



Protecting Workloads on GCP / GCVE

NetApp Solutions

NetApp
October 30, 2024

This PDF was generated from <https://docs.netapp.com/us-en/netapp-solutions/ehc/gcp-app-dr-sc-cvs-veeam.html> on October 30, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Protecting Workloads on GCP / GCVE 1
 - Application Consistent Disaster Recovery with NetApp SnapCenter and Veeam Replication 1
 - Application Disaster Recovery with SnapCenter, Cloud Volumes ONTAP and Veeam Replication 4
 - Using Veeam Replication and Google Cloud NetApp Volumes datastore for disaster recovery to Google Cloud VMware Engine 8

Protecting Workloads on GCP / GCVE

Application Consistent Disaster Recovery with NetApp SnapCenter and Veeam Replication

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 Google Cloud.

Authors: Suresh Thoppay, NetApp


Overview

Many customers are looking for an effective disaster recovery solution for their application VMs hosted on VMware vSphere. Many of them use their existing backup solution to perform recovery during disaster. Many times that solution increases the RTO and doesn't meet their expectations. To reduce the RPO and RTO, Veeam VM replication can be utilized even from on-prem to GCVE as long as network connectivity and environment with appropriate permissions are available.

NOTE: Veeam VM Replication doesn't protect VM guest connected storage devices like iSCSI or NFS mounts inside the guest VM. Need to protect those separately.

For application consistent replication for SQL VM and to reduce the RTO, we used SnapCenter to orchestrate snapmirror operations of SQL database and log volumes.

This document provides a step-by-step approach for setting up and performing disaster recovery that uses NetApp SnapMirror, Veeam, and the Google Cloud VMware Engine (GCVE).



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 Google Cloud network, use the connectivity options like dedicated interconnect or Cloud VPN. Segments should be created based on the on-premises VLAN design.



There are multiple options for connecting on-premises datacenters to Google Cloud, which prevents us from outlining a specific workflow in this document. Refer to the Google Cloud documentation for the appropriate on-premises-to-Google 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 BlueXP 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 Veeam software and start replicating virtual machines to Google Cloud VMware Engine instance.
4. During a disaster event, break the SnapMirror relationship using BlueXP and trigger failover of virtual machines with Veeam.
 - a. Reconnect the iSCSI LUNs and NFS mounts for the application VMs.
 - b. Bring up applications online.
5. Invoke failback to the protected site by reverse resyncing SnapMirror after the primary site has been recovered.

Deployment Details

Configure CVO on Google Cloud and replicate volumes to CVO

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



For sample step-by-step instructions on setting up SnapCenter and replicating the data, Refer to [Setup Replication with SnapCenter](#)

[Review of SQL VM protection with SnapCenter](#)

Configure GCVE hosts and CVO data access

Two important factors to consider when deploying the SDDC are the size of the SDDC cluster in the GCVE 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.

NetApp Cloud Volume Service for NFS Datastore and Cloud Volumes ONTAP for SQL databases and log can be deployed to any VPC and GCVE should have private connection to that VPC to mount NFS datastore and have VM connect to iSCSI LUNs.

To configure GCVE SDDC, see [Deploy and configure the Virtualization Environment on Google Cloud Platform \(GCP\)](#). As a prerequisite, verify that the guest VMs residing on the GCVE hosts are able to consume data from Cloud Volumes ONTAP after connectivity has been established.

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

Install Veeam Components

Based on deployment scenario, the Veeam backup server, backup repository and backup proxy that needs to be deployed. For this use case, there is no need to deploy object store for Veeam and Scale-out repository also not required.

[Refer to the Veeam documentation for the installation procedure](#)

For additional information, please refer [Migration with Veeam Replication](#)

Setup VM Replication with Veeam

Both on-premises vCenter and GCVE vCenter needs to be registered with Veeam. [Setup vSphere VM Replication Job](#) At the Guest Processing step of wizard, select disable application processing as we will be utilizing SnapCenter for application aware backup and recovery.

<https://netapp.hosted.panopto.com/Panopto/Pages/Embed.aspx?id=8b7e4a9b-7de1-4d48-a8e2-b01200f00692>

Failover of Microsoft SQL Server VM

<https://netapp.hosted.panopto.com/Panopto/Pages/Embed.aspx?id=9762dc99-081b-41a2-ac68-b01200f00ac0>

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.
- Veeam Replication allows changing VM IP addresses on DR site.

Application Disaster Recovery with SnapCenter, Cloud Volumes ONTAP and Veeam Replication


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 Google Cloud.

Authors: Suresh Thoppay, NetApp

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 Veeam 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, Veeam, and the Google Cloud VMware Engine (GCVE).



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 Google Cloud network, use the connectivity options like dedicated interconnect or Cloud VPN. Segments should be created based on the on-premises VLAN design.



There are multiple options for connecting on-premises datacenters to Google Cloud, which prevents us from outlining a specific workflow in this document. Refer to the Google Cloud documentation for the appropriate on-premises-to-Google 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 Veeam software and start replicating virtual machines to Google Cloud VMware Engine instance.
4. During a disaster event, break the SnapMirror relationship using Cloud Manager and trigger failover of virtual machines with Veeam.
 - a. Reconnect the iSCSI LUNs and NFS mounts for the application VMs.
 - b. Bring up applications online.
5. Invoke failback to the protected site by reverse resyncing SnapMirror after the primary site has been recovered.

Deployment Details

Configure CVO on Google Cloud and replicate volumes to CVO

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



For sample step-by-step instructions on setting up SnapCenter and replicating the data, Refer to [Setup Replication with SnapCenter](#)

[Setup Replication with SnapCenter](#)

Configure GCVE hosts and CVO data access

Two important factors to consider when deploying the SDDC are the size of the SDDC cluster in the GCVE 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.

Cloud Volumes ONTAP can be deployed to any VPC and GCVE should have private connection to that VPC to have VM connect to iSCSI LUNs.

To configure GCVE SDDC, see [Deploy and configure the Virtualization Environment on Google Cloud Platform \(GCP\)](#). As a prerequisite, verify that the guest VMs residing on the GCVE hosts are able to consume data from Cloud Volumes ONTAP after connectivity has been established.

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

Install Veeam Components

Based on deployment scenario, the Veeam backup server, backup repository and backup proxy that needs to be deployed. For this use case, there is no need to deploy object store for Veeam and Scale-out repository also not required.

[Refer to the Veeam documentation for the installation procedure](#)

Setup VM Replication with Veeam

Both on-premises vCenter and GCVE vCenter needs to be registered with Veeam. [Setup vSphere VM Replication Job](#) At the Guest Processing step of wizard, select disable application processing as we will be utilizing SnapCenter for application aware backup and recovery.

