



Protecting Workloads on Azure / AWS

NetApp Solutions

NetApp
July 31, 2024

Table of Contents

- Protecting Workloads on Azure / AVS 1
 - Disaster Recovery with ANF and JetStream 1
 - Disaster Recovery with CVO and AVS (guest-connected storage) 13
 - TR-4955: Disaster Recovery with Azure NetApp Files (ANF) and Azure VMware Solution (AVS) 39
 - Using Veeam Replication and Azure NetApp Files datastore for disaster recovery to Azure VMware Solution 53

Protecting Workloads on Azure / AVS

Disaster Recovery with ANF and JetStream

Disaster recovery to cloud is a resilient and cost-effective way of protecting the workloads against site outages and data corruption events (for example, ransomware). Using the VMware VAIO framework, on-premises VMware workloads can be replicated to Azure Blob storage and recovered, enabling minimal or close to no data loss and near-zero RTO.

JetStream DR can be used to seamlessly recover the workloads replicated from on-premises to AVS and specifically to Azure NetApp Files. It enables cost-effective disaster recovery by using minimal resources at the DR site and cost-effective cloud storage. JetStream DR automates recovery to ANF datastores via Azure Blob Storage. JetStream DR recovers independent VMs or groups of related VMs into recovery site infrastructure according to network mapping and provides point-in-time recovery for ransomware protection.

This document provides an understanding of the JetStream DR principles of operations and its main components.

Solution deployment overview

1. Install JetStream DR software in the on-premises data center.
 - a. Download the JetStream DR software bundle from Azure Marketplace (ZIP) and deploy the JetStream DR MSA (OVA) in the designated cluster.
 - b. Configure the cluster with the I/O filter package (install JetStream VIB).
 - c. Provision Azure Blob (Azure Storage Account) in the same region as the DR AVS cluster.
 - d. Deploy DRVA appliances and assign replication log volumes (VMDK from existing datastore or shared iSCSI storage).
 - e. Create protected domains (groups of related VMs) and assign DRVAs and Azure Blob Storage/ANF.
 - f. Start protection.
2. Install JetStream DR software in the Azure VMware Solution private cloud.
 - a. Use the Run command to install and configure JetStream DR.
 - b. Add the same Azure Blob container and discover domains using the Scan Domains option.
 - c. Deploy required DRVA appliances.
 - d. Create replication log volumes using available vSAN or ANF datastores.
 - e. Import protected domains and configure RocVA (recovery VA) to use ANF datastore for VM placements.
 - f. Select the appropriate failover option and start continuous rehydration for near-zero RTO domains or VMs.
3. During a disaster event, trigger failover to Azure NetApp Files datastores in the designated AVS DR site.
4. Invoke failback to the protected site after the protected site has been recovered. Before starting, make sure that the prerequisites are met as indicated in this [link](#) and also run the Bandwidth Testing Tool (BWT) provided by JetStream Software to evaluate the potential performance of Azure Blob storage and its replication bandwidth when used with JetStream DR software. After the pre-requisites, including connectivity, are in place, set up and subscribe to JetStream DR for AVS from the [Azure Marketplace](#). After the software bundle is downloaded, proceed with the installation process described above.

When planning and starting protection for a large number of VMs (for example, 100+), use the Capacity Planning Tool (CPT) from the JetStream DR Automation Toolkit. Provide a list of VMs to be protected together with their RTO and recovery group preferences, and then run CPT.

CPT performs the following functions:

- Combining VMs into protection domains according to their RTO.
- Defining the optimal number of DRVAs and their resources.
- Estimating required replication bandwidth.
- Identifying replication log volume characteristics (capacity, bandwidth, and so on).
- Estimating required object storage capacity, and more.




The number and content of domains prescribed depend upon various VM characteristics such as average IOPS, total capacity, priority (which defines failover order), RTO, and others.

Install JetStream DR in On-Premises Datacenter


JetStream DR software consists of three major components: JetStream DR Management Server Virtual Appliance (MSA), DR Virtual Appliance (DRVA), and host components (I/O Filter packages). MSA is used to install and configure host components on the compute cluster and then to administer JetStream DR software. The following list provides a high-level description of the installation process:

How to install JetStream DR for on-premises

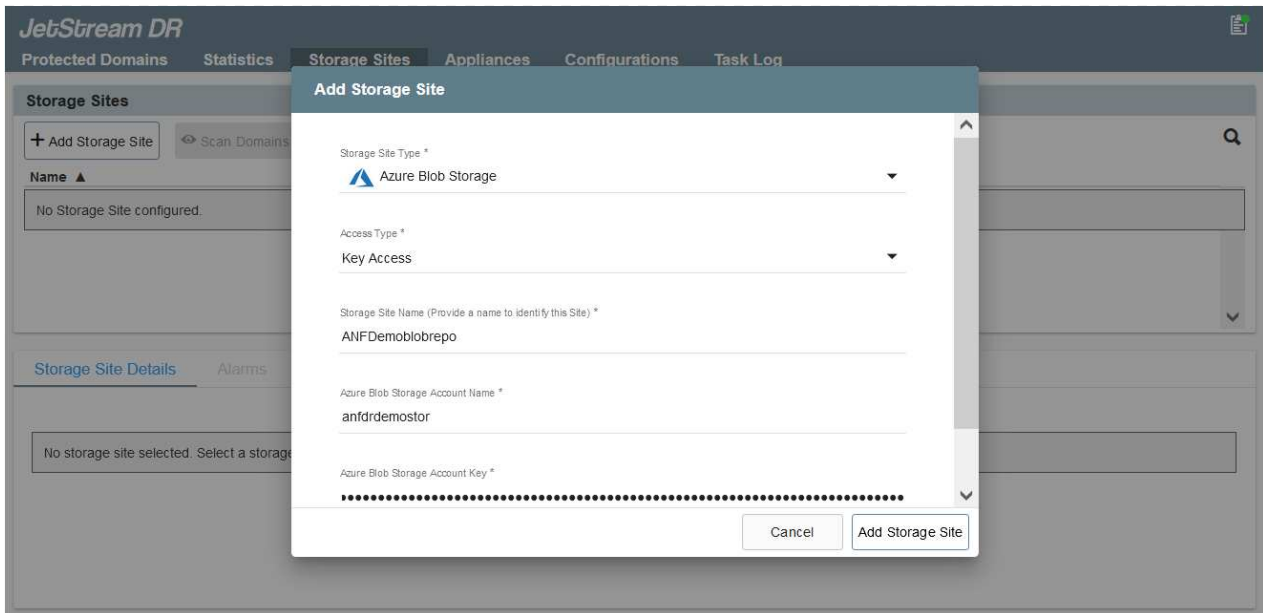
1. Check prerequisites.
2. Run the Capacity Planning Tool for resource and configuration recommendations (optional but recommended for proof-of-concept trials).
3. Deploy the JetStream DR MSA to a vSphere host in the designated cluster.
4. Launch the MSA using its DNS name in a browser.
5. Register the vCenter server with the MSA. To perform the installation, complete the following detailed steps:
6. After JetStream DR MSA has been deployed and the vCenter Server has been registered, access the JetStream DR plug-in using the vSphere Web Client. This can be done by navigating to Datacenter > Configure > JetStream DR.



7. From the JetStream DR interface, select the appropriate cluster.



8. Configure the cluster with the I/O filter package.



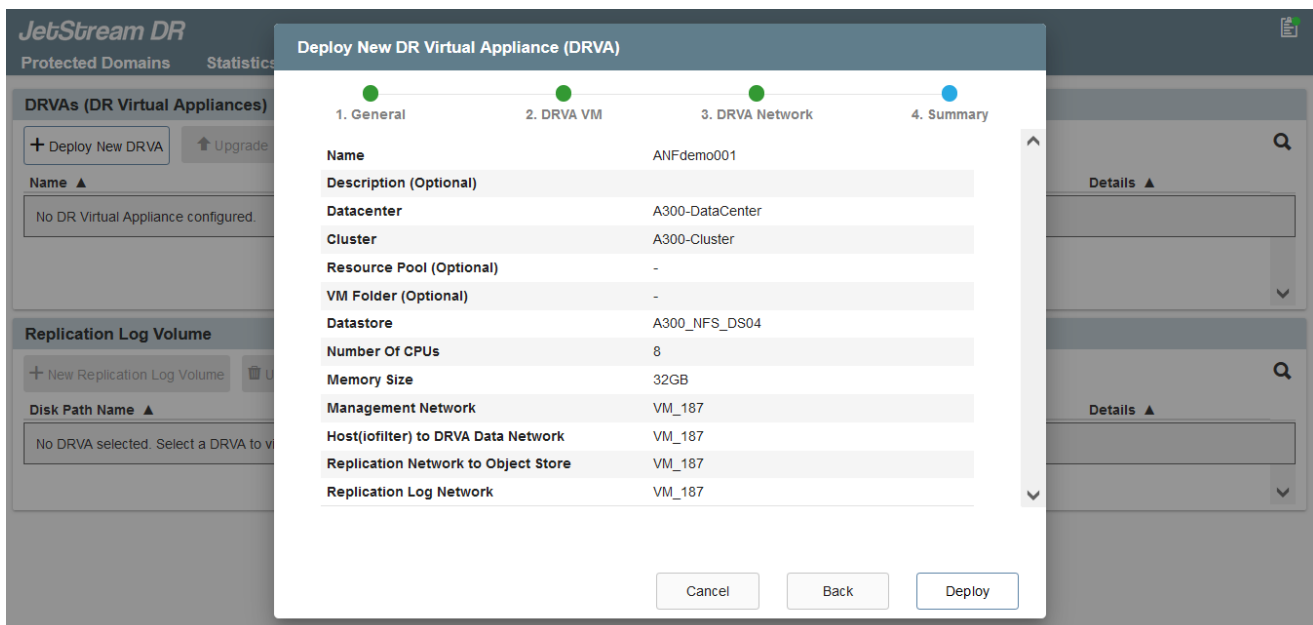
9. Add Azure Blob Storage located at the recovery site.

10. Deploy a DR Virtual Appliance (DRVA) from the Appliances tab.



DRVAs can be automatically created by CPT, but for POC trials we recommend configuring and running the DR cycle manually (start protection > failover > failback).

The JetStream DRVA is a virtual appliance that facilitates key functions in the data replication process. A protected cluster must contain at least one DRVA, and typically one DRVA is configured per host. Each DRVA can manage multiple protected domains.

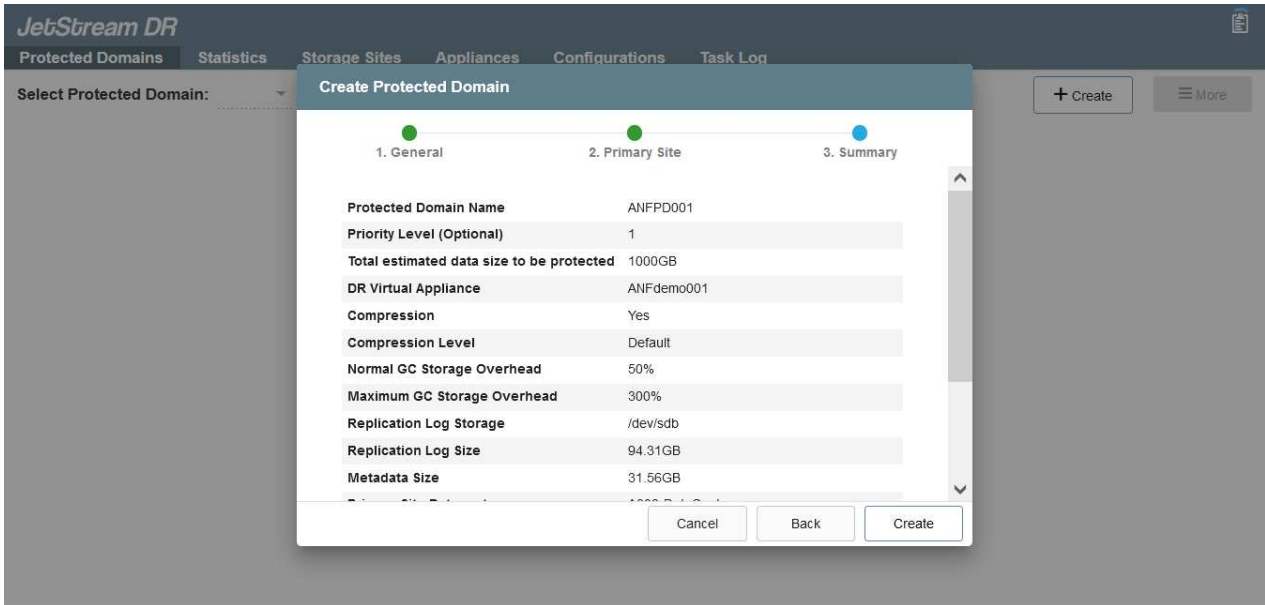


In this example, four DRVA's were created for 80 virtual machines.

1. Create replication log volumes for each DRVA using VMDK Log from the datastores available or independent shared iSCSI storage pools.

2. From the Protected Domains tab, create the required number of protected domains using information

about the Azure Blob Storage site, DRVA instance, and replication log. A protected domain defines a specific VM or set of VMs within the cluster that are protected together and assigned a priority order for failover/failback operations.




3. Select VMs you want to protect and start VM protection of the protected domain. This begins data replication to the designated Blob Store.



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



Write- Back(VMDK) mode can offer higher performance.



Verify that replication log volumes are placed on high performance storage.



Failover run books can be configured to group the VMs (called Recovery Group), set boot order sequence, and modify the CPU/memory settings along with IP configurations.

Install JetStream DR for AVS in an Azure VMware Solution private cloud using the Run command

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 items:

- 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 moreJetStream DR supports 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.



Network configuration including segment creation should be configured on the AVS cluster to match on-premises requirements.

Depending on the SLA and RTO requirements, continuous failover or regular (standard) failover mode can be used. For near-zero RTO, continuous rehydration should be started at the recovery site.

How to install JetStream DR for AVS in a private cloud


To install JetStream DR for AVS on an Azure VMware Solution private cloud, complete the following steps:

1. From the Azure portal, go to the Azure VMware solution, select the private cloud, and select Run command > Packages > JSDR.Configuration.



The default CloudAdmin user in Azure VMware Solution doesn't have sufficient privileges to install JetStream DR for AVS. 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

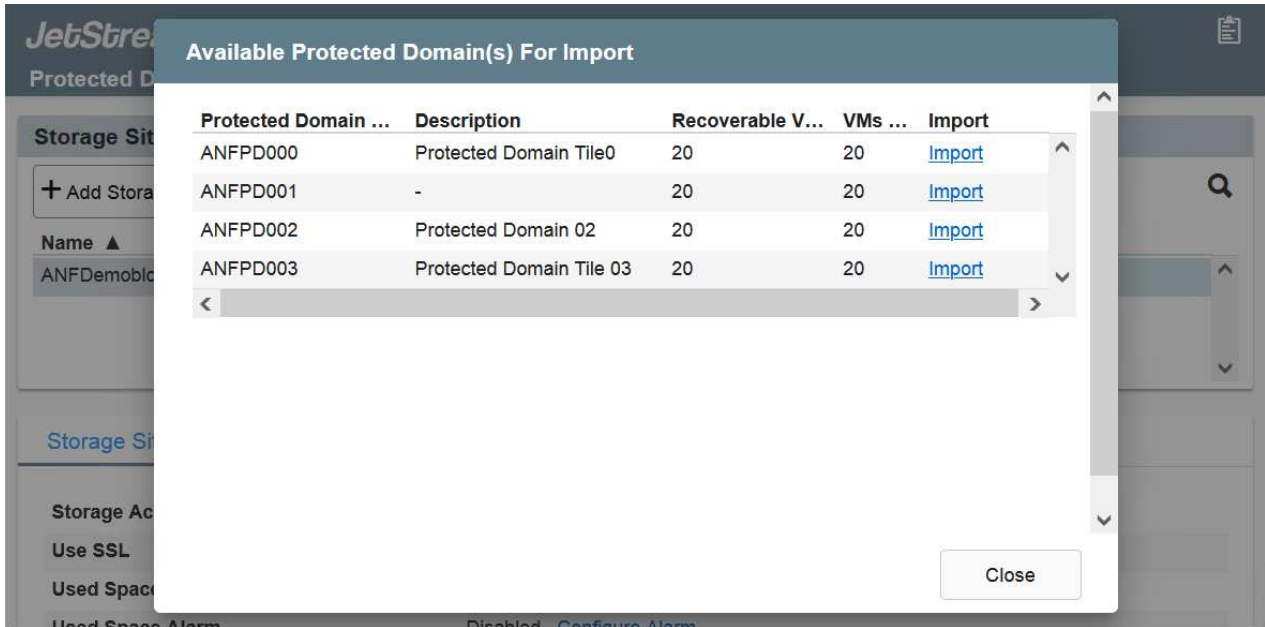
Site Details [Alarm Settings](#)

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

- From the JetStream DR interface, 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.

