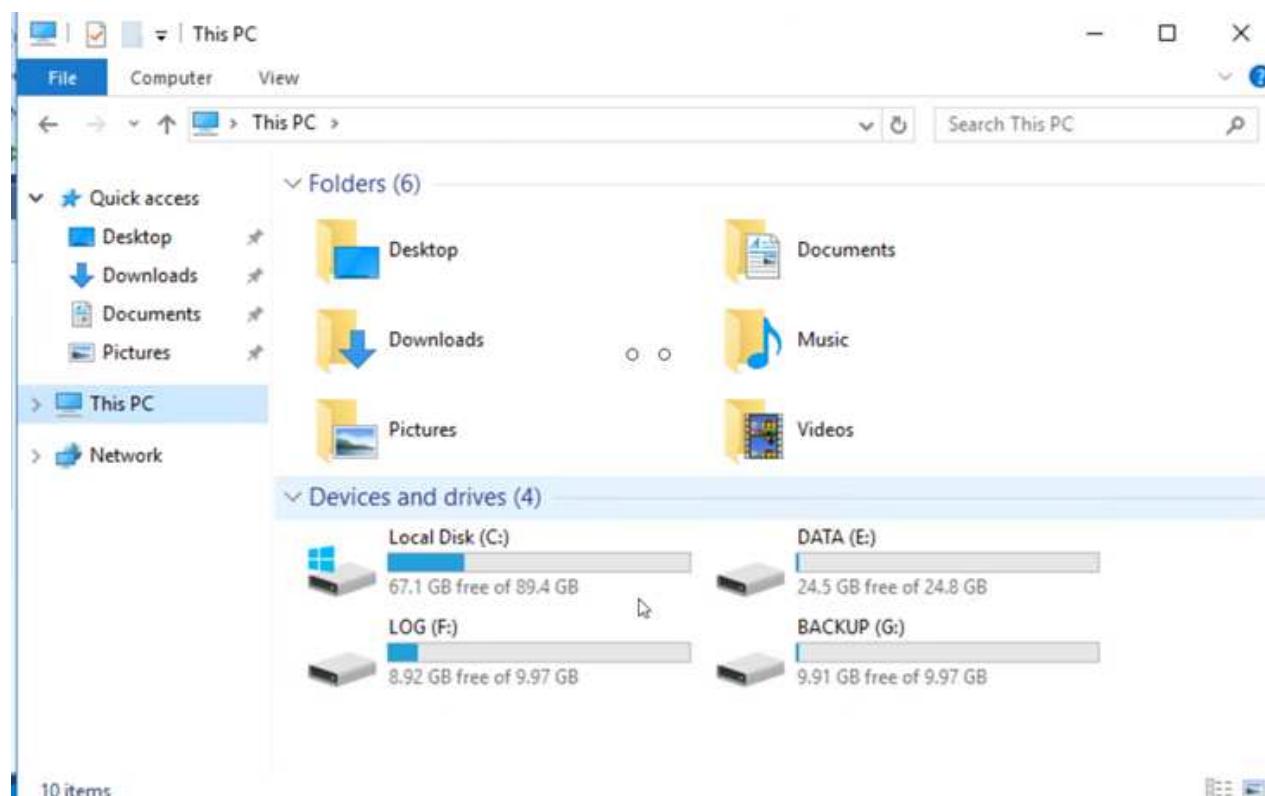
Name

Status

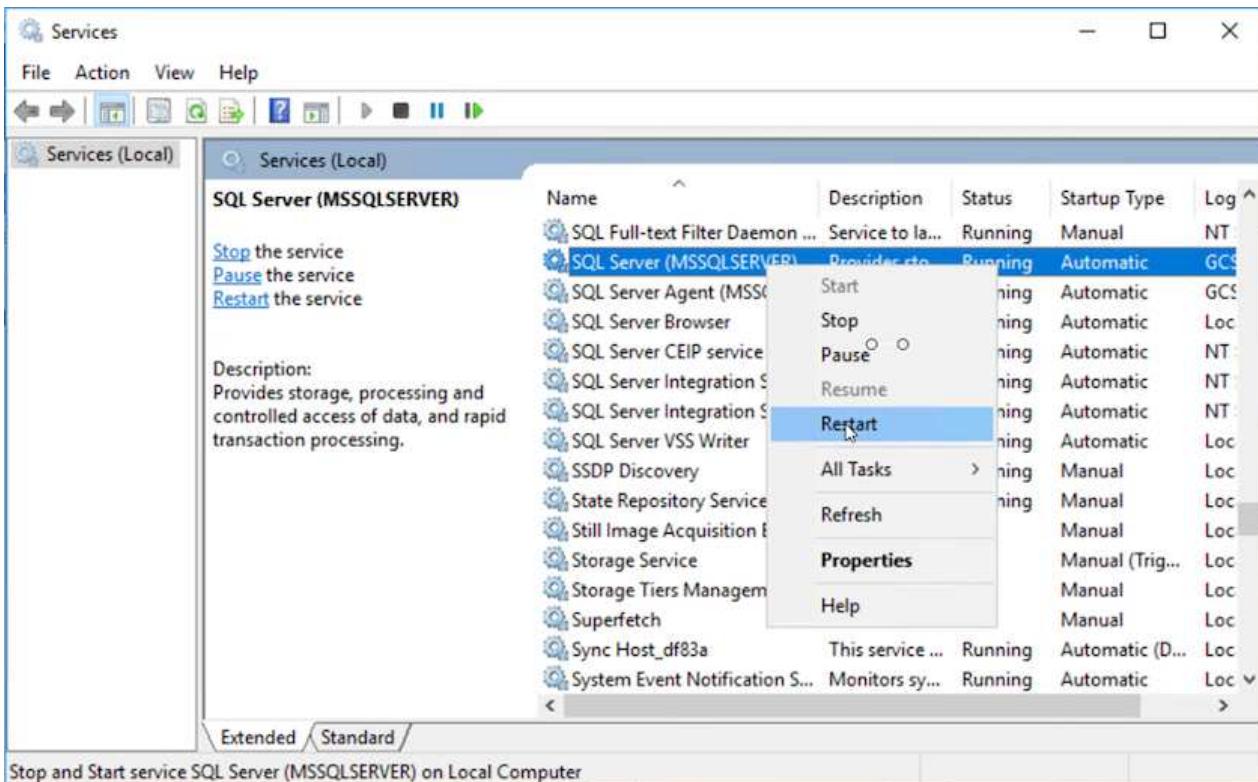
iqn.1992-08.com.netapp:sn.547772ccc47811ecbb62000... Connected

iqn.1992-08.com.netapp:sn.aeab78ab720011ec939800... Reconnecting...

9. Make sure that all the disks are connected using the same drive letters that were used prior to DR.



10. Restart the MSSQL server service.



11. Make sure that the SQL resources are back online.

The screenshot shows the Microsoft SQL Server Management Studio (SSMS) interface. The title bar indicates the query is for 'SQLQuery1.sql - GCS-DR-SQL03.CarDB (GCSDC\adminnnimo (66)) - Microsoft SQL Server Management Studio (Administrator)'. The Object Explorer on the left shows the database structure, including 'CarDB' and its 'dbo.Cars' table. The central pane displays a T-SQL script for a 'SELECT TOP (1000)' query on the 'Cars' table, which retrieves columns 'Id', 'Name', and 'Price'. The results pane shows the following data:

	Id	Name	Price
1	1	Car-1	1000
2	2	Car-2	2000
3	3	Car-3	3000
4	4	Car-4	4000
5	5	Car-5	5000

The status bar at the bottom right of the results pane shows 'Query executed successfully.'



In the case of NFS, attach the volumes using the mount command and update the /etc/fstab entries.

At this point, operations can be run and business continues normally.



On the NSX-T end, a separate dedicated tier-1 gateway can be created for simulating failover scenarios. This ensures that all workloads can communicate with each other but that no traffic can route in or out of the environment, so that any triage, containment, or hardening tasks can be performed without risk of cross-contamination. This operation is outside of the scope of this document, but it can easily be achieved for simulating isolation.

After the primary site is up and running again, you can perform failback. VM protection is resumed by Jetstream and the SnapMirror relationship must be reversed.

1. Restore the on-premises environment. Depending on the type of disaster incident, it might be necessary to restore and/or verify the configuration of the protected cluster. If necessary, JetStream DR software might need to be reinstalled.
2. Access the restored on-premises environment, go to the Jetstream DR UI, and select the appropriate protected domain. After the protected site is ready for failback, select the Failback option in the UI.