[Setup vSphere VM Replication Job](#)

Failover of Microsoft SQL Server VM

[Failover of Microsoft SQL Server VM](#)

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.
- Veeam Replication allows changing VM IP addresses on DR site.

Using Veeam Replication and Google Cloud NetApp Volumes datastore for disaster recovery to Google Cloud VMware Engine

A comprehensive disaster recovery plan is critical for businesses in times of crisis. Many organizations leverage cloud computing for daily operations and disaster recovery. This proactive approach can reduce or eliminate expensive business disruptions.

This article describes how to use Veeam Backup & Replication to set up disaster recovery for on-premises VMware VMs to Google Cloud VMware Engine (GCVE) with Google Cloud NetApp Volumes (NetApp Volumes).

Overview

Google Cloud NetApp Volumes is a storage service from Google and NetApp that is available for Google Cloud. NetApp Volumes service provides high performance NFS/SMB storage. VMware certified NetApp Volumes NFS storage can be used as an external datastore for ESXi hosts in GCVE. Users are required to make a peering connection between their GCVE private cloud and NetApp Volumes project. There are no network charges resulting from storage access within a region. Users can create NetApp Volumes volumes in the Google Cloud console and enable deletion protection before mounting volumes as datastores to their ESXi hosts.

NetApp Volumes 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 NetApp Volumes datastores, it enables cost optimized deployment instead of building a Google Cloud VMware Engine (GCVE) based SDDC with a large number of ESXi hosts to accommodate the storage. This approach is called a “Pilot Light Cluster”. A pilot light cluster is a minimal GCVE host configuration (3 x GCVE ESXi hosts) along with NetApp Volumes datastores capacity to allow for independent scaling to meet capacity requirements.

The objective is to sustain a cost-effective infrastructure with just the core components to manage a failover. A pilot light cluster can expand and add more GCVE hosts in the event of a failover. Once the failover is resolved and normal operations resume, the pilot light cluster can reduce its scale, returning to a low-cost operational mode.

Purposes of this document

This article describes how to use a Google Cloud NetApp Volumes datastore with Veeam Backup & Replication to set up disaster recovery for on-premises VMware VMs to GCVE 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 will create an exact copy of the VMs in the native VMware vSphere format on the target GCVE 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 GCVE 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 known good restore points. This enables ransomware recovery or isolated testing as needed. Veeam Backup & Replication offers multiple options to handle different disaster recovery scenarios.

Solution Overview

This solution covers the following high level steps:

1. Create an NFS volume using Google Cloud NetApp Volumes
2. Follow GCP process to create a GCVE datastore from the NetApp Volumes NFS volume.
3. Set up a replication job to create VM replicas using Veeam Backup & Replication.
4. Create a failover plan and perform failover.
5. Switch back to production VMs once the disaster event is complete and the primary site is up.

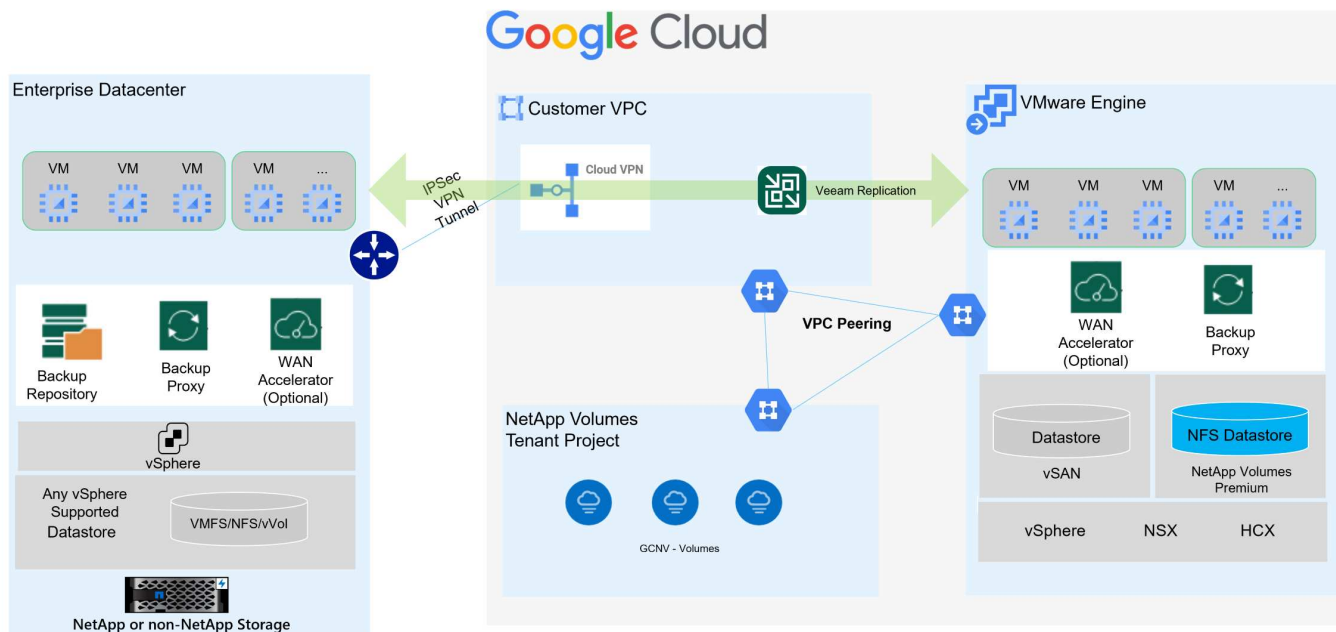


When creating a volume in NetApp Volumes, for use as a GCVE datastore, only NFS v3 is supported.

For more information on using NetApp Volumes NFS volumes as datastores for GCVE, check out [Using NFS volume as vSphere datastore hosted by Google Cloud NetApp Volumes](#).

Architecture

The following diagram shows the architecture of the solution presented in this documentation. A recommended best practice is to have a Veeam Backup & Replication server located at both the on-premises site and in the GCVE SDDC. Backup and recovery is performed and managed by the Veeam server on-premises, and replication is managed by the Veeam server in the GCVE SDDC. This architecture provides the highest availability when a failure occurs in the primary datacenter.



Pre-requisites for Veeam Replication to GCVE and NetApp Volumes datastores

This solution requires the following components and configurations:

1. NetApp Volumes has a Storage Pool available with enough free capacity to accommodate the NFS volume to be created.
2. Veeam Backup and Replication software is running in an on-premises environment with appropriate network connectivity.
3. Ensure the Veeam Backup & Replication backup VM is connected to the source as well as the target GCVE SDDC clusters.
4. Ensure the Veeam Backup & Replication backup VM is connected to the Veeam Proxy server VMs at both the source and target GCVE clusters.
5. The backup server must be able to resolve short names and connect to source and target vCenters.