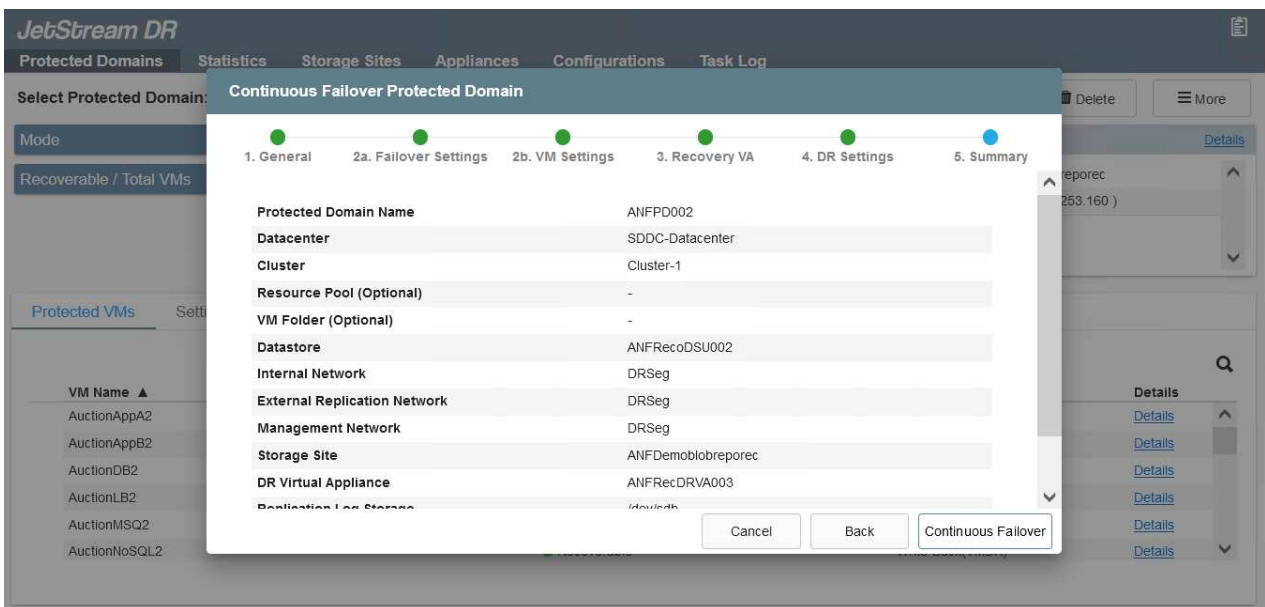


- After the protected domains are imported, deploy DRVA appliances. In this example, continuous rehydration is started manually from the recovery site using the JetStream DR UI.



These steps can also be automated using CPT created plans.

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





Make sure that DHCP is enabled on the selected segment and 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.

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

JetStream DR

Protected Domains | Statistics | Storage Sites | Appliances | Configurations | Task Log

Select Protected Domain: ANFPD000 [View all](#) + Create Delete More

Mode: Imported

Recoverable / Total VMs: 20 / 20

Configurations

- Restore
- Failover
- Continuous Failover
- Test Failover

Protected VMs | Settings | Alarms

VM Name ▲	Protection Status ▲	Protection Mode ▲	Details
AuctionAppA0	✓ Recoverable	Write-Back(VMDK)	Details ^
AuctionAppB0	✓ Recoverable	Write-Back(VMDK)	Details

Performing Failover / Failback

How to perform a Failover / Failback



1. After a disaster occurs in the protected cluster of the on-premises environment (partial or full failure), trigger the failover.



CPT can be used to execute the failover plan to recover the VMs from Azure Blob Storage into the AVS cluster recovery site.




After failover (for continuous or standard rehydration) when the protected VMs have been started in AVS, protection is automatically resumed and JetStream DR continues to replicate their data into the appropriate/original containers in Azure Blob Storage.



The task bar shows progress of failover activities.

- When the task is complete, access the recovered VMs and business continues as normal.



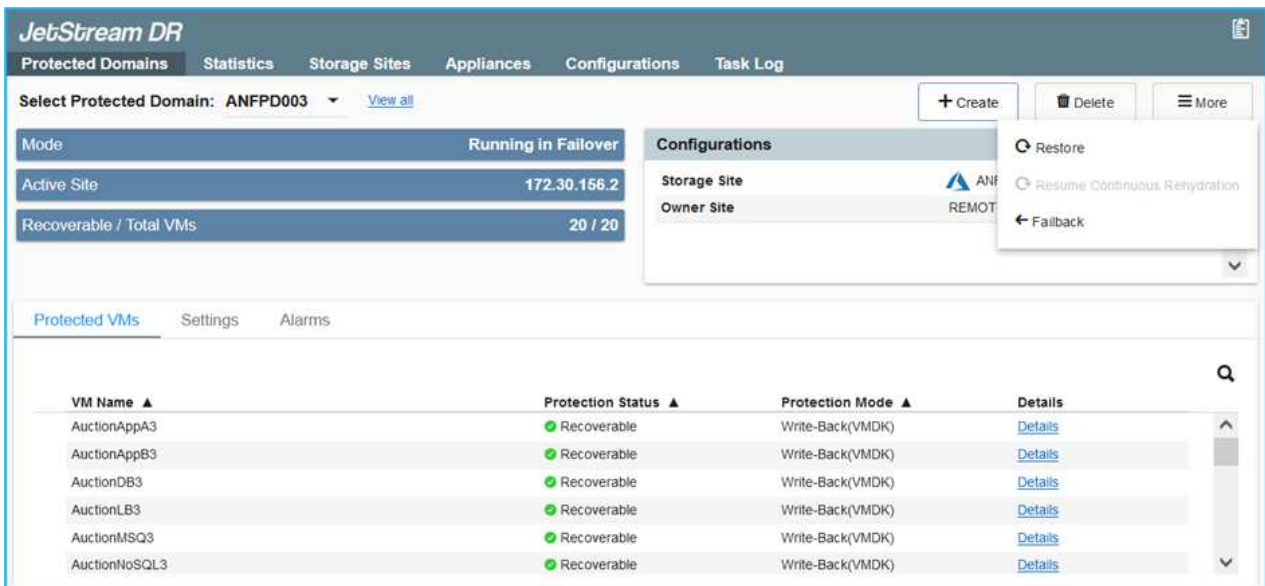
After the primary site is up and running again, failback can be performed. VM protection is resumed and data consistency should be checked.

- Restore the on-premises environment. Depending upon 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.



Note: The `recovery_utility_prepare_failback` script provided in the Automation Toolkit can be used to help clean the original protected site of any obsolete VMs, domain information, and so on.

- 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.



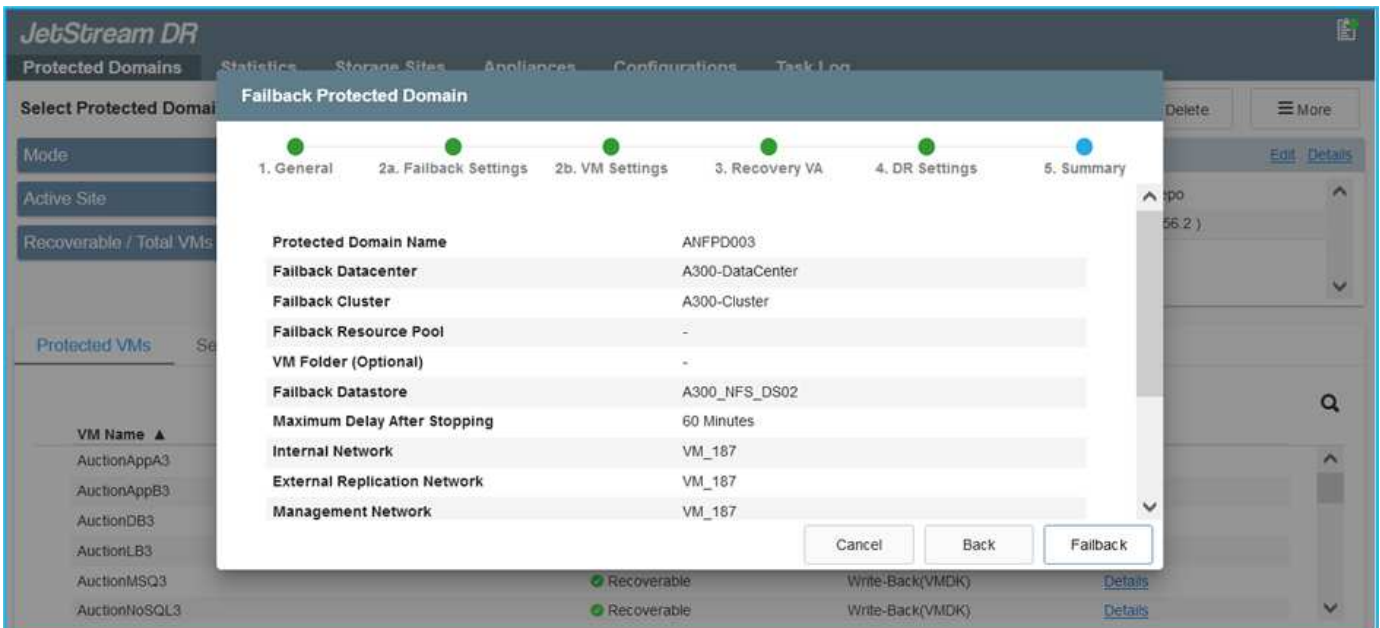
Specify the maximum delay after pausing VMs in the recovery site and restarting in the protected site. This time includes completing replication after stopping failover VMs, the time to clean recovery site, and the time to recreate VMs in protected site. The NetApp recommended value is 10 minutes.

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

Ransomware Recovery

Recovering from ransomware can be a daunting task. Specifically, it can be hard for IT organizations to determine the safe point of return and, once determined, how to ensure that recovered workloads are safeguarded from the attacks reoccurring (from sleeping malware or through vulnerable applications).

JetStream DR for AVS together with Azure NetApp Files datastores can address these concerns by allowing organizations to recover from available points in time, so that workloads are recovered to a functional, isolated network if required. Recovery allows applications to function and communicate with each other while not exposing them to north-south traffic, thereby giving security teams a safe place to perform forensics and other necessary remediation.



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

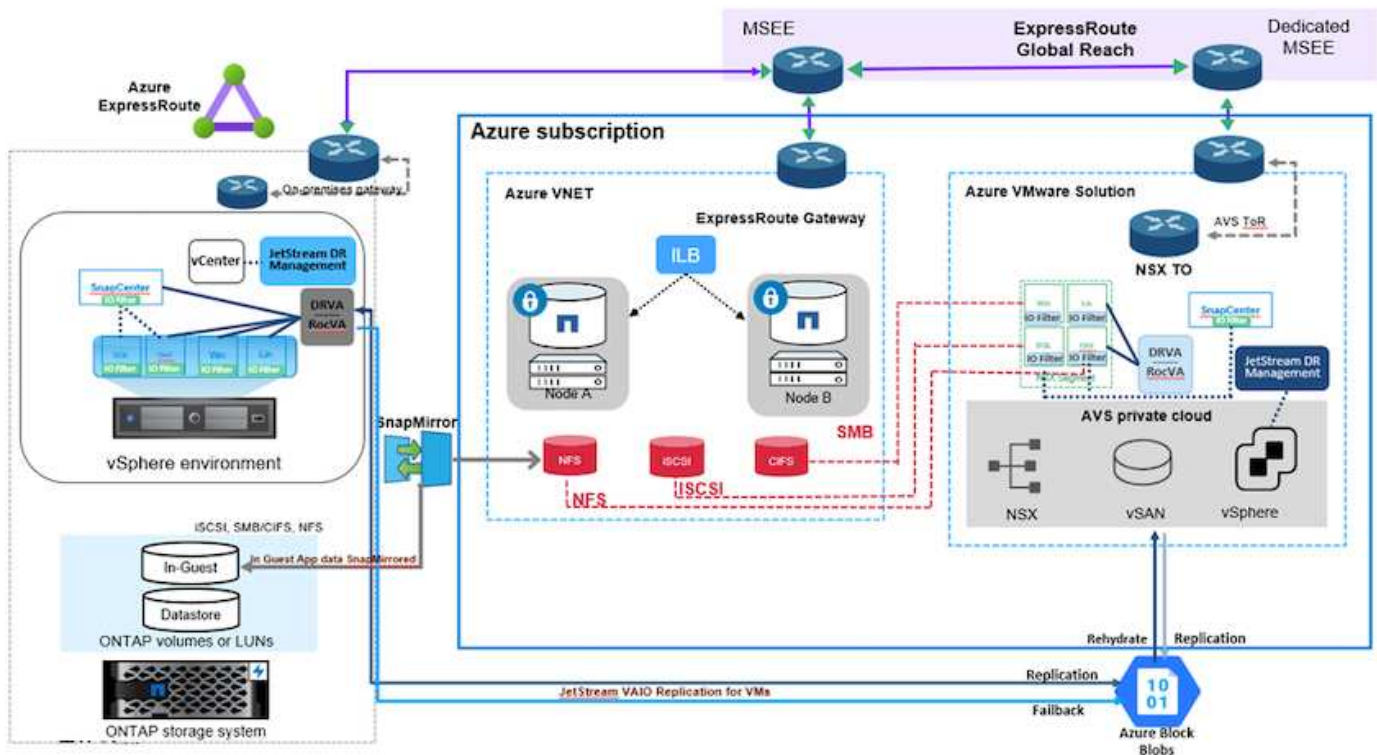
Disaster recovery to cloud is a resilient and cost-effective way of protecting workloads against site outages and data corruption events such as ransomware. With NetApp SnapMirror, on-premises VMware workloads that use guest-connected storage can be replicated to NetApp Cloud Volumes ONTAP running in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

Overview

Authors: Ravi BCB and Niyaz Mohamed, NetApp

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
	gcsdrsqidb_sc46 ntaphci-a300e9u25	gcsdrsqidb_sc46_copy ANFCVODRDemo	17 seconds	idle	snapmirrored	May 6, 2022, 11:43:18 AM 105.06 KiB
	gcsdrsqlihd_sc46_copy ANFCVODRDemo	gcsdrsqlihd_sc46 ntaphci-a300e9u25	7 seconds	idle	snapmirrored	May 6, 2022, 11:42:20 AM 7.22 MiB
	gcsdrsqliog_sc46 ntaphci-a300e9u25	gcsdrsqliog_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.