The CPT-generated failback plan can also be used to initiate the return of the VMs and their data from the object store back to the original VMware environment.

VM Name	Protection Status	Protection Mode	Details
GCS-DR-DC	Recoverable	Write-Back(VMDK)	Details
GCS-DR-Lin\M01	Recoverable	Write-Back(VMDK)	Details
GCS-DR-SCA	Recoverable	Write-Back(VMDK)	Details
GCS-DR-SQL01	Recoverable	Write-Back(VMDK)	Details
GCS-DR-Win\M01	Recoverable	Write-Back(VMDK)	Details



Specify the maximum delay after pausing the VMs in the recovery site and restarting them in the protected site. The time need to complete this process includes the completion of replication after stopping failover VMs, the time needed to clean the recovery site, and the time needed to recreate VMs in the protected site. NetApp recommends 10 minutes.

Fallback Protected Domain

1. General 2a. Fallback Settings 2b. VM Settings 3. Recovery VA 4. DR Settings 5. Summary

Fallback Datacenter	A300-DataCenter
Fallback Cluster	A300-Cluster
Fallback Resource Pool	-
VM Folder (Optional)	-
Fallback Datastore	A300_NFS_vMotion
Maximum Delay After Stopping	10 Minutes
Internal Network	VM_187
External Replication Network	VM_187
Management Network	VM_187
Storage Site	ANFCVODR
DR Virtual Appliance	GCSDRVA002
Replication Log Storage	/dev/sdb

Buttons: Cancel, Back, Fallback

3. Complete the failback process and then confirm the resumption of VM protection and data consistency.

JetStream DR

Protected Domains Statistics Storage S...

Select Protected Domain: GCSDRPD002

Recoverable / Total VMs

Replication Status

Remaining Background Data

Current RPO

Protected VMs Settings Alarms

Fallback Task Result

Task Completed Successfully

Protected Domain	GCSDRPD002
VMs Recovery Status	Success
Total VMs Recovered	4
GCSR03 Status:	
Pre-script Execution Status	Not defined
Runbook Execution Status	Success
Post-script Execution Status	Not defined

4. After the VMs are recovered, disconnect the secondary storage from the host and connect to the primary storage.

Health Status	Source Volume	Target Volume	Total Transfer Time	Status	Mirror State	Last Successful Transfer	More
✓	gcsdrsqldb_sc46_ntaphci-a300e9u25	gcsdrsqldb_sc46_copy_ANFCVODRDemo	6 minutes 41 seconds	idle	broken-off	May 5, 2022, 12:08:34 PM 33.66 kB	...
✓	gcsdrsqlhld_sc46_ntaphci-a300e9u25	gcsdrsqlhld_sc46_copy_ANFCVODRDemo	4 minutes 56 seconds	idle	broken-off		Information
✓	gcsdrsqllog_sc46_ntaphci-a300e9u25	gcsdrsqllog_sc46_copy_ANFCVODRDemo	10 minutes 18 seconds	idle	broken-off		Resync Reverse Resync Edit Schedule Edit Max Transfer Rate Delete

5. Restart the MSSQL server service.
6. Verify that the SQL resources are back online.



To failback to the primary storage, make sure that the relationship direction remains the same as it was before the failover by performing a reverse resync operation.



To retain the roles of primary and secondary storage after the reverse resync operation, perform the reverse resync operation again.

This process is applicable to other applications like Oracle, similar database flavors, and any other applications using guest-connected storage.

As always, test the steps involved for recovering the critical workloads before porting them into production.

Benefits of this solution

- Uses the efficient and resilient replication of SnapMirror.
- Recovers to any available points in time with ONTAP snapshot retention.
- Full automation is available for all required steps to recover hundreds to thousands of VMs, from the storage, compute, network, and application validation steps.
- SnapCenter uses cloning mechanisms that do not change the replicated volume.
 - This avoids the risk of data corruption for volumes and snapshots.
 - Avoids replication interruptions during DR test workflows.
 - Leverages the DR data for workflows beyond DR, such as dev/test, security testing, patch and upgrade testing, and remediation testing.
- CPU and RAM optimization can help lower cloud costs by enabling recovery to smaller compute clusters.

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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