Users are required to make a peering connection between their GCVE private cloud and NetApp Volumes project using the VPC Network peering or Private connections pages within the VMware Engine Cloud console UI.



Veeam requires a GCVE solution user account with elevated privileges when adding the GCVE vCenter server to the Veeam Backup and Replication inventory. For more information refer to the Google Cloud Platform (GCP) documentation, [Elevating VMware Engine privileges](#).

For additional information refer to [Considerations and Limitations](#) in the Veeam Backup & Replication documentation.

Deployment steps

The following sections outline the deployment steps to create and mount an NFS datastore using Google Cloud NetApp Volumes, and use Veeam Backup and Replication to implement a full disaster recovery solution between an on-premises datacenter and Google Cloud VMware Engine.

Create NetApp Volumes NFS volume and datastore for GCVE

Refer to [Using NFS volume as vSphere datastore hosted by Google Cloud NetApp Volumes](#) for an overview of how to Google Cloud NetApp Volumes as a datastore for GCVE.

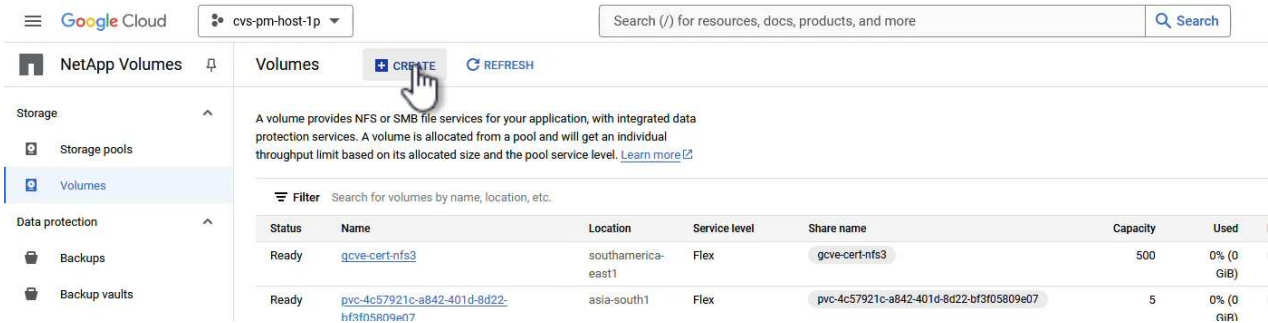
Complete the following steps to create and use an NFS datastore for GCVE using NetApp Volumes:

Create NetApp Volumes NFS volume

Google Cloud NetApp Volumes is accessed from the Google Cloud Platform (GCP) console.

Refer to [Create a volume](#) in the Google Cloud NetApp Volumes documentation for detailed information on this step.

1. In a web browser, navigate to <https://console.cloud.google.com/> and log into your GCP console. Search for **NetApp Volumes** to get started.
2. In the **NetApp Volumes** management interface, click on **Create** to get started creating an NFS volume.



3. In the **Create a volume** wizard, fill out all required information:
 - A name for the volume.
 - The Storage Pool on which to create the volume.
 - A share name used when mounting the NFS volume.
 - The capacity of the volume in GiB.
 - The storage protocol to be used.
 - Check the box to **Block volume from deletion when clients are connected** (required by GCVE when mounting as a datastore).
 - The export rules for accessing the volume. This is the IP addresses of the ESXi adapters on the NFS network.
 - A snapshot schedule used to protect the volume using local snapshots.
 - Optionally, choose to backup the volume and/or create labels for the volume.



When creating a volume in NetApp Volumes, for use as a GCVE datastore, only NFS v3 is supported.

Google Cloud

cvr-pn-host-1p

Search (/) for resources, docs, prod...

NetApp Volumes

Create a volume

Storage

Storage pools

Volumes

Data protection

Backups

Backup vaults

Policies

Active Directory policies

CMEIX policies

Backup policies

A volume provides NFS or SMB file services for your application with integrated data protection services. A volume is allocated from a storage pool and gets an individual or shared throughput limit based on its allocated capacity and storage pool service level.

[Learn more](#)

Volume name *
gcnv-dt-plan

Choose a permanent. Must be unique to the region. Use lowercase letters, numbers, hyphens and underscores. Start with a letter.

Description

Storage pool details

Select a storage pool in which to create the volume.

[SELECT STORAGE POOL](#) [CREATE NEW STORAGE POOL](#)

Volume details

Share name *
Must be unique to a location

Capacity *
GB

Capacity must be between 100 GB and 102,400 GB. Increments of 1 GB.

Protocol(s) *
NFSv3

Configuration for selected protocol(s)

☐ Block volume from deletion when clients are connected.
Required for volumes used as GCVE datastores. Choice is permanent.

Export rules

Snapshot configuration

[CREATE](#) [CANCEL](#)

Select a storage pool

Storage pools

Name	Location	Available capacity	Service level	VPC	Active Directory	LDAP enabled	Encry
<input checked="" type="radio"/> asize1-gve	asia-southeast1	1548 GiB	Premium	shared-vpc-prod		No	
<input type="radio"/> asize1-gve-extreme	asia-southeast1	0 GiB	Extreme	shared-vpc-prod	asia-southeast1.ad	No	
<input type="radio"/> gve-data-pool	asia-south1	1014 GiB	Flex	shared-vpc-prod		No	
<input type="radio"/> gve-cent-noraml	southamerica-east1	524 GiB	Flex	shared-vpc-prod		No	
<input type="radio"/> montreal-premium	northamerica-northeast1	1148 GiB	Premium	shared-vpc-prod	montreal.ad	No	
<input type="radio"/> ok-at-pool	northamerica-northeast1	998 GiB	Premium	shared-vpc-prod	montreal.ad	No	
<input type="radio"/> rnvind-db-perfext	asia-south1-a	1536 GiB	Flex	shared-vpc-prod		No	
<input type="radio"/> rnvind-sb21	asia-southeast1	1948 GiB	Standard	shared-vpc-prod		No	
<input type="radio"/> rnvind-sb22	australia-southeast1	1748 GiB	Standard	shared-vpc-prod		No	gcp.italy
<input type="radio"/> rnvind-vertxual	asia-south1	769 GiB	Flex	shared-vpc-prod		No	
<input type="radio"/> vp-1p-sa-s1-gve-ds22	southamerica-east1-a	0 GiB	Flex	shared-vpc-prod		No	
<input type="radio"/> test	me-west1-b	1024 GiB	Flex	shared-vpc-prod		No	
<input type="radio"/> yashwan-pool1	northamerica-northeast1	1792 GiB	Premium	shared-vpc-prod	montreal.ad	No	

[SELECT](#) [CANCEL](#)

Google Cloud

cvs-pm-host-1p

Search (/) for resources, docs, or images

NetApp Volumes

Storage pools

Volumes

Backups

Backup vaults

Active Directory policies

CMEK policies

Backup policies

Create a volume

SELECT STORAGE POOL

CREATE NEW STORAGE POOL

Volume details

Share name *
gcnv-dr-plan
Must be unique to a location

Capacity *
1000
GiB
Capacity must be between 100 GiB and 102,400 GiB. Increments of 1 GiB

Protocol(s) *
NFSv3

Configuration for selected protocol(s)

☒ Block volume from deletion when clients are connected
Required for volumes used as GCVE datastores. Choice is permanent.