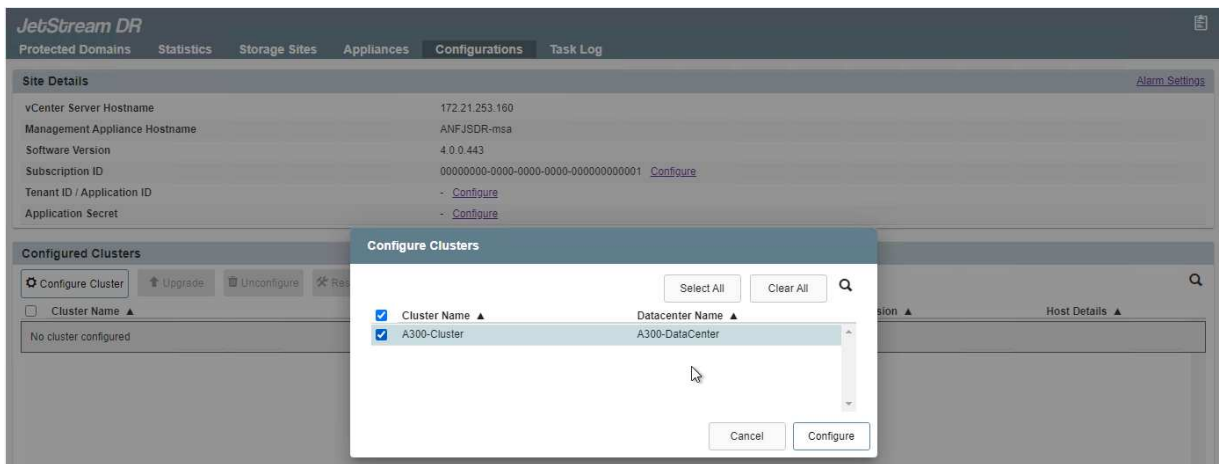
Install JetStream DR in on-premises datacenter

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.



7. From the JetStream DR interface, complete the following tasks:
 - a. Configure the cluster with the I/O filter package.



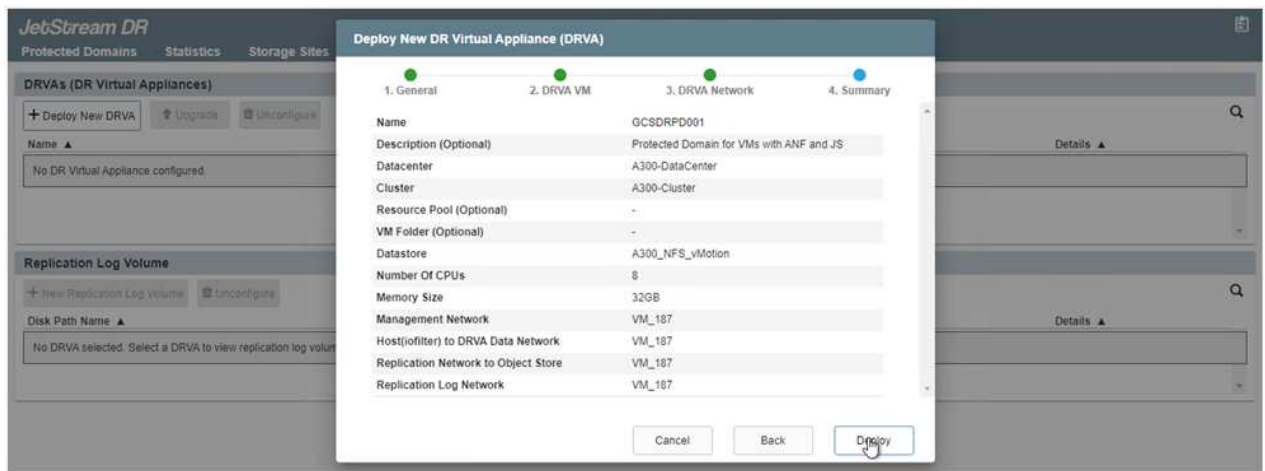
- b. Add the Azure Blob storage located at the recovery site.




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




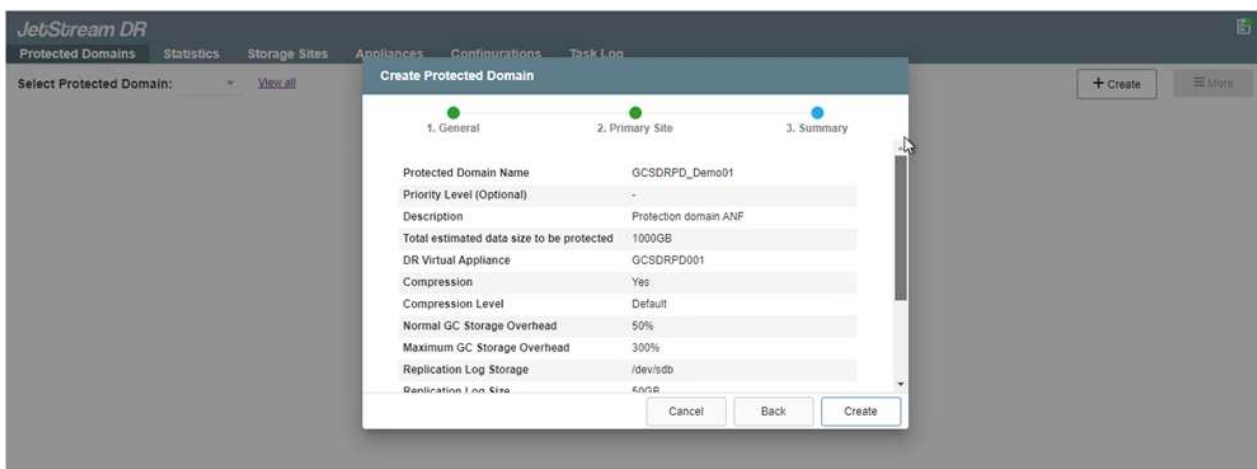
Use the capacity planning tool to estimate the number of DRVAs required.



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



- 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.




- 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.

JetStream DR

Protected Domains | Statistics | Storage Sites | Appliances | Configurations | Task Log

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

Recoverable / Total VMs: 0 / 5

Replication Status: OK

Remaining Background Data: 0 B

Current RPO: -

Configurations

- Storage Site: ANFDRD
- Owner Site: LOCAL (172.2
- Datacenter \ Cluster: A300-DataCen
- Point-in-time Recovery: Disabled

Running Tasks

- Start Protection (GCS-DR-SCA) 50%
- Start Protection (GCS-DR-Win...) 50%
- Start Protection (GCS-DR-Lin...) 50%
- Start Protection (GCS-DR-DC) 50%
- Start Protection (GCS-DR-SQ...) 50%
- Configure VMDK Re... Completed

Close

Protected VMs | Settings | Alarms

+ Start Protection | Stop Protection

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.

JetStream DR

Protected Domains | Statistics | Storage Sites | Appliances | Configurations | Task Log

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

Recoverable / Total VMs: 5 / 5

Replication Status: OK

Remaining Background Data: 0 B

Current RPO: 0s

Configurations

- Storage Site: ANFDRDemoFailoverSite
- Owner Site: LOCAL (172.21.253.160)
- Datacenter \ Cluster: A300-DataCenter \ A300-Cluster
- Point-in-time Recovery: Disabled

Protected VMs | Settings | Alarms

+ Start Protection | Stop Protection

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.

JetStream DR

Protected Domains | Statistics | Storage Sites | Appliances | Configurations | Task Log

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

Recoverable / Total VMs: 5 / 5

Replication Status: OK

Remaining Background Data: 0 B

Current RPO: 0s

Configurations

- Storage Site: ANFDRDemoFailoverSite
- Owner Site: LOCAL (172.21.253.160)
- Datacenter \ Cluster: A300-DataCenter \ A300-Cluster
- Point-in-time Recovery: Disabled

Protected VMs | Settings | Alarms

Failover Runbook: Not Configured [Configure](#)

Test Failover Runbook: Not Configured [Configure](#)

Fallback Runbook: Not Configured [Configure](#)

Memory Setting: Not Configured [Configure](#)

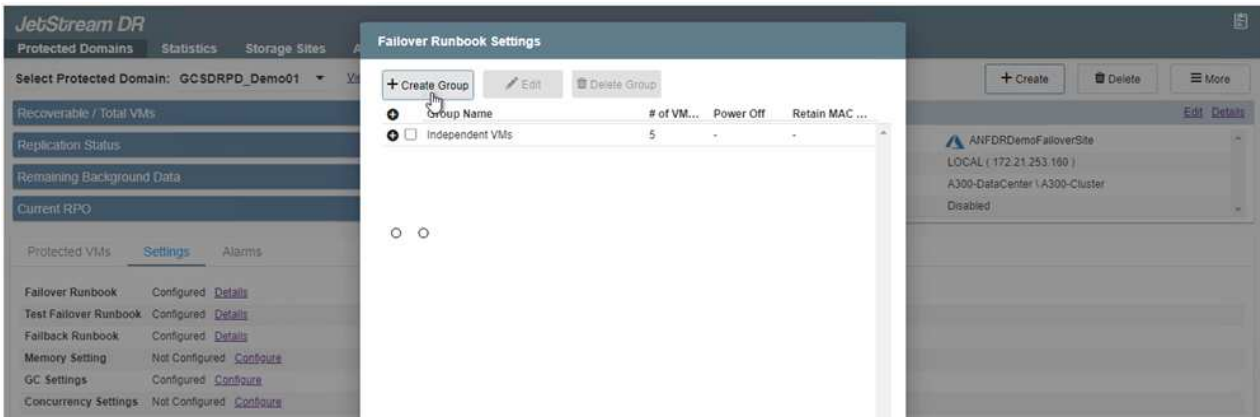
GC Settings: Configured [Configure](#)

Concurrency Settings: Not Configured [Configure](#)

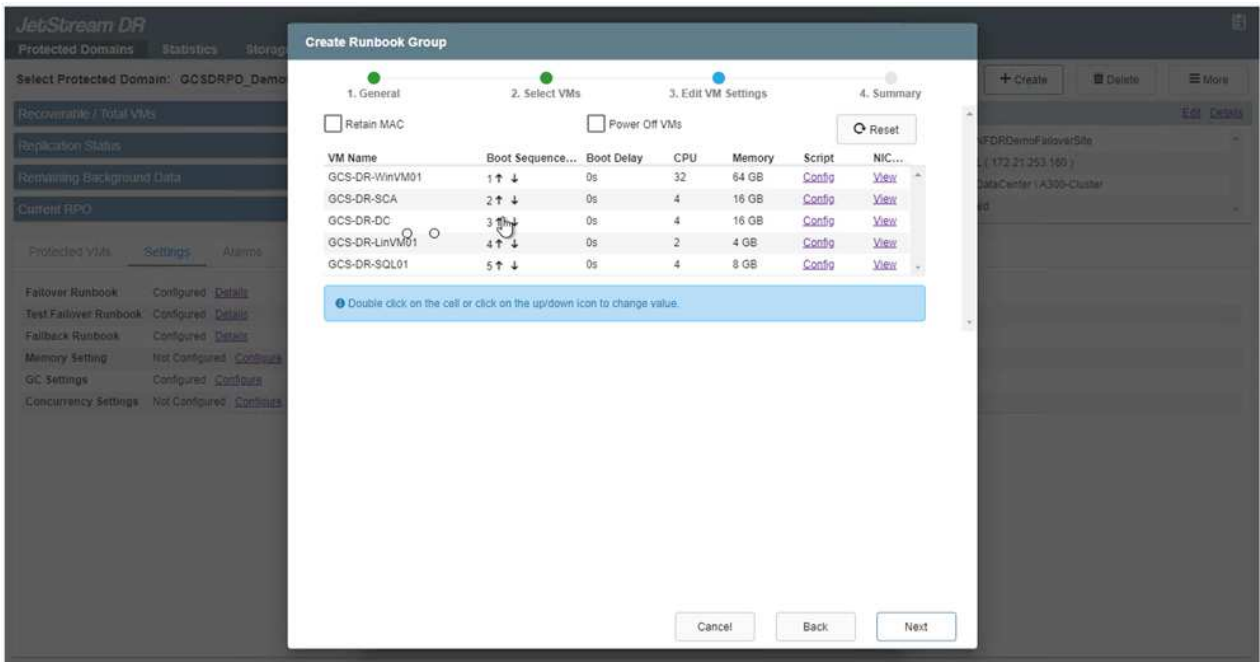
16. Click the Create Group button to begin creating a new runbook group.



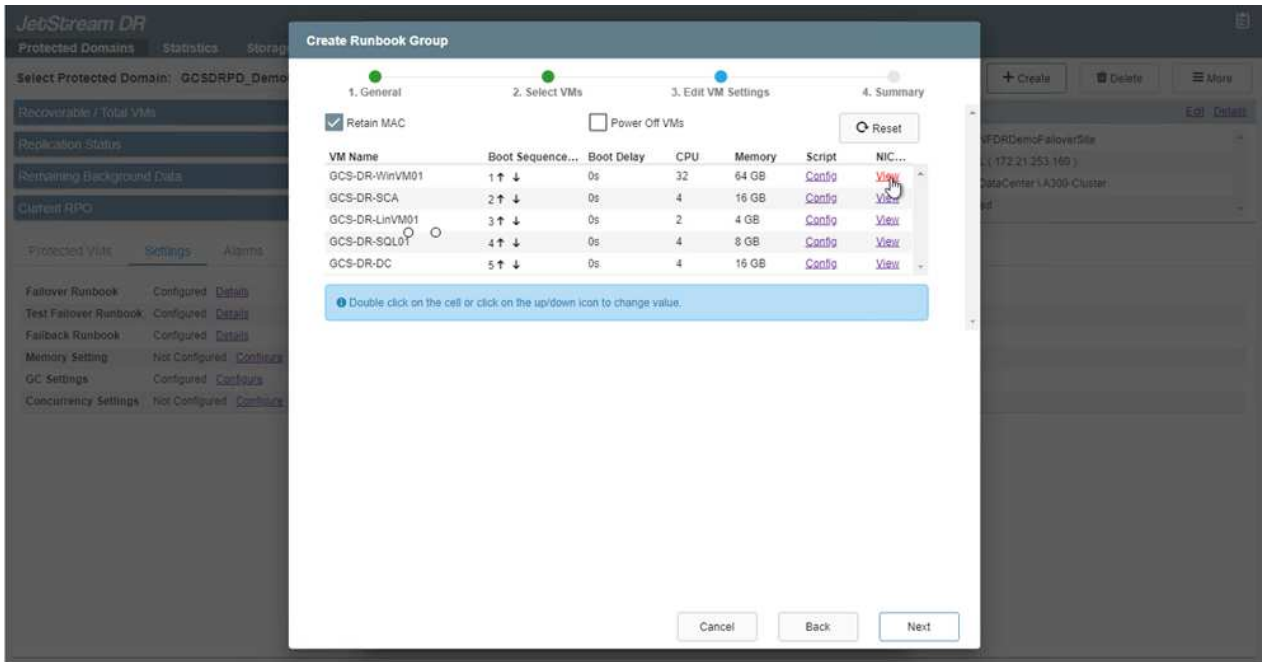
If needed, in the lower portion of the screen, apply custom pre-scripts and post-scripts to automatically run prior to and following operation of the runbook group. Make sure that the Runbook scripts are residing on the management server.



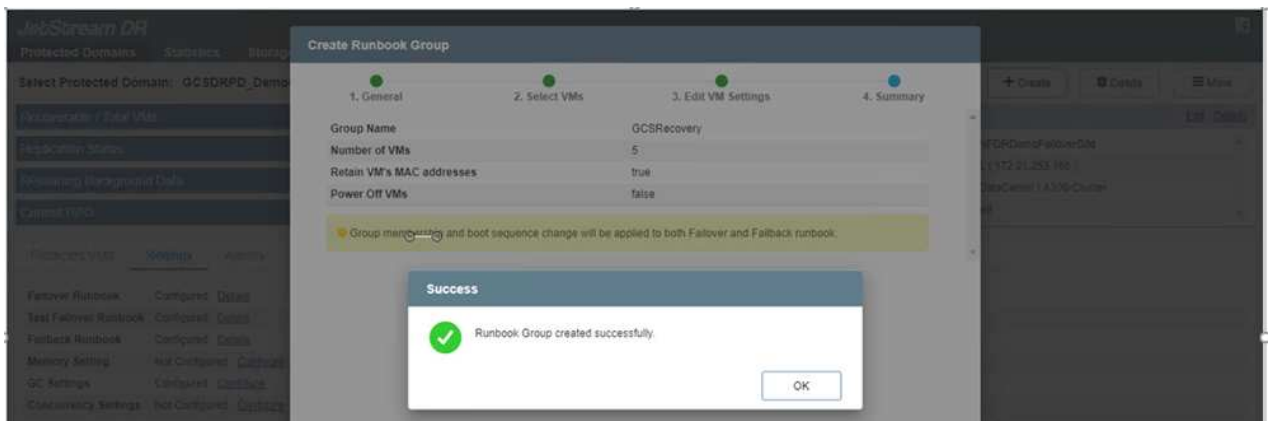

17. Edit the VM settings as required. Specify the parameters for recovering the VMs, including the boot sequence, the boot delay (specified in seconds), the number of CPUs, and the amount of memory to allocate. Change the boot sequence of the VMs by clicking the up or down arrows. Options are also provided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



Network configuration including segment creation should be configured on the AVS cluster to match on-premises requirements.



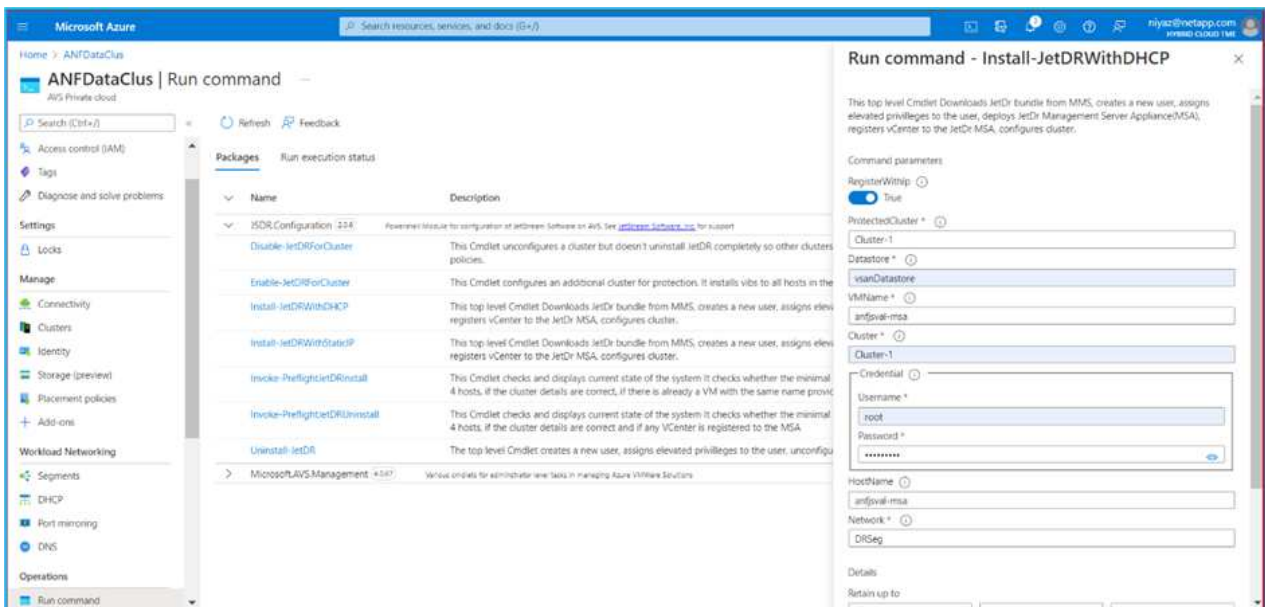
Depending on the SLA and RTO requirements, you can use continuous failover or regular (standard) failover mode. For near-zero RTO, you should start continuous rehydration at the recovery site.

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.

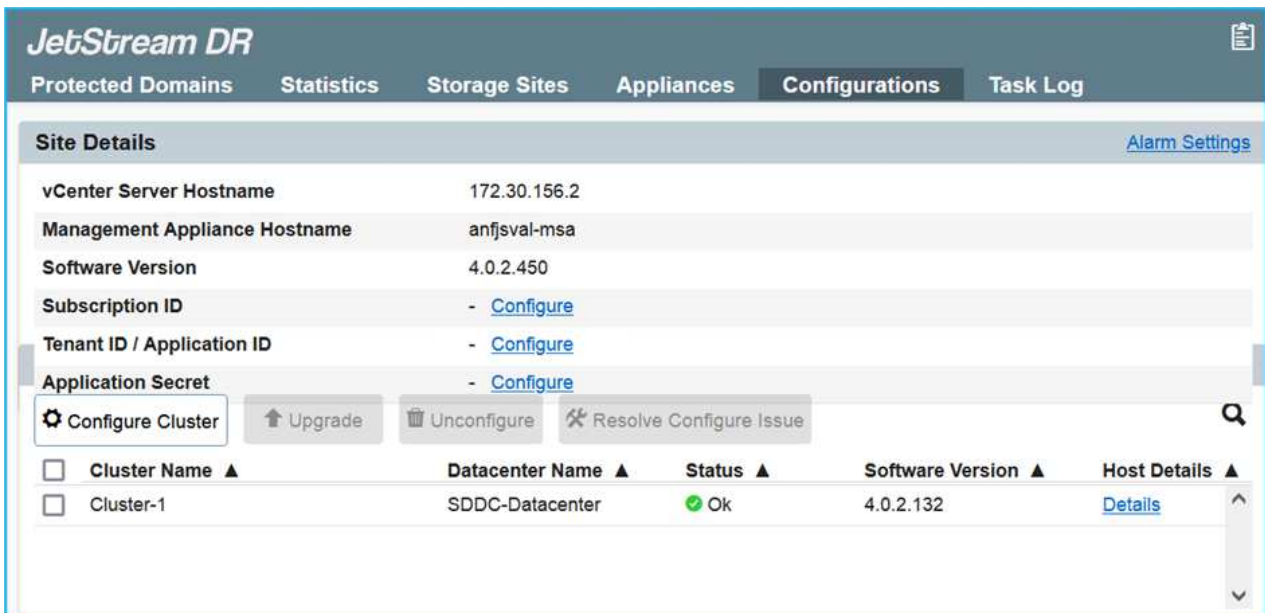


The default CloudAdmin user of the Azure VMware Solution doesn't have sufficient privileges to install JetStream DR for AVS. The Azure VMware Solution enables simplified and automated installation of JetStream DR by invoking the Azure VMware Solution Run command for JetStream DR.

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.

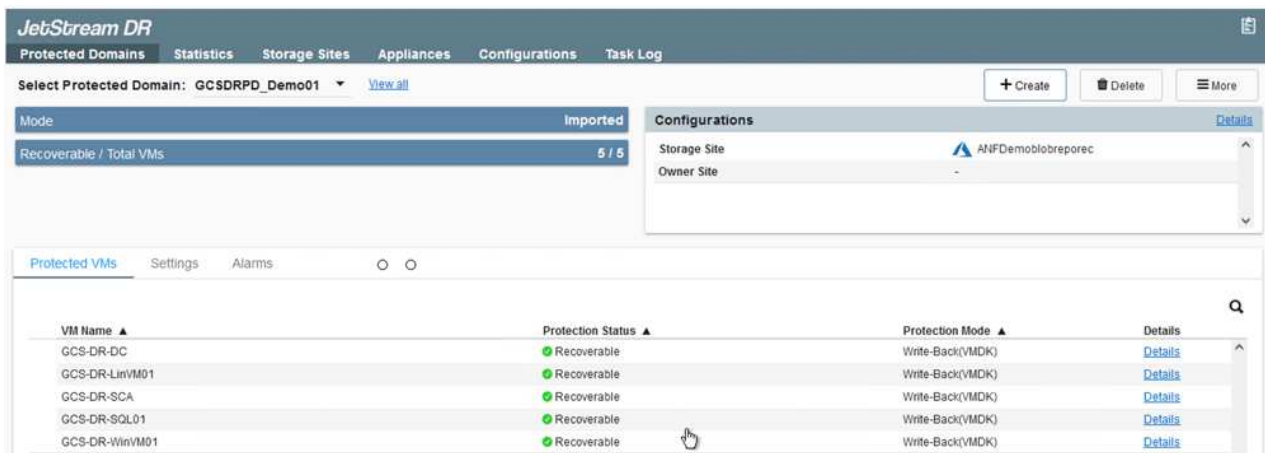


3. From the JetStream DR interface, complete the following tasks:

- a. 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.
- b. 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).



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.

The screenshot displays the Replication management interface. At the top, a summary bar shows 3 Volume Relationships, 4.78 GiB Replicated Capacity, 0 Currently Transferring, 3 Healthy, and 0 Failed. Below this, a table lists three volume relationships, all with a 'snapmirrored' state and 'idle' status. A context menu is open over the first row, with the 'Break' option selected. A 'Break Relationship' dialog box is overlaid on the table, asking for confirmation to break the relationship between 'gcsdrsqldb_sc46' and 'gcsdrsqldb_sc46_copy'. The dialog has 'Break' and 'Cancel' buttons.

Health Status	Source Volume	Target Volume	Total Transfer Time	Status	Mirror State	Last Successful Transfer
✓	gcsdrsqldb_sc46 ntaphci-a300e9u25	gcsdrsqldb_sc46_copy ANFCVODRDemo	6 minutes 41 seconds	idle	snapmirrored	May 5, 2022, 12:08:34 PM 33.66 KiB
✓	gcsdrsqhld_sc46 ntaphci-a300e9u25	gcsdrsqhld_sc46_copy ANFCVODRDemo	4 minutes 56 seconds	idle	snapmirrored	
✓	gcsdrsqlog_sc46 ntaphci-a300e9u25	gcsdrsqlog_sc46_copy ANFCVODRDemo	10 minutes 18 seconds	idle	snapmirrored	



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](#) + Create Failover More

Mode: Continuous Rehydration in Progress

Recoverable / Total VMs: 4 / 4

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

Current Step: Recover VMs' data from Storage Site

Configurations

- Storage Site: ANFDemotobreporec
- Owner Site: REMOTE (172.21.253.160)
- Datacenter \ Cluster: SDDC-Datacenter \ Cluster-1
- Point-in-time Recovery: Disabled

Protected VMs | Settings | Alarms

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

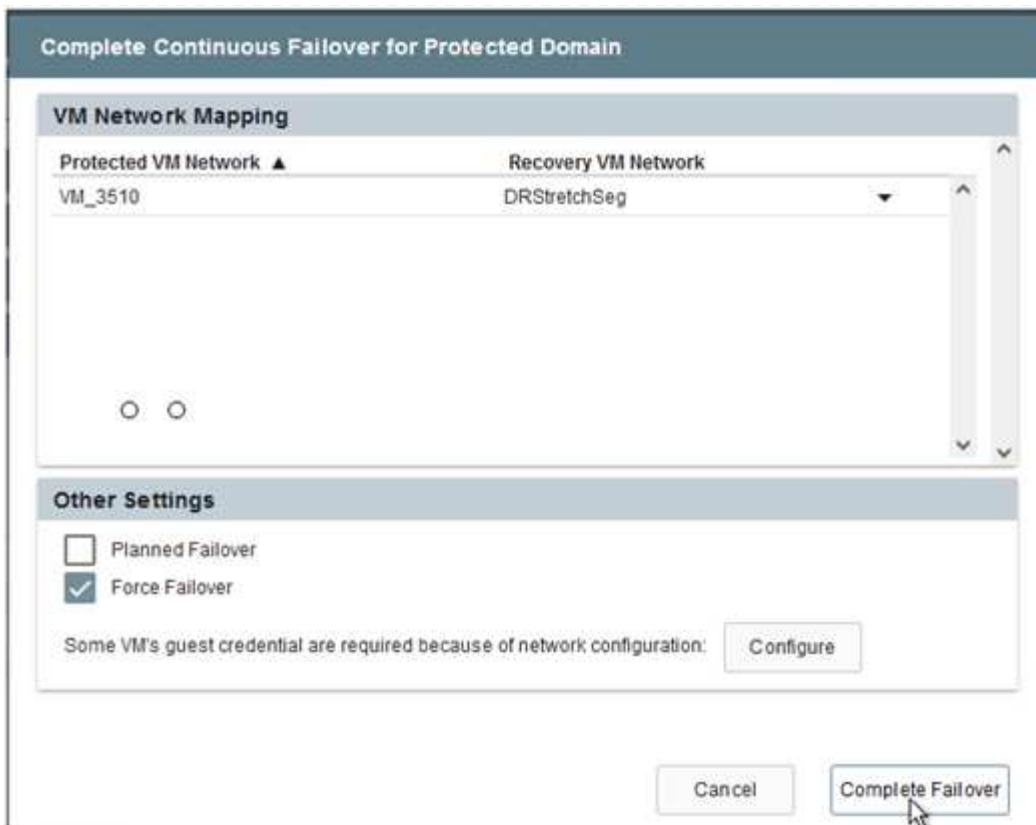
Cancel Complete Failover

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.



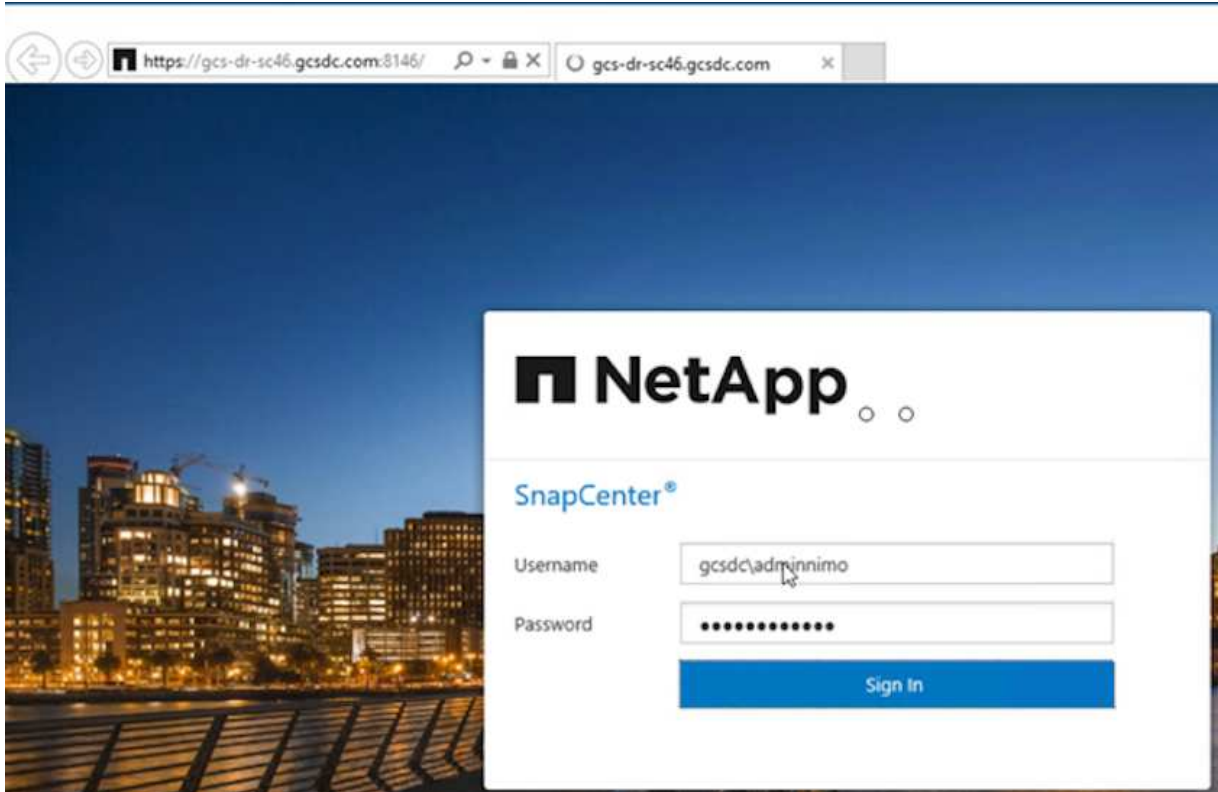
To verify the failover configuration and infrastructure, JetStream DR can be operated in test mode (Test Failover option) to observe the recovery of virtual machines and their data from the object store into a test recovery environment. When a failover procedure is executed in test mode, its operation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



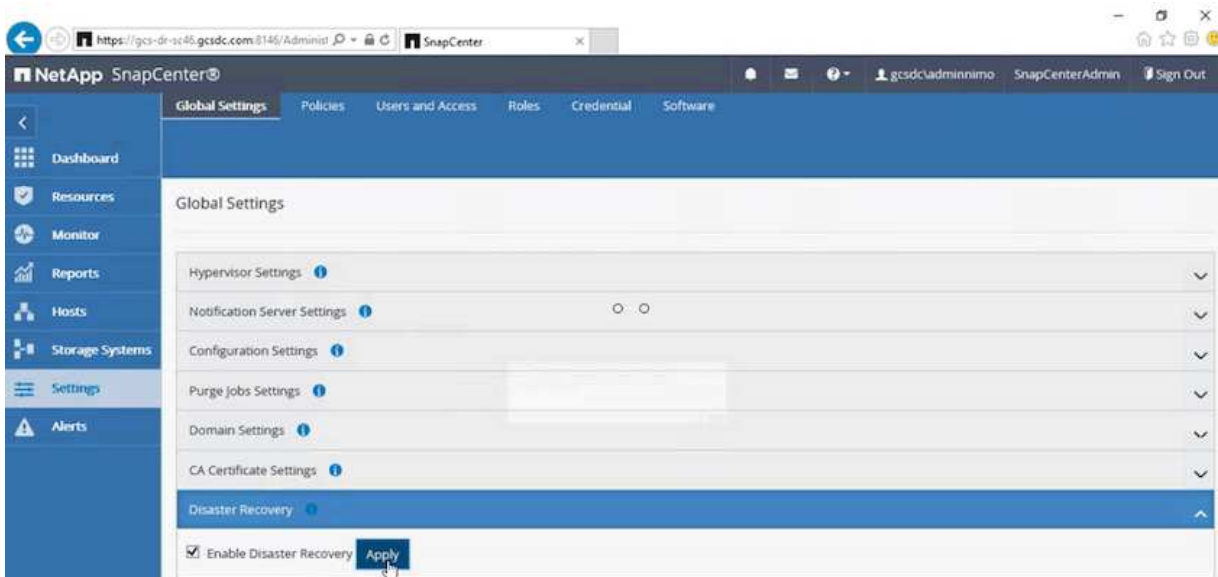
4. After the virtual machines are recovered, use storage disaster recovery for in-guest storage. To

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 browserN.