Export rules

Rules are evaluated in order. First matching rule applies.

Rules

New Rule

Allowed Clients *
192.168.100.15,192.168.100.16,192.168.100.18
Comma-separated list of IPv4 addresses or CIDRs (up to 4096 characters).

Access *
☒ Read & Write
☐ Read Only

Root Access (no_root_squash)
☒ On
☐ Off

CREATE CANCEL

Click on **Create** to finish creating the volume.

4. Once the volume is created, the NFS export path required to mount the volume can be viewed from the volume's properties page.

Google Cloud

cvs-pm-host-1p

Search (/) for resources, docs, products,

NetApp Volumes

gcnv-dr-plan

EDIT

REVERT

MOUNT INSTRUCTIONS

DELETE

Storage

Storage pools

Volumes

Data protection

Backups

Backup vaults

Policies

Active Directory policies

CMEK policies

Backup policies

Resource type

Volume

State

Ready

State details

Available for use

Description

OVERVIEW

SNAPSHOTS

BACKUPS

REPLICATION

A volume provides NFS or SMB file services for your application with integrated data protection services. A volume is allocated from a storage pool and gets an individual or shared throughput limit based on its allocated capacity and storage pool service level.

Share name

NFS export path

Used to mount this file share on a linux client VM. Run the mount command with the following remote target on the VM's local directory.

```
$ 10.165.128.100:/gcnv-dr-plan
```

Name	gcnv-dr-plan
Capacity	1000 GiB
Used	0% (0 GiB)
Protocol(s)	NFSV3
Storage pool	asiase1-gcve
Location	asia-southeast1
Service level	Premium
VPC	shared-vpc-prod
Active directory policy	No value
LDAP enabled	No
Encryption	Google-managed
Block volume from deletion when clients are connected	Yes
Make snapshot directory visible	No
Allow scheduled backups	No

Mount the NFS datastore in GCVE

At the time of this writing the process to mount a datastore in GCVE requires opening a GCP support ticket to have the volume mounted as an NFS datastore.

Refer to [Using NFS volume as vSphere datastore hosted by Google Cloud NetApp Volumes](#) for more information.


Replicate VMs to GCVE and execute failover plan, and failback

Replicate VMs to NFS datastore in GCVE

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, complete the following steps:

1. Open the Veeam Backup & Replication Console.
2. On the **Home** tab, click on **Replication Job > Virtual machine...**



3. On the **Name** page of the **New Replication Job** wizard, specify a job name and select the appropriate advanced control checkboxes.
 - Select the Replica seeding check box if connectivity between on-premises and GCP has restricted bandwidth.
 - Select the Network remapping (for GCVE SDDC sites with different networks) check box if segments on the GCVE SDDC do not match that of the on-premises site networks.
 - Select the Replica re-IP (for DR sites with different IP addressing scheme) check box if the IP addressing scheme in the on-premises production site differs from the scheme in the target GCVE site.

New Replication Job ✕


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

Name	Name: <input type="text" value="DR_Replication_on-prem_GCVE"/>
Virtual Machines	Description: <input type="text" value="Created by VEEAMREPLICATIO\Administrator at 9/5/2024 5:04 PM."/>
Destination	Show advanced controls: <input type="checkbox"/> Replica seeding (for low bandwidth DR sites) <input checked="" type="checkbox"/> Network remapping (for DR sites with different virtual networks) <input checked="" type="checkbox"/> Replica re-IP (for DR sites with different IP addressing scheme)
Network	
Re-IP	
Job Settings	<input checked="" type="checkbox"/> 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.
Data Transfer	
Guest Processing	
Schedule	
Summary	

- On the **Virtual Machines** page, select the VMs to be replicated to the NetApp Volumes datastore attached to a GCVE SDDC. Click **Add**, then in the **Add Object** window select the necessary VMs or VM containers and click **Add**. Click **Next**.




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 vSAN cluster will be limited. The rest of the data can be easily placed on Google Cloud NetApp Volumes datastores so that the VMs can be recovered, and the cluster can later be expanded to meet the CPU/mem requirements.



5. On the **Destination** page, select the destination as the GCVE SDDC cluster / hosts and the appropriate resource pool, VM folder and GCNV datastore for the VM replicas. Click **Next** to continue.

New Replication Job ✕

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

Name	Host or cluster:	<input type="text" value="cluster"/>	<input type="button" value="Choose..."/>
Virtual Machines			
Destination	Resource pool:	<input type="text" value="Resources"/>	<input type="button" value="Choose..."/>
Network		Pick resource pool for selected replicas	
Re-IP	VM folder:	<input type="text" value="Replicas"/>	<input type="button" value="Choose..."/>
Job Settings		Pick VM folder for selected replicas	
Data Transfer	Datastore:	<input type="text" value="gcnvdatastore1"/>	<input type="button" value="Choose..."/>
Guest Processing		Pick datastore for selected virtual disks	
Schedule			
Summary			

- On the **Network** page, create the mapping between source and target virtual networks as needed. Click **Next** to continue.

New Replication Job X

Network
Specify how virtual networks map to each other between production and DR sites.

Name

Virtual Machines

Destination

Network

Re-IP

Job Settings

Data Transfer

Guest Processing

Schedule

Summary

Network mapping:

Source network	Target network
VLAN 3376 (DSwitch)	WorkloadVM-Segment (Datacent...

Add...

Edit...

Remove


< Previous

Next >

Finish

Cancel

- On the **Re-IP** page, click on the **Add...** button to add a new re-ip rule. Fill out the source and target VM ip ranges to specify the networking that will be applied to the source VM's in the case of a failover. Use asterisks to specify a range of addresses is indicated for that octet. Click **Next** to continue.