- 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.


3 Volume Relationships

Health Status	Source Volume	Target Volume	Total Transfer Time	Status	Mirror State	Last Successful Transfer
✓	gcsdrsqldb_sc46 ntaphci-a300e9u25	gcsdrsqldb_sc46_copy ANFCVODRDemo	6 minutes 41 seconds	idle	broken-off	May 5, 2022, 12:08:34 PM 33.66 KiB
✓	gcsdrsqlhd_sc46 ntaphci-a300e9u25	gcsdrsqlhd_sc46_copy ANFCVODRDemo	4 minutes 56 seconds	idle	broken-off	May 5, 2022, 12:09:15 PM 69.84 KiB
✓	gcsdrsqllog_sc46 ntaphci-a300e9u25	gcsdrsqllog_sc46_copy ANFCVODRDemo	10 minutes 18 seconds	idle	broken-off	May 5, 2022, 12:08:34 PM 104.34 KiB


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

Volume	Layout	Type	File System	Status	Capacity	Free Spa...	% Free
---	Simple	Basic		Healthy (R...	450 MB	450 MB	100 %
---	Simple	Basic		Healthy (E...	99 MB	99 MB	100 %
(C:)	Simple	Basic	NTFS	Healthy (B...	89.45 GB	67.03 GB	75 %
BACKUP (G:)	Simple	Basic	NTFS	Healthy (P...	9.97 GB	9.92 GB	99 %
DATA (E:)	Simple	Basic	NTFS	Healthy (P...	24.88 GB	24.57 GB	99 %
LOG (F:)	Simple	Basic	NTFS	Healthy (P...	9.97 GB	8.93 GB	90 %


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




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.



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.



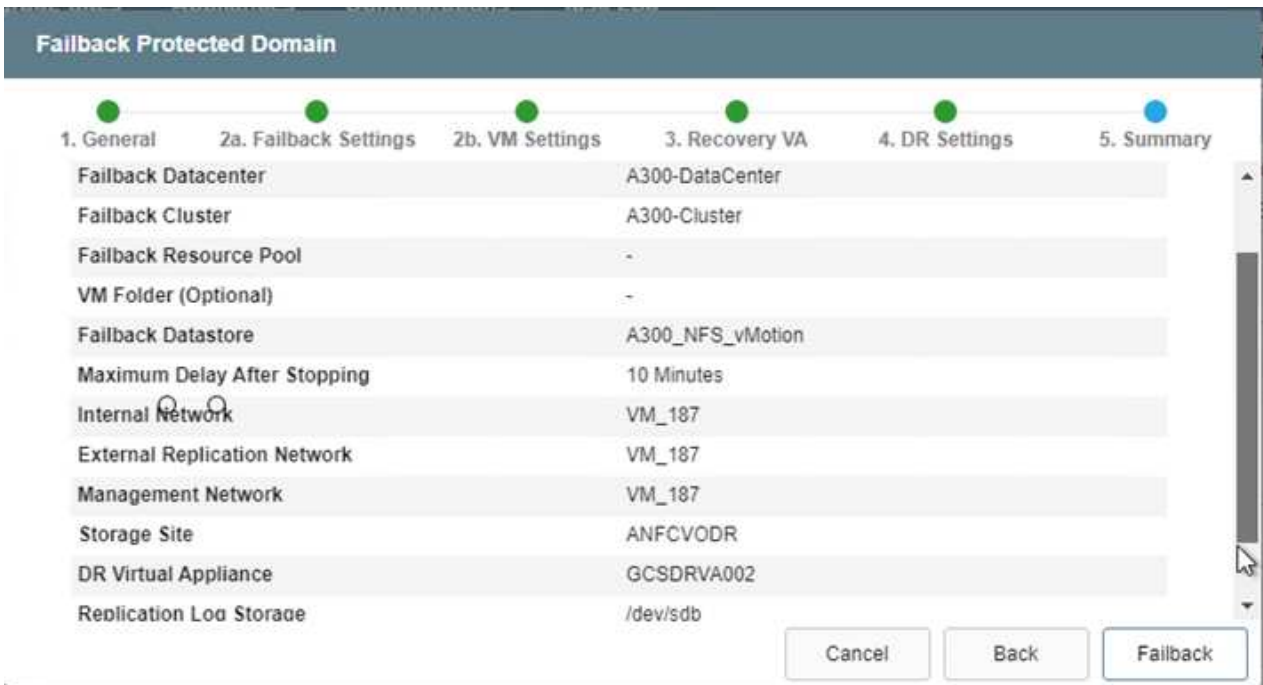
The CPT-generated failback plan can also be used to initiate the return of the VMs and their data from the object store back to the original VMware environment.

The screenshot shows the JetStream DR web interface. At the top, there are navigation tabs: Protected Domains, Statistics, Storage Sites, Appliances, Configurations, and Task Log. Below the tabs, the 'Protected Domains' section is active, showing 'Select Protected Domain: GCSDRPD_Demo01'. A table displays the current state: Mode is 'Running in Failover', Active Site is '172.30.156.2', and Recoverable / Total VMs is '4 / 4'. A 'Configurations' table shows 'Storage Site: ANFCVODR' and 'Owner Site: REMOTE (172.30.156.2)'. A dropdown menu is open over the 'Configurations' table, with the 'Failback' option selected. Below the configurations, there are tabs for 'Protected VMs', 'Settings', and 'Alarms'. The 'Protected VMs' tab is active, showing a table with the following data:

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




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.



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



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
✓	gcsdrsqldb_sc46 ntaphci-a300e9u25	gcsdrsqldb_sc46_copy ANFCVODRDemo	19 seconds	idle	snapmirrored	May 6, 2022, 11:03:08 AM 5.73 MiB
✓	gcsdrsqlhd_sc46_copy ANFCVODRDemo	gcsdrsqlhd_sc46 ntaphci-a300e9u25	1 minute 46 seconds	idle	snapmirrored	May 6, 2022, 11:01:39 AM 800.76 MiB
✓	gcsdrsqllog_sc46 ntaphci-a300e9u25	gcsdrsqllog_sc46_copy ANFCVODRDemo	51 seconds	idle	snapmirrored	May 6, 2022, 11:03:15 AM 785.8 MiB

- Restart the MSSQL server service.
- 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.

TR-4955: Disaster Recovery with Azure NetApp Files (ANF) and Azure VMware Solution (AVS)


Disaster recovery using block-level replication between regions within the cloud is a resilient and cost-effective way of protecting the workloads against site outages and data corruption events (for example, ransomware).

Author(s): Niyaz Mohamed, NetApp Solutions Engineering

Overview

With Azure NetApp files (ANF) cross-region volume replication, VMware workloads running on an Azure VMware Solution (AVS) SDDC site using Azure NetApp files volumes as an NFS datastore on the primary AVS site can be replicated to a designated secondary AVS site in the target recovery region.

Disaster Recovery Orchestrator (DRO) (a scripted solution with a UI) can be used to seamlessly recover workloads replicated from one AVS SDDC to another. DRO automates recovery by breaking replication peering and then mounting the destination volume as a datastore, through VM registration to AVS, to network mappings directly on NSX-T (included with all AVS private clouds).



Prerequisites and general recommendations

- Verify that you have enabled cross-region replication by creating replication peering. See [Create volume replication for Azure NetApp Files](#).
- You must configure ExpressRoute Global Reach between the source and target Azure VMware Solution private clouds.
- You must have a service principal that can access resources.
- The following topology is supported: primary AVS site to secondary AVS site.
- Configure the [replication](#) schedule for each volume appropriately based on business needs and the data-change rate.



Cascading and fan- in and fan- out topologies are not supported.

Getting started

Deploy Azure VMware Solution

The [Azure VMware Solution](#) (AVS) is a hybrid cloud service that provides fully functional VMware SDDCs within a Microsoft Azure public cloud. AVS is a first-party solution fully managed and supported by Microsoft and verified by VMware that uses Azure infrastructure. Therefore, customers get VMware ESXi for compute virtualization, vSAN for hyper-converged storage, and NSX for networking and security, all while taking advantage of Microsoft Azure's global presence, class-leading data-center facilities, and proximity to the rich ecosystem of native Azure services and solutions. A combination of Azure VMware Solution SDDC and Azure NetApp Files provides the best performance with minimal network latency.

To configure an AVS private cloud on Azure, follow the steps in this [link](#) for NetApp documentation and in this [link](#) for Microsoft documentation. A pilot- light environment set up with a minimal configuration can be used for DR purposes. This setup only contains core components to support critical applications, and it can scale out

and spawn more hosts to take the bulk of the load if a failover occurs.




In the initial release, DRO supports an existing AVS SDDC cluster. On-demand SDDC creation will be available in an upcoming release.

Provision and configure Azure NetApp Files

[Azure NetApp Files](#) is a high-performance, enterprise-class, metered file- storage service. Follow the steps in this [link](#) to provision and configure Azure NetApp Files as a NFS datastore to optimize AVS private cloud deployments.

Create volume replication for Azure NetApp Files-powered datastore volumes

The first step is to set up cross- region replication for the desired datastore volumes from the AVS primary site to the AVS secondary site with the appropriate frequencies and retentions.



Follow the steps in this [link](#) to set up cross-region replication by creating replication peering. The service level for the destination capacity pool can match that of the source capacity pool. However, for this specific use case, you can select the standard service level and then [modify the service level](#) in the event of a real disaster or DR simulations.



A cross- region replication relationship is a prerequisite and must be created beforehand.

DRO installation

To get started with DRO, use the Ubuntu operating system on the designated Azure virtual machine and make sure you meet the prerequisites. Then install the package.

Prerequisites:

- Service principal that can access resources.
- Make sure that appropriate connectivity exists to the source and destination SDDC and Azure NetApp Files instances.
- DNS resolution should be in place if you are using DNS names. Otherwise, use IP addresses for vCenter.

OS requirements:

- Ubuntu Focal 20.04 (LTS)The following packages must be installed on the designated agent virtual machine:
- Docker
- Docker- compose
- JqChange `docker.sock` to this new permission: `sudo chmod 666 /var/run/docker.sock.`



The `deploy.sh` script executes all required prerequisites.

The steps are as follows:

1. Download the installation package on the designated virtual machine:

```
git clone https://github.com/NetApp/DRO-Azure.git
```



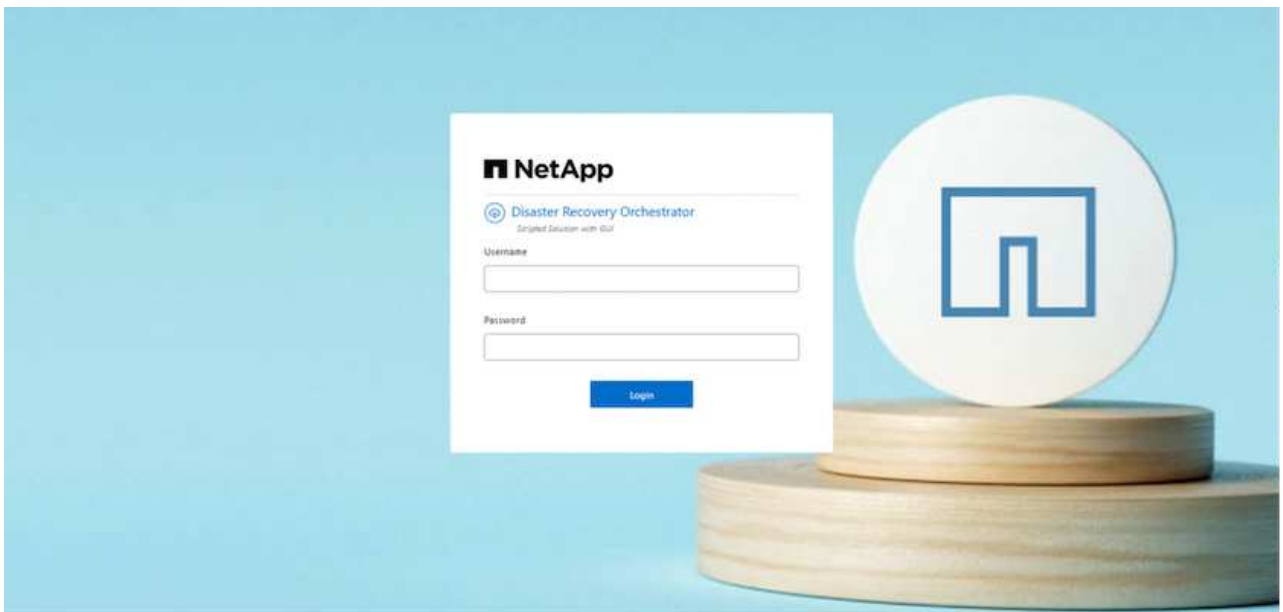
The agent must be installed in the secondary AVS site region or in the primary AVS site region in a separate AZ than the SDDC.

2. Unzip the package, run the deployment script, and enter the host IP (for example, 10.10.10.10).

```
tar xvf draas_package.tar  
Navigate to the directory and run the deploy script as below:  
sudo sh deploy.sh
```

3. Access the UI using the following credentials:

- Username: admin
- Password: admin



DRO configuration

After Azure NetApp Files and AVS have been configured properly, you can begin configuring DRO to automate the recovery of workloads from the primary AVS site to the secondary AVS site. NetApp recommends deploying the DRO agent in the secondary AVS site and configuring the ExpressRoute gateway connection so that the DRO agent can communicate via the network with the appropriate AVS and Azure NetApp Files components.