8. On the **Job Settings** page, specify the backup repository that will store metadata for VM replicas, the retention policy and select the button at the bottom for **Advanced...** button at the bottom for additional job settings. Click **Next** to continue.
9. On the **Data Transfer**, select the proxy servers that reside at the source and targets sites, and keep the Direct option selected. WAN accelerators can also be selected here, if configured. Click **Next** to continue.

**Data Transfer**

Choose how VM data should be transferred to the target site.

Name	When replicating between remote sites, we highly recommended that you deploy at least one backup proxy server locally in both sites to allow for direct access to storage.
Virtual Machines	Source proxy: veeamproxyccloud.sddc.netapp.com; veeamproxyccloud2.sddc.netapp.com Choose...
Destination	Target proxy: veeamproxy1.cvsdemo.internal; veeamproxy2.cvsdemo.internal Choose...
Network	
Re-IP	<input checked="" type="radio"/> Direct Best for local and off-site replication over fast links.
Job Settings	<input type="radio"/> Through built-in WAN accelerators Best for off-site replication over slow links due to significant bandwidth savings.
Data Transfer	Source WAN accelerator: <input type="text"/>
Guest Processing	Target WAN accelerator: <input type="text"/>
Schedule	
Summary	

< Previous Next > Finish Cancel

10. On the **Guest Processing** page, check the box for **Enable application-aware processing** as needed and select the **Guest OS credentials**. Click **Next** to continue.

**Guest Processing**

Choose guest OS processing options available for running VMs.

Name	<input checked="" type="checkbox"/> 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: <input type="text" value="Automatic selection"/> Choose...
Re-IP	Guest OS credentials: <input type="text" value="administrator (administrator, last edited: 1 day ago)"/> Add... Manage accounts
Job Settings	Customize guest OS credentials for individual machines and operating systems Credentials...
Data Transfer	Verify network connectivity and credentials for each machine included in the job Test Now
Guest Processing	
Schedule	
Summary	

< Previous **Next >** Finish Cancel

11. On the **Schedule** page, define the times and frequency at which the replication job will run. Click **Next** to continue.

**Schedule**

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

Name	<input checked="" type="checkbox"/> Run the job automatically
Virtual Machines	<input checked="" type="radio"/> Daily at this time: 09:00 AM Everyday Days...
Destination	<input type="radio"/> Monthly at this time: 10:00 PM Fourth Saturday Months...
Network	<input type="radio"/> Periodically every: 1 Hours Schedule...
Re-IP	<input type="radio"/> After this job: [dropdown]
Job Settings	Automatic retry
Data Transfer	<input checked="" type="checkbox"/> Retry failed items processing: 3 times
Guest Processing	Wait before each retry attempt for: 10 minutes
Schedule	Backup window
Summary	<input type="checkbox"/> Terminate the job outside of the allowed backup window Window...
	Long running or accidentally started jobs will be terminated to prevent impact on your production infrastructure during busy hours.
<div>< Previous Next > Finish Cancel</div>	

12. Finally, review the job setting on the **Summary** page. Check the box to **Run the job when I click Finish**, and click on **Finish** to complete creating the replication job.

13. Once run, the replication job can be viewed in the job status window.

DR_Replication_on-prem_GCVE (Full)

Job progress:

0%

0 of 17 VMs

SUMMARY

Duration: 01:47

Processing rate: N/A

Bottleneck: Detecting

DATA

Processed: 0 B (0%)

Read: 0 B

Transferred: 0 B

STATUS

Success: 0

Warnings: 0

Errors: 0

THROUGHPUT (LAST 5 MIN)

Name	Status	Action	Duration
OracleSrv_01	0%	Queued for processing at 9/10/2024 12:47:14 PM	
OracleSrv_02	0%	Required backup infrastructure resources have been assigned	00:00
OracleSrv_03	0%	VM processing started at 9/10/2024 12:47:19 PM	
OracleSrv_04	0%	VM size: 100 GB (21.1 GB used)	
OracleSrv_05	0%	Discovering replica VM	00:00
OracleSrv_06	0%	Resetting CBT per job settings for active fulls	00:31
OracleSrv_07	0%	Getting VM info from vSphere	00:03
OracleSrv_08	0%		
SQLSRV-01	0%		
SQLSRV-02	Pending		
SQLSRV-03	Pending		
SQLSRV-04	Pending		
SQLSRV-05	Pending		

Hide Details

OK




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

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 dependent VMs are already running.

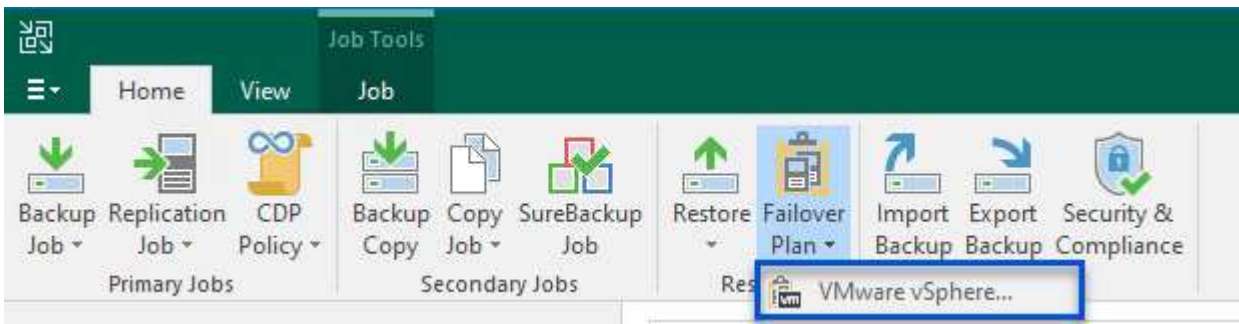
After completing the initial replication or seeding, create a failover plan. This plan serves as a strategic blueprint for orchestrating the failover of dependent VMs, either individually or as a group. It defines the processing order of VMs, incorporates necessary boot delays, and ensures that critical dependent VMs are operational before others. By implementing a well-structured failover plan, organizations can streamline their disaster recovery process, minimizing downtime and maintaining the integrity of interdependent systems during a failover event.

When creating the plan, Veeam Backup & Replication automatically identifies and uses the most recent restore points to initiate the VM replicas.

-  The failover plan can only be created 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**, complete the following steps:

1. On the **Home** view, Click on the **Failover Plan** button in the **Restore** section. In the drop down, select **VMware vSphere...**



2. On the **General** page of the **New Failover Plan** wizard, provide a name and a description to the plan. Pre and Post-failover scripts can be added as required. For instance, run a script to shutdown VMs before starting the replicated VMs.

New Failover Plan



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:

SQL Server DR Plan

Description:

Created by VEEAMREPLICATIO\Administrator at 9/17/2024 6:38 AM.

☐ Pre-failover script:

Browse...

☐ Post-failover script:

Browse...

< Previous

Next >

Finish

Cancel

- On the **Virtual Machines** page, click the button to **Add VM** and select **From replicas...**. Choose the VMs to be part of the failover plan, and then modify the VM boot order and any required boot delays to meet application dependencies.

New Failover Plan



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

Add VM

From infrastructure...

From replicas...

Get Endpoints...

**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
SQLSRV-04	60 sec	less than a day ago (6:1...
SQLSRV-05	60 sec	less than a day ago (5:4...
SQLSRV-01	120 sec	less than a day ago (5:4...
SQLSRV-02	90 sec	less than a day ago (5:4...
SQLSRV-03	60 sec	less than a day ago (5:4...
SQLSRV-06	60 sec	less than a day ago (5:4...
SQLSRV-07	60 sec	less than a day ago (5:4...
SQLSRV-08	60 sec	less than a day ago (5:4...

Add VM

Remove

Set Delay...

↑ Up

↓ Down

< Previous

Apply

Finish

Cancel

Click on **Apply** to continue.

4. Finally review all of the failover plan settings and click on **Finish** to create the failover plan.

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

Run the failover plan

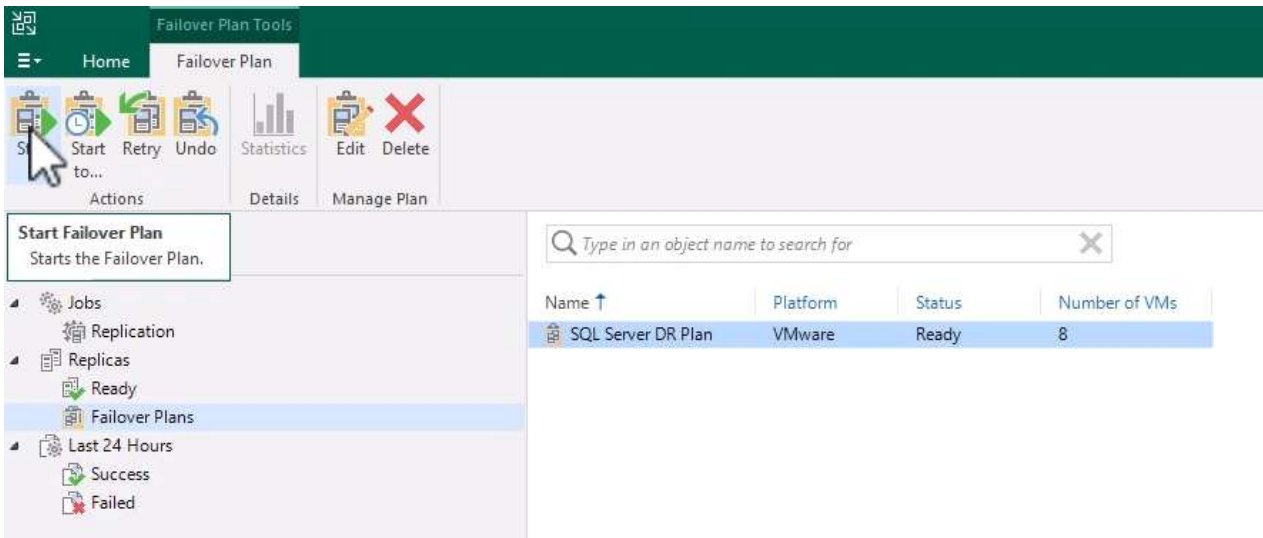
During failover, the source VM in the production site switches over to its replica at the disaster recovery site. As part of the process, Veeam Backup & Replication restores the VM replica to the required restore point and transfers all I/O activities from the source VM to its replica. Replicas serve not only for actual disasters but also for simulating DR drills. In failover simulation, the source VM continues running. Upon completion of necessary tests, the failover can be undone, returning operations to normal.



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

Complete the follow steps to start the failover plan:

1. To get started, in the **Home** view, click on **Replicas > Failover Plans** in the left-hand menu and then on the **Start** button. Alternately, the **Start to...** button can be used to failover to a prior restore point.



2. Monitor the progress of the failover in the **Executing failover plan** window.

Name: **SQL Server DR Plan**Status: **In progress**

Restore type: Failover Plan

Start time: 9/17/2024 10:35:19 AM

Initiated by: VEEAMREPLICATIO\Administrator

[Cancel restore task](#)

VM name	Status
SQLSRV-04	Success
SQLSRV-05	Success
SQLSRV-01	Success
SQLSRV-02	Success
SQLSRV-03	Processing
SQLSRV-06	Success
SQLSRV-07	Processing
SQLSRV-08	Processing

Log

Message	Duration
Performing failover to the latest state	
Building list of machines to process	
Processing VM: SQLSRV-04	0:05:11
Waiting 60 sec before the next VM	0:01:00
Processing VM: SQLSRV-05	0:02:27
Waiting 60 sec before the next VM	0:01:00
Processing VM: SQLSRV-01	0:01:28
Waiting 120 sec before the next VM	0:02:00
Processing VM: SQLSRV-02	0:00:29
Waiting 90 sec before the next VM	0:01:30
Processing VM: SQLSRV-03	0:03:21
Waiting 60 sec before the next VM	0:01:00
Processing VM: SQLSRV-06	0:01:29
Waiting 60 sec before the next VM	0:01:00
Processing VM: SQLSRV-07	0:01:21
Waiting 60 sec before the next VM	0:01:00
Processing VM: SQLSRV-08	0:00:21

Close



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](#).

Failback to the production site

Conducting a failover is considered an intermediate step and needs to be finalized based on the requirement. The options include the following:

- **Failback to production** - Revert to the original VM and synchronize all modifications made during the replica's active period back to the source VM.