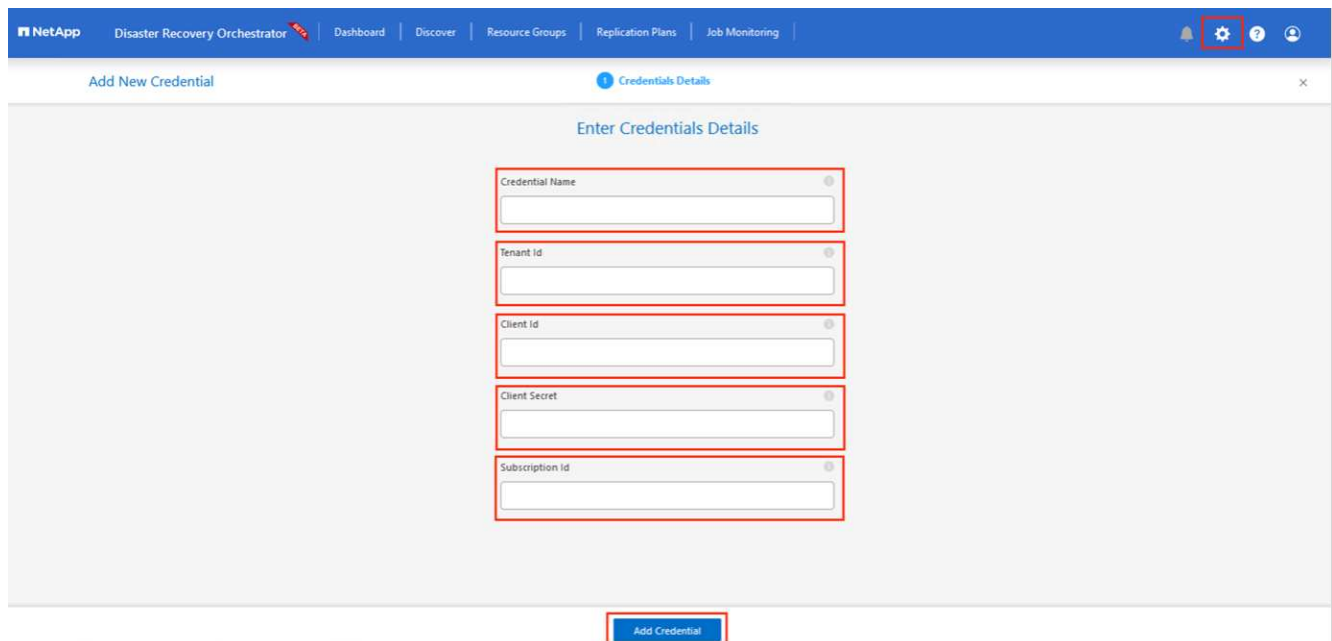
The first step is to Add credentials. DRO requires permission to discover Azure NetApp Files and the Azure VMware Solution. You can grant the required permissions to an Azure account by creating and setting up an Azure Active Directory (AD) application and by obtaining the Azure credentials that DRO needs. You must bind the service principal to your Azure subscription and assign it a custom role that has the relevant required permissions. When you add source and destination environments, you are prompted to select the credentials associated with the service principal. You need to add these credentials to DRO before you can click Add New Site.

To perform this operation, complete the following steps:

1. Open DRO in a supported browser and use the default username and password (admin/admin). The password can be reset after the first login using the Change Password option.
2. In the upper right of the DRO console, click the **Settings** icon, and select **Credentials**.
3. Click Add New Credential and follow the steps in the wizard.
4. To define the credentials, enter information about the Azure Active Directory service principal that grants the required permissions:
 - Credential name
 - Tenant ID
 - Client ID
 - Client secret
 - Subscription ID

You should have captured this information when you created the AD application.

5. Confirm the details about the new credentials and click Add Credential.

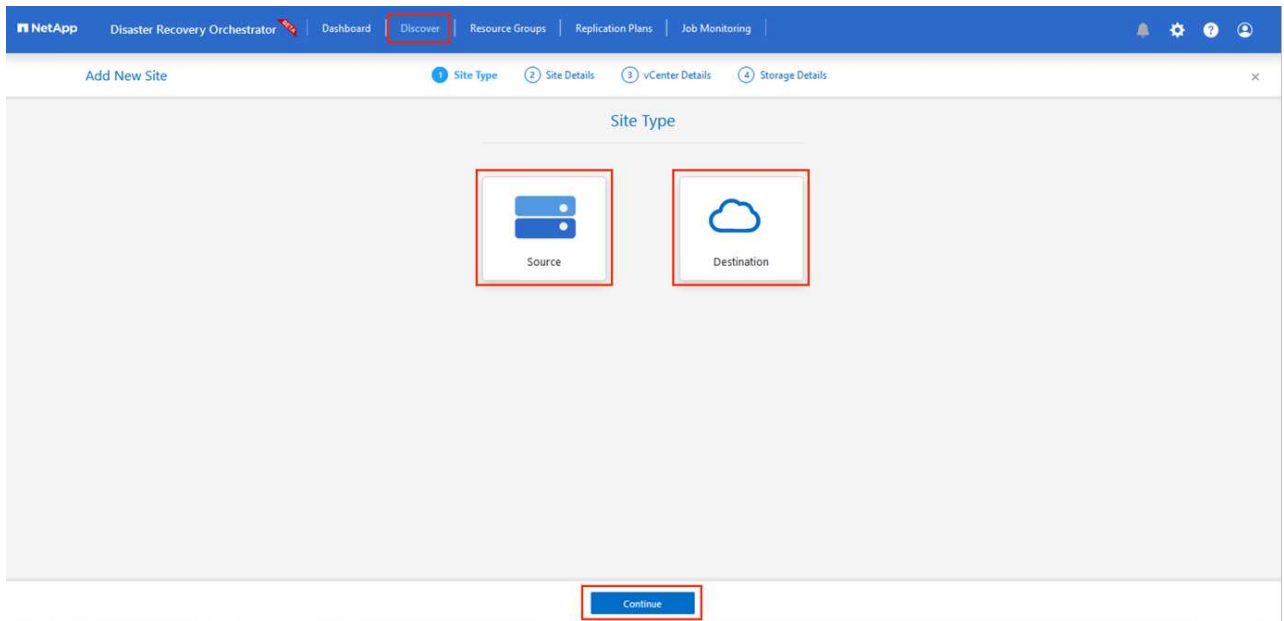


The screenshot shows the NetApp Disaster Recovery Orchestrator (DRO) console. The top navigation bar includes the NetApp logo, the product name 'Disaster Recovery Orchestrator', and several menu items: Dashboard, Discover, Resource Groups, Replication Plans, and Job Monitoring. On the right side of the navigation bar, there are icons for a notification bell, a settings gear (highlighted with a red box), a help question mark, and a user profile. Below the navigation bar, the main content area is titled 'Add New Credential' and 'Credentials Details'. The 'Enter Credentials Details' section contains five input fields, each with a red box around it: 'Credential Name', 'Tenant Id', 'Client Id', 'Client Secret', and 'Subscription Id'. At the bottom of the form, there is a blue 'Add Credential' button, also highlighted with a red box.

After you add the credentials, it's time to discover and add the primary and secondary AVS sites (both vCenter and the Azure NetApp files storage account) to DRO. To add the source and destination site, complete the following steps:

6. Go to the **Discover** tab.

7. Click **Add New Site**.
8. Add the following primary AVS site (designated as **Source** in the console).
 - SDDC vCenter
 - Azure NetApp Files storage account
9. Add the following secondary AVS site (designated as **Destination** in the console).
 - SDDC vCenter
 - Azure NetApp Files storage account

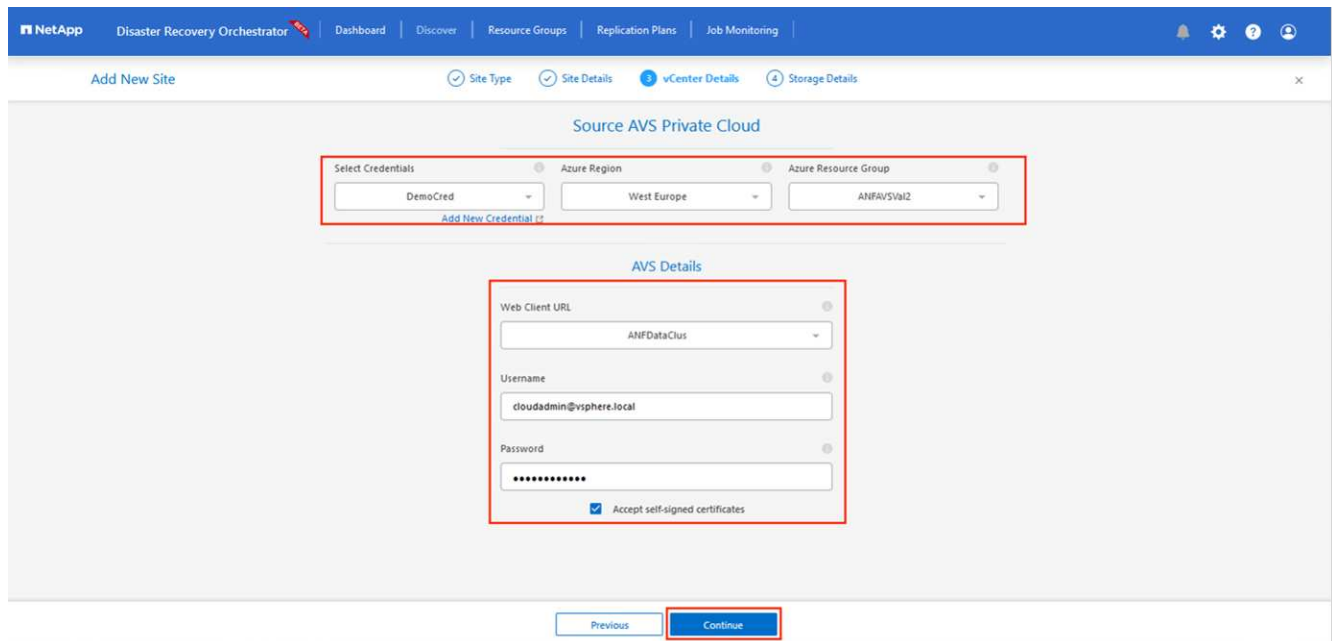


10. Add site details by clicking **Source**, entering a friendly site name, and select the connector. Then click **Continue**.




For demonstration purposes, adding a source site is covered in this document.


11. Update the vCenter details. To do this, select the credentials, Azure region, and resource group from the dropdown for the primary AVS SDDC.
12. DRO lists all the available SDDCs within the region. Select the designated private cloud URL from the dropdown.
13. Enter the `cloudadmin@vsphere.local` user credentials. This can be accessed from Azure Portal. Follow the steps mentioned in this [link](#). Once done, click **Continue**.



14. Select the Source Storage details (ANF) by selecting the Azure Resource group and NetApp account.
15. Click **Create Site**.



Once added, DRO performs automatic discovery and displays the VMs that have corresponding cross-region replicas from the source site to the destination site. DRO automatically detects the networks and segments used by the VMs and populates them.




The next step is to group the required VMs into their functional groups as resource groups.

Resource groupings

After the platforms have been added, group the VMs you want to recover into resource groups. DRO resource groups allow you to group a set of dependent VMs into logical groups that contain their boot orders, boot delays, and optional application validations that can be executed upon recovery.


To start creating resource groups, click the **Create New Resource Group** menu item.

1. Access **Resource Groups** and click **Create New Resource Group**.




2. Under New Resource Group, select the source site from the dropdown and click **Create**.
3. Provide the resource group details and click **Continue**.
4. Select appropriate VMs using the search option.
5. Select the **Boot Order** and **Boot Delay** (secs) for all the selected VMs. Set the order of the power-on sequence by selecting each virtual machine and setting up the priority for it. The default value for all virtual machines is 3. The options are as follows:
 - The first virtual machine to power on
 - Default

- The last virtual machine to power on



6. Click **Create Resource Group**.




Replication plans


You must have a plan to recover applications in the event of a disaster. Select the source and destination vCenter platforms from the drop down, pick the resource groups to be included in this plan, and also include the grouping of how applications should be restored and powered on (for example, domain controllers, tier-1, tier-2, and so on). Plans are often called blueprints as well. To define the recovery plan, navigate to the Replication Plan tab, and click **New Replication Plan**.

To start creating a replication plan, complete the following steps:


1. Navigate to **Replication Plans** and click **Create New Replication Plan**.



2. On the **New Replication Plan**, provide a name for the plan and add recovery mappings by selecting the Source Site, associated vCenter, Destination Site, and associated vCenter.



3. After recovery mapping is complete, select the **Cluster Mapping**.



4. Select **Resource Group Details** and click **Continue**.


5. Set the execution order for the resource group. This option enables you to select the sequence of operations when multiple resource groups exist.

6. Once done, set network mapping to the appropriate segment. The segments should already be provisioned on the secondary AVS cluster, and, to map the VMs to those, select the appropriate segment.

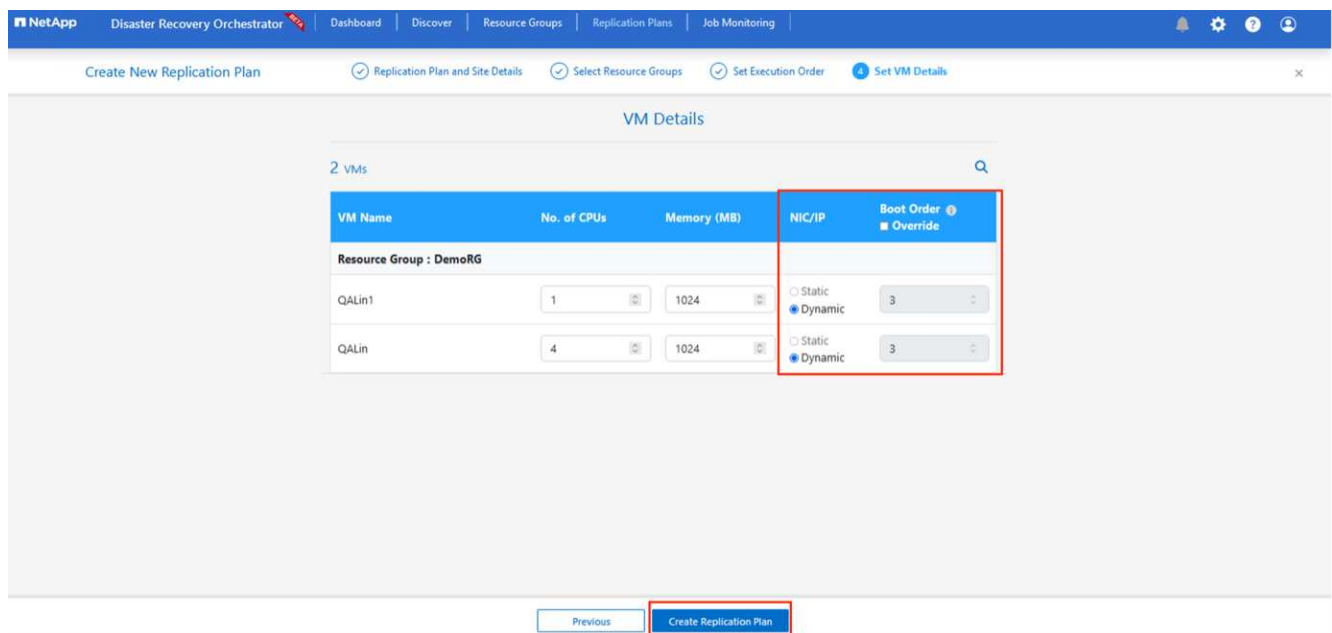
7. Datastore mappings are automatically selected based on the selection of VMs.




Cross- region replication (CRR) is at the volume level. Therefore, all VMs residing on the respective volume are replicated to the CRR destination. Make sure to select all VMs that are part of the datastore, because only virtual machines that are part of the replication plan are processed.




8. Under VM details, you can optionally resize the VMs CPU and RAM parameters. This can be very helpful when you are recovering large environments to smaller target clusters or when you are conducting DR tests without having to provision a one-to-one physical VMware infrastructure. Also, modify the boot order and boot delay (secs) for all the selected VMs across the resource groups. There is an additional option to modify the boot order if any changes are required from what you selected during resource- group boot-order selection. By default, the boot order selected during resource- group selection is used, however any modifications can be performed at this stage.