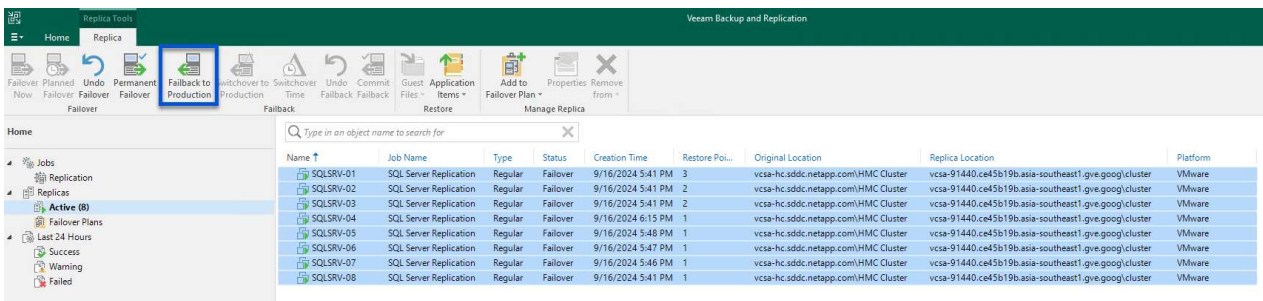
During failback, changes are transferred but not immediately applied. Select **Commit failback** once the original VM's functionality is verified. Alternatively, choose **Undo failback** to revert to the VM replica if the original VM exhibits unexpected behavior.

- **Undo failover** - Revert to the original VM, discarding all changes made to the VM replica during its operational period.
- **Permanent Failover** - Permanently switch from the original VM to its replica, establishing the replica as the new primary VM for ongoing operations.

In this scenario, the "Failback to production" option was selected.


Complete the following steps to perform a failback to the production site:

1. From the **Home** view, click on **Replicas > Active** in the left-hand menu. Select the VMs to be included and click on the **Failback to Production** button in the top menu.



2. On the **Replica** page of the **Failback** wizard, select the replicas to include in the failback job.
3. On the **Destination** page, select **Failback to the original VM** and click on **Next** to continue.

Failback

 **Destination**
Choose the destination for failback operation.

Replica

Destination

Failback Mode

Summary

☒ **Failback to the original VM**
Use if your production site is restored without any infrastructure changes, and the original VM is still present at the same location. Only differences between existing virtual disks and their actual state on replica will be transferred over the network.

☐ **Failback to the original VM restored in a different location**
Use if you have restored the original VM from backup to a location that is different from original. Only differences between existing virtual disks and their actual state on replica will be transferred over the network.


☐ **Failback to the specified location (advanced)**
Use if you do not have original VM remains available anywhere in the failback destination site. Actual state of entire replica's virtual disks will be transferred to the destination site, resulting in significant network traffic.
[Pick backup proxies for data transfer](#)

☐ **Quick rollback (sync changed blocks only)**
Accelerates failback from failovers triggered by a software problem or a user error. Do not use this option if the disaster was caused by a hardware or storage issue, or by a power loss.

< Previous Next > Finish Cancel

4. On the **Failback Mode** page, select **Auto** to start the failback as soon as possible.

Failback

 **Failback Mode**
Specify how and when the failback process should be initiated.

Replica

Destination

Failback Mode

Summary

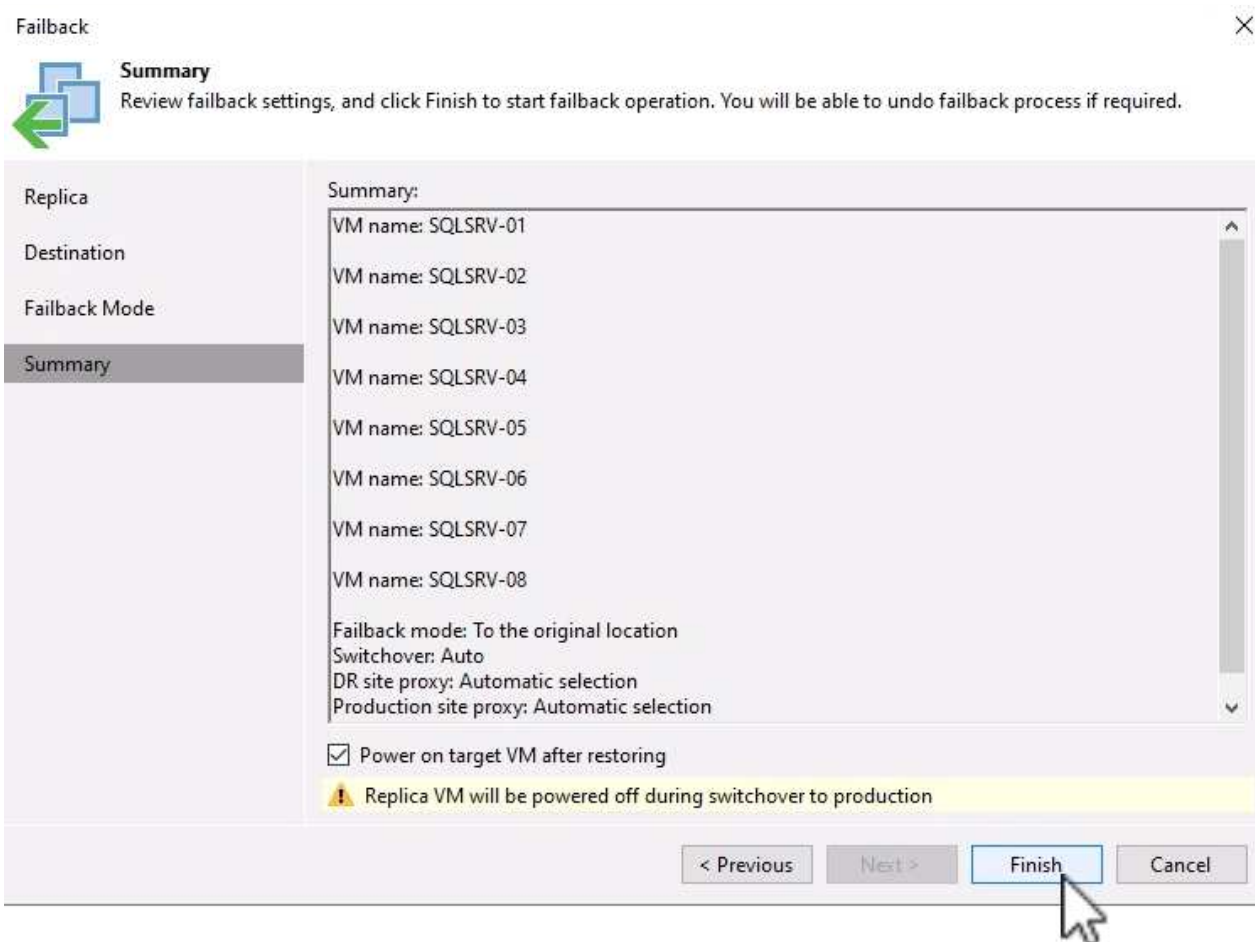
☒ **Auto**
Replicated VMs will be failed over to the production site as soon as they are ready.

☐ **Scheduled**
Perform failover automatically during the scheduled downtime at: 11:45 AM

☐ **Manual**
We will wait for you to issue the failover command manually.

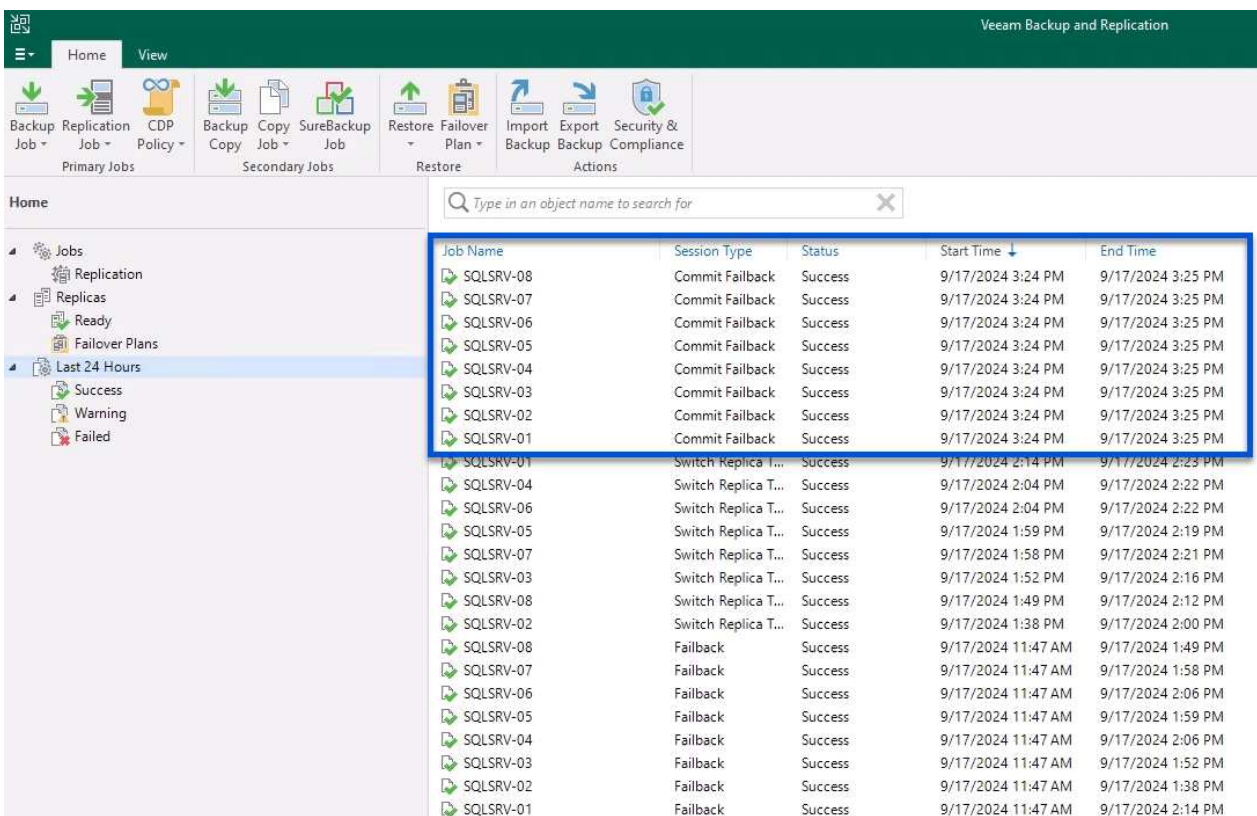
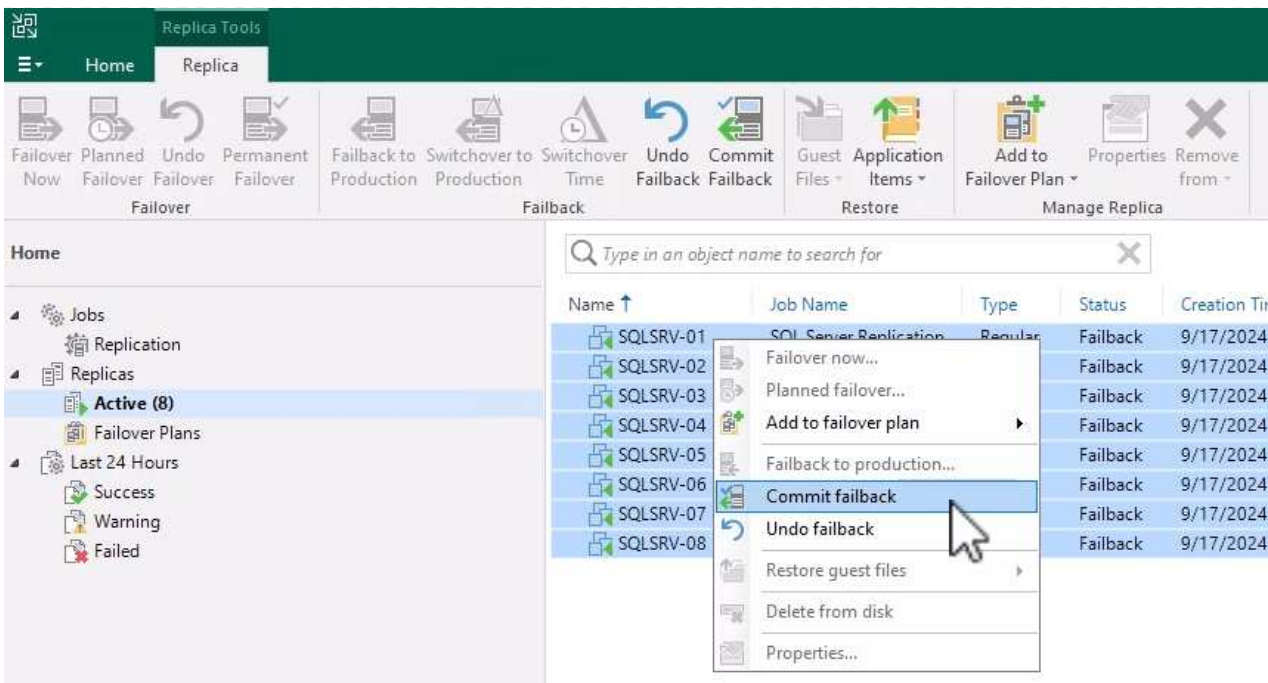
< Previous Next > Finish Cancel

5. On the **Summary** page, choose whether to **Power on target VM after restoring** and then click on Finish to start the failback job.



Failback commit finalizes the failback operation, confirming the successful integration of changes to the production VM. Upon commit, Veeam Backup & Replication resumes regular replication activities for the restored production VM. This changes the status of the restored replica from *Failback* to *