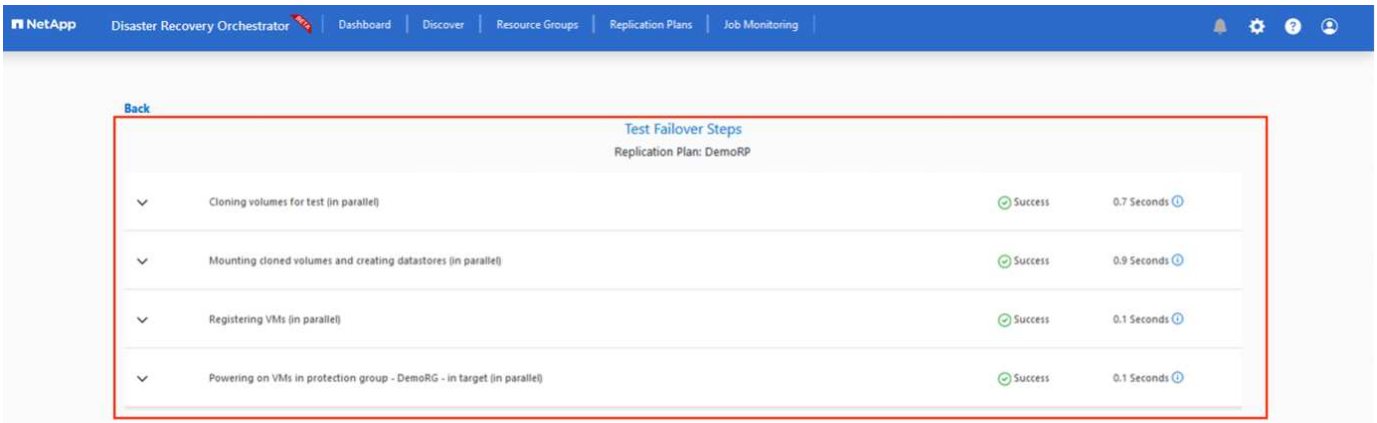
9. Click **Create Replication Plan**. After the replication plan is created, you can exercise the failover, test failover, or migrate options depending on your requirements.




During the failover and test failover options, the most recent snapshot is used, or a specific snapshot can be selected from a point-in-time snapshot. The point-in-time option can be very beneficial if you are facing a corruption event like ransomware, where the most recent replicas are already compromised or encrypted. DRO shows all available time points.



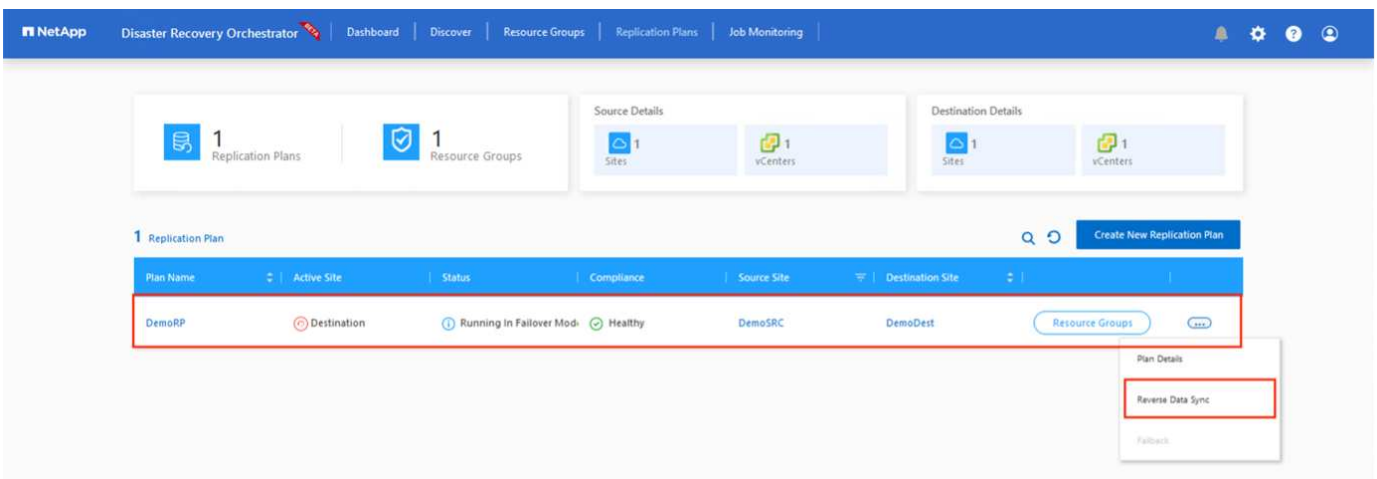
To trigger failover or test failover with the configuration specified in the replication plan, you can click **Failover** or **Test Failover**. You can monitor the replication plan in the task menu.





After failover is triggered, the recovered items can be seen in the secondary site AVS SDDC vCenter (VMs, networks, and datastores). By default, the VMs are recovered to Workload folder.



Failback can be triggered at the replication plan level. In case of test failover, the tear down option can be used to roll back the changes and remove the newly created volume. Failbacks related to failover are a two-step process. Select the replication plan and select **Reverse Data sync**.



After this step is complete, trigger failback to move back to the primary AVS site.



From the Azure portal, we can see that the replication health has been broken off for the appropriate volumes that were mapped to the secondary site AVS SDDC as read/write volumes. During test failover, DRO does not map the destination or replica volume. Instead, it creates a new volume of the required cross-region replication snapshot and exposes the volume as a datastore, which consumes additional physical capacity from the capacity pool and ensures that the source volume is not modified. Notably, replication jobs can continue during DR tests or triage workflows. Additionally, this process makes sure that the recovery can be cleaned up without the risk of the replica being destroyed if errors occur or corrupted data is recovered.

Ransomware recovery

Recovering from ransomware can be a daunting task. Specifically, it can be difficult for IT organizations to pinpoint what the safe point of return is, and, once that's determined, how to ensure that recovered workloads are safeguarded from the attacks reoccurring (for example, from sleeping malware or through vulnerable applications).

DRO addresses these concerns by allowing organizations to recover from any available point-in-time. Workloads are then recovered to functional and yet isolated networks, so that applications can function and communicate with each other but are not exposed to any north-south traffic. This process gives security teams a safe place to conduct forensics and identify any hidden or sleeping malware.

Conclusion

The Azure NetApp Files and Azure VMware disaster recovery solution provide you with the following benefits:

- Leverage efficient and resilient Azure NetApp Files cross- region replication.
- Recover to any available point-in-time with snapshot retention.
- Fully automate all required steps to recover hundreds to thousands of VMs from the storage, compute, network, and application validation steps.
- Workload recovery leverages the “Create new volumes from the most recent snapshots” process, which doesn’t manipulate the replicated volume.
- Avoid any risk of data corruption on the volumes or snapshots.
- Avoid replication interruptions during DR test workflows.
- Leverage DR data and cloud compute resources 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 allowing recovery to smaller compute clusters.

Where to find additional information

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

- Create volume replication for Azure NetApp Files

<https://learn.microsoft.com/en-us/azure/azure-netapp-files/cross-region-replication-create-peering>

- Cross-region replication of Azure NetApp Files volumes

<https://learn.microsoft.com/en-us/azure/azure-netapp-files/cross-region-replication-introduction#service-level-objectives>

- Azure VMware Solution

<https://learn.microsoft.com/en-us/azure/azure-vmware/introduction>

- Deploy and configure the Virtualization Environment on Azure

[Setup AVS on Azure](#)

- Deploy and configure Azure VMware Solution

<https://learn.microsoft.com/en-us/azure/azure-vmware/deploy-azure-vmware-solution?tabs=azure-portal>

Using Veeam Replication and Azure NetApp Files datastore for disaster recovery to Azure VMware Solution

Azure NetApp Files (ANF) datastores decouples storage from compute and unlocks the flexibility needed for any organisation to take their workloads to the cloud. It provides customers with flexible, high-performance storage infrastructure that scales independently of compute resources. Azure NetApp Files datastore’s simplifies and

optimizes the deployment alongside Azure VMware Solution (AVS) as a disaster recovery site for on premises VMWare environments.

Author: Niyaz Mohamed - NetApp Solutions Engineering

Overview

Azure NetApp Files (ANF) volume based NFS datastores can be used to replicate data from on-premises using any validated third-party solution that provides VM replication capability. By adding Azure NetApp Files datastores, it will enable cost optimised deployment vs building an Azure VMware Solution SDDC with enormous amount of ESXi hosts to accommodate the storage. This approach is called a “Pilot Light Cluster”. A pilot light cluster is a minimal AVS host configuration (3 x AVS nodes) along with Azure NetApp Files Datastore capacity.

The objective is to maintain a low-cost infrastructure with all the core components to handle a failover. A pilot light cluster can scale out and provision more AVS hosts if a failover does occur. And once the failover is complete and normal operations are restored, the pilot light cluster can scale back down to low-cost mode of operations.


Purposes of this document

This article describes how to use Azure NetApp Files datastore with Veeam Backup and replication to set up disaster recovery for on-premises VMware VMs to (AVS) using the Veeam VM replication software functionality.

Veeam Backup & Replication is a backup and replication application for virtual environments. When virtual machines are replicated, Veeam Backup & Replication is replicated from on AVS, the software will create an exact copy of the VMs in the native VMware vSphere format on the target AVS SDDC cluster. Veeam Backup & Replication will keep the copy synchronized with the original VM. Replication provides the best recovery time objective (RTO) as there is a mounted copy of a VM at the DR site in a ready-to-start state.

This replication mechanism ensures that the workloads can quickly start in a AVS SDDC in the case of a disaster event. The Veeam Backup & Replication software also optimizes traffic transmission for replication over WAN and slow connections. In addition, it also filters out duplicate data blocks, zero data blocks, swap files, and “excluded VM guest OS files”. The software will also compress the replica traffic. To prevent replication jobs from consuming the entire network bandwidth, WAN accelerators and network throttling rules can be utilized.

The replication process in Veeam Backup & Replication is job driven which means replication is performed by configuring replication jobs. In the case of a disaster event, failover can be triggered to recover the VMs by failing over to its replica copy. When failover is performed, a replicated VM takes over the role of the original VM. Failover can be performed to the latest state of a replica or to any of its good known restore points. This enables ransomware recovery or isolated testing as needed. Veeam Backup & Replication offers multiple options to handle different disaster recovery scenarios.



Solution Deployment

High level steps

1. Veeam Backup and Replication software is running in an on-premises environment with appropriate network connectivity.
2. [Deploy Azure VMware Solution \(AVS\)](#) private cloud and [attach Azure NetApp Files datastores](#) to Azure VMware Solution hosts.

A pilot-light environment set up with a minimal configuration can be used for DR purposes. VMs will fail over to this cluster in the event of an incident, and additional nodes can be added).

3. Set up replication job to create VM replicas using Veeam Backup and Replication.
4. Create failover plan and perform failover.
5. Switch back to production VMs once the disaster event is complete and primary site is Up.

Pre-requisites for Veeam VM Replication to AVS and ANF datastores

1. Ensure the Veeam Backup & Replication backup VM is connected to the source as well as the target AVS SDDC clusters.
2. The backup server must be able to resolve short names and connect to source and target vCenters.
3. The target Azure NetApp Files datastore must have enough free space to store VMDKs of replicated VMs.

For additional information, refer to "Considerations and Limitations" covered [here](#).

Deployment Details

Step 1: Replicate VMs

Veeam Backup & Replication leverages VMware vSphere snapshot capabilities/During replication, Veeam Backup & Replication requests VMware vSphere to create a VM snapshot. The VM snapshot is the point-in-time copy of a VM that includes virtual disks, system state, configuration and metadata. Veeam Backup & Replication uses the snapshot as a source of data for replication.

To replicate VMs, follow the below steps:

1. Open the Veeam Backup & Replication Console.
2. On the Home view. Right click the jobs node and select Replication Job > Virtual machine.
3. Specify a job name and select the appropriate advanced control checkbox. Click Next.
 - Select the Replica seeding check box if connectivity between on-premises and Azure has restricted bandwidth.
*Select the Network remapping (for AVS SDDC sites with different networks) check box if segments on Azure VMware Solution SDDC do not match that of on-premises site networks.
 - If the IP addressing scheme in on-premises production site differs from the scheme in the target AVS site, select the Replica re-IP (for DR sites with different IP addressing scheme) check box.

Name
Specify the name and description for this job, and provide information on your DR site.

Name:
AVS_20230522_RepJob01

Description:
Created by VEEAMBKPSRV05\Administrator at 5/21/2023 10:52 PM.

Show advanced controls:

- Replica seeding (for low bandwidth DR sites)
- Network remapping (for DR sites with different virtual networks)
- Replica re-IP (for DR sites with different IP addressing scheme)

High priority
Backup infrastructure resources are offered to high priority jobs first. Use this option for jobs sensitive to the start time, or jobs with strict RPO requirements.

< Previous **Next >** Finish Cancel

4. Select the VMs to be replicated to Azure NetApp Files datastore attached to a Azure VMware Solution SDDC in the **Virtual Machines*** step. The Virtual machines can be placed on vSAN to fill the available vSAN datastore capacity. In a pilot light cluster, the usable capacity of a 3-node cluster will be limited. The rest of the data can be easily placed on Azure NetApp Files datastores so that the VMs can recovered, and cluster can be expanded to meet the CPU/mem requirements. Click **Add**, then in the **Add Object** window select the necessary VMs or VM containers and click **Add**. Click **Next**.

Virtual Machines
Select one or more VMs to replicate. Use exclusion settings to exclude specific VMs and virtual disks from replication.

Virtual machines to replicate:

Name	Type	Size
TestVeeam21	Virtual Machine	873 MB
TestVeeam22	Virtual Machine	890 MB
TestVeeam23	Virtual Machine	883 MB
TestVeeam24	Virtual Machine	879 MB
TestVeeam25	Virtual Machine	885 MB
TestVeeam26	Virtual Machine	883 MB
TestVeeam27	Virtual Machine	879 MB
TestVeeam28	Virtual Machine	880 MB
TestVeeam29	Virtual Machine	878 MB
TestVeeam30	Virtual Machine	876 MB
TestVeeam31	Virtual Machine	888 MB
TestVeeam32	Virtual Machine	881 MB
TestVeeam33	Virtual Machine	877 MB
TestVeeam34	Virtual Machine	875 MB
TestVeeam35	Virtual Machine	882 MB
WinSQL401	Virtual Machine	20.3 GB
WinSQL405	Virtual Machine	24.2 GB

Total size: 120 GB

< Previous **Next >** Finish Cancel

- After that, select the destination as Azure VMware Solution SDDC cluster / host and the appropriate resource pool, VM folder and FSx for ONTAP datastore for VM replicas. Then click **Next**.

Edit Replication Job [AVS_20230522_RepJob01]

Destination
Specify where replicas should be created in the DR site.

Host or cluster:
Cluster-1 Choose...

Resource pool:
Resources Choose...
[Pick resource pool](#) for selected replicas

VM folder:
vm Choose...
[Pick VM folder](#) for selected replicas

Datastore:
ds001 [152.6 GB free] ds001 is an ANF Datastore Choose...
[Pick datastore](#) for selected virtual disks

< Previous **Next >** Finish Cancel

6. In the next step, create the mapping between source and destination virtual network as needed.

The screenshot shows a software interface for configuring a network. At the top left, there is a green arrow icon pointing right, followed by the heading "Network" and the instruction "Select how virtual networks map to each other between production and DR sites." Below this is a vertical sidebar with menu items: "Name", "Virtual Machines", "Destination", "Network" (which is highlighted), "Job Settings", "Data Transfer", "Guest Processing", "Schedule", and "Summary". The main area is titled "Network mapping:" and contains a table with two columns: "Source network" and "Target network". The table has two rows of data: the first row shows "VM_3508 (vDS-Switch0)" under Source and "SepSeg" under Target; the second row shows "VM_3510 (vDS-Switch0)" under Source and "SegmentTemp" under Target. To the right of the table are three buttons: "Add...", "Edit...", and "Remove". At the bottom of the interface are four buttons: "< Previous", "Next >" (highlighted with a blue border), "Finish", and "Cancel".

Source network	Target network
VM_3508 (vDS-Switch0)	SepSeg
VM_3510 (vDS-Switch0)	SegmentTemp

7. In the **Job Settings** step, specify the backup repository that will store metadata for VM replicas, retention policy and so on.

8. Update the **Source** and **Target** proxy servers in the **Data Transfer** step and leave **Automatic** selection (default) and keep **Direct** option selected and click **Next**.

9. At the **Guest Processing** step, select **Enable application-aware processing** option as needed. Click **Next**.

Guest Processing
Choose guest OS processing options available for running VMs.

Name **Enable application-aware processing**
Detects and prepares applications for consistent backup, performs transaction logs processing, and configures the OS to perform required application restore steps upon first boot.

Customize application handling options for individual machines and applications [Applications...](#)

Virtual Machines

Destination

Network
Guest interaction proxy: Automatic selection [Choose...](#)

Job Settings
Guest OS credentials: [Add...](#)

Data Transfer

Guest Processing [Manage accounts](#)

Customize guest OS credentials for individual machines and operating systems [Credentials...](#)

Schedule
Verify network connectivity and credentials for each machine included in the job [Test Now](#)

Summary

< Previous **Next >** Finish Cancel

10. Choose the replication schedule to run the replication job to run on a regular basis.

Schedule
Specify the job scheduling options. If you do not set the schedule, the job will need to be controlled manually.

Name **Run the job automatically**

Daily at this time: 10:00 PM [Days...](#) Everyday

Monthly at this time: 10:00 PM [Months...](#) Fourth Saturday

Periodically every: 1 [Schedule...](#) Hours

After this job: Replication Job 2 (Created by VEEAMBKPSRV05\Administrator at 6/6/)

Virtual Machines

Destination

Network

Job Settings

Data Transfer

Guest Processing

Schedule

Summary

Automatic retry

Retry failed items processing: 3 times

Wait before each retry attempt for: 10 minutes

Backup window


Terminate job if it exceeds allowed backup window [Window...](#)

If the job does not complete within allocated backup window, it will be terminated to prevent snapshot commit during production hours.


< Previous **Apply** Finish Cancel

11. At the **Summary** step of the wizard, review details of the replication job. To start the job right after the

wizard is closed, select the **Run the job when I click Finish** check box, otherwise leave the check box unselected. Then click **Finish** to close the wizard.



Once the replication job starts, the VMs with the suffix specified will be populated on the destination AVS SDDC cluster / host.



For additional information for Veeam replication, refer [How Replication Works](#)

Step 2: Create a failover plan

When the initial replication or seeding is complete, create the failover plan. Failover plan helps in performing failover for dependent VMs one by one or as a group automatically. Failover plan is the blueprint for the order in which the VMs are processed including the boot delays. The failover plan also helps to ensure that critical dependant VMs are already running.

To create the plan, navigate to the new sub section called **Replicas** and select **Failover Plan**. Choose the appropriate VMs. Veeam Backup & Replication will look for the closest restore points to this point in time and use them to start VM replicas.



The failover plan can only be added once the initial replication is complete and the VM replicas are in Ready state.



The maximum number of VMs that can be started simultaneously when running a failover plan is 10



During the failover process, the source VMs will not be powered off

To create the **Failover Plan**, do the following:

1. On the Home view. Right click the Replicas node and select Failover Plans > Failover Plan > VMware vSphere.

The screenshot shows the Veeam Backup & Replication software interface. At the top, there are several icons for different actions: Backup Job, Replication Job, CDP Policy, Backup Copy, Copy Job, Restore, Failover Plan (highlighted with a red box), Import Backup, and Export Backup. Below this is a search bar and a table of jobs. The table has columns for Name, Type, Objects, Status, and Last Run. The table contains two rows of data.

Name	Type	Objects	Status	Last Run
Replication Job 2	VMware Replication	1	Stopped	63 days ago
AVS_20230522_RepJob01	VMware Replication	20	Stopped	8 days ago

2. Next provide a name and a description to the plan. Pre and Post-failover script can be added as required. For instance, run a script to shutdown VMs before starting the replicated VMs.

Edit Failover Plan [ANF_AVS_FP01] X

General
Type in name and description for this failover plan, and optionally specify scripts to trigger before and after the failover.

General
Virtual Machines
Summary

Name: ANF_AVS_FP01

Description: Created by VEEAMBKPSRV05\Administrator at 5/24/2023 9:08 AM.

Pre-failover script:
 Browse...

Post-failover script:
 Browse...

< Previous Next > Finish Cancel

3. Add the VMs to the plan and modify the VM boot order and boot delays to meet the application dependencies.



Virtual Machines

Add virtual machines to be failed over as a part of this plan. Use VM order and delays to ensure all application dependencies are met.

General

Virtual Machines

Summary

Virtual machines:

Name	Delay	Replica state
TestVeeam21	2 sec	63 days ago (5:52 AM T...
TestVeeam23	2 sec	7 days ago (10:12 AM T...
TestVeeam24	2 sec	7 days ago (10:20 AM T...
TestVeeam22	2 sec	7 days ago (10:10 AM T...
WinSQL401	2 sec	7 days ago (3:52 AM Tu...
WinSQL405	2 sec	8 days ago (4:05 PM Mo...
TestVeeam25	2 sec	7 days ago (10:14 AM T...
TestVeeam26	2 sec	7 days ago (10:17 AM T...
TestVeeam27	2 sec	7 days ago (10:18 AM T...
TestVeeam28	2 sec	7 days ago (10:14 AM T...
TestVeeam29	2 sec	7 days ago (10:18 AM T...
TestVeeam30	2 sec	7 days ago (10:15 AM T...
TestVeeam31	2 sec	7 days ago (10:21 AM T...
TestVeeam32	2 sec	7 days ago (10:13 AM T...
TestVeeam33	2 sec	7 days ago (10:15 AM T...
TestVeeam34	2 sec	7 days ago (10:14 AM T...
TestVeeam35	2 sec	7 days ago (10:20 AM T...

Add VM

Remove

Set Delay...

↑ Up

↓ Down

< Previous

Apply

Finish

Cancel

For additional information for creating replication jobs, refer [Creating Replication Jobs](#).

Step 3: Run the failover plan

During failover, the source VM in the production site is switched over to its replica at the disaster recovery site. As part of the failover process, Veeam Backup & Replication restores the VM replica to the required restore point and moves all I/O activities from the source VM to its replica. Replicas can be used not only in case of a disaster, but also to simulate DR drills. During failover simulation, the source VM remains running. Once all the necessary tests have been conducted, you can undo the failover and return to normal operations.




Make sure network segmentation is in place to avoid IP conflicts during failover.

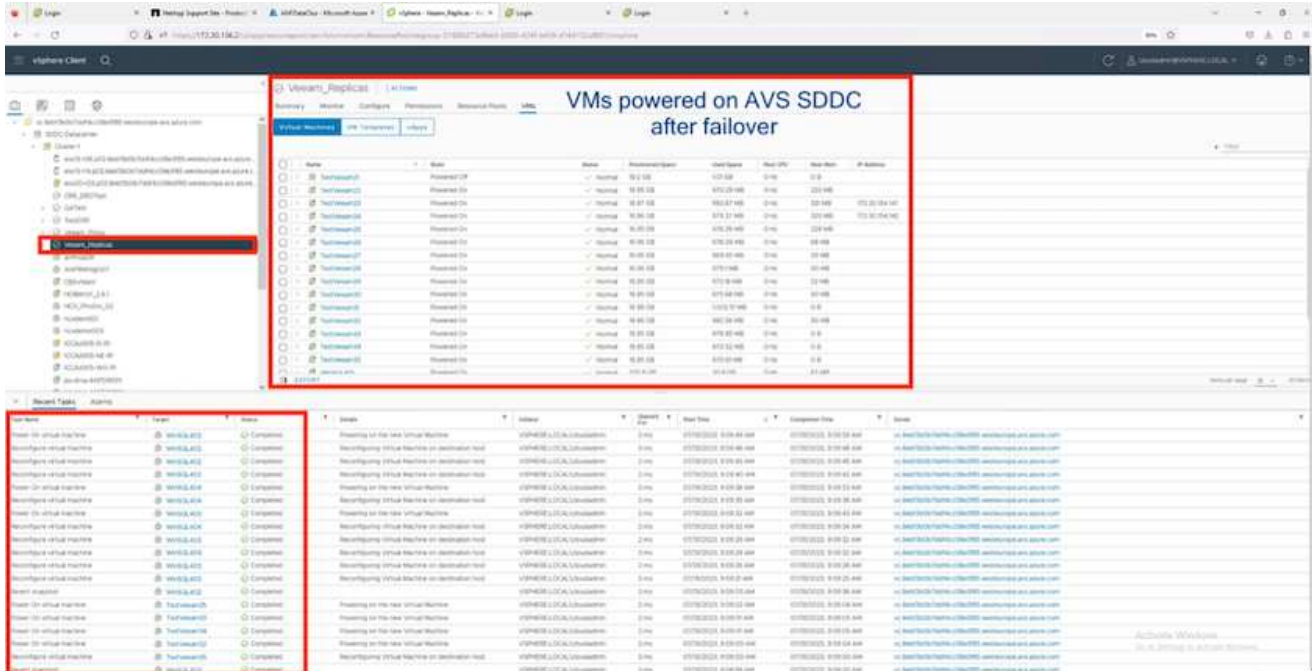
To start the failover plan, simply click in **Failover Plans** tab and right click on your failover plan. Select ***Start**. This will failover using the latest restore points of VM replicas. To fail over to specific restore points of VM replicas, select **Start to**.

The screenshot displays the Veeam Backup & Replication console interface. At the top, there is a toolbar with icons for Start, Start to..., Retry, Undo, Statistics, Edit, and Delete. Below the toolbar, the left sidebar shows a navigation tree with 'Failover Plans' selected. The main pane shows a table of failover plans. The table has columns for Name, Platform, Status, and Number of VMs. One plan, 'ANF_AVS_FP01', is listed with a 'VMware' platform, 'Completed' status, and '20' VMs. A context menu is open over this plan, showing options: Start, Start to..., Undo, Statistics, Delete, and Edit... The 'Start' and 'Start to...' options are highlighted with red boxes.

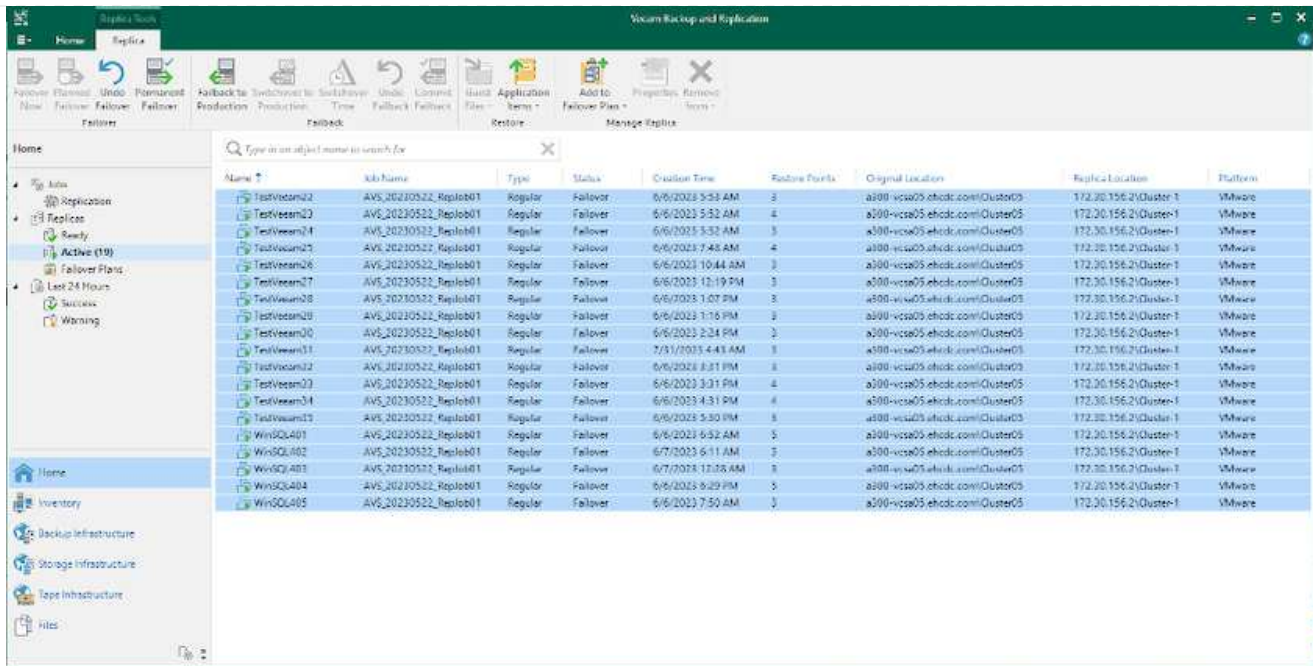
Name	Platform	Status	Number of VMs
ANF_AVS_FP01	VMware	Completed	20



The state of the VM replica changes from Ready to Failover and VMs will start on the destination Azure VMware Solution (AVS) SDDC cluster / host.



Once the failover is complete, the status of the VMs will change to "Failover".



Veeam Backup & Replication stops all replication activities for the source VM until its replica is returned to the Ready state.

For detailed information about failover plans, refer [Failover Plans](#).

Step 4: Failback to the Production site

When the failover plan is running, it is considered as an intermediate step and needs to be finalized based on the requirement. The options include the following:


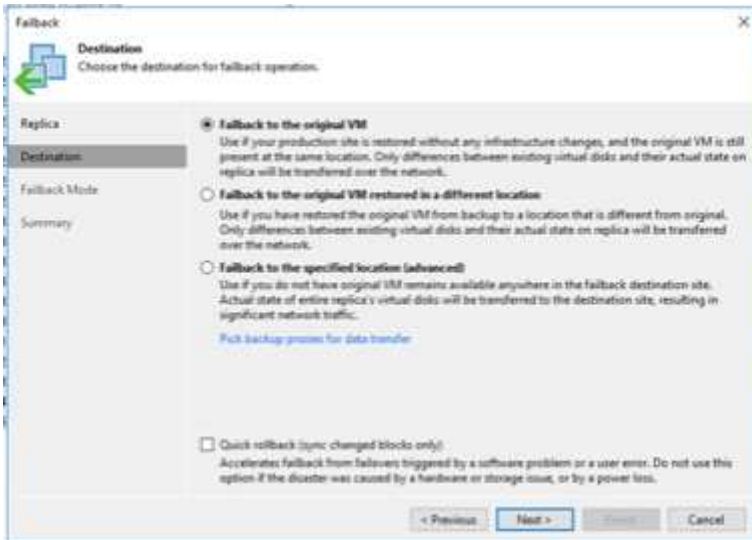
- **Failback to production** - switch back to the original VM and transfer all changes that took place while the VM replica was running to the original VM.





When you perform failback, changes are only transferred but not published. Choose **Commit failback** (once the original VM is confirmed to work as expected) or Undo failback to get back to the VM replica If the original VM is not working as expected.

- **Undo failover** - switch back to the original VM and discard all changes made to the VM replica while it was running.
- **Permanent Failover** - permanently switch from the original VM to a VM replica and use this replica as the original VM.


In this demo, Failback to production VM was chosen. Failback to the original VM was selected during the Destination step of the wizard and “Power on VM after restoring” check box was enabled.






Failback commit is one of the ways to finalize failback operation. When failback is committed, it confirms that the changes sent to the VM which is failed back (the production VM) are working as expected. After the commit operation, Veeam Backup & Replication resumes replication activities for the production VM.

For detailed information about the failback process, refer Veeam documentation for [Failover and Failback for replication](#).



After failback to production is successful, the VMs are all restored back to the original production site.



Conclusion

Azure NetApp Files datastore capability enables Veeam or any validated third-party tool to provide a low-cost DR solution by leveraging Pilot light clusters instead of standing up a large cluster only to accommodate VM replicas. This provides an efficacious way to handle a tailored, customized disaster recovery plan and to reuse existing backup products in house for DR, enabling cloud-based disaster recovery by exiting on-premises DR datacenters. It is possible to failover by clicking a button in case of disaster or to failover automatically if a

disaster occurs.

To learn more about this process, feel free to follow the detailed walkthrough video.

<https://netapp.hosted.panopto.com/Panopto/Pages/Embed.aspx?id=2855e0d5-97e7-430f-944a-b061015e9278>

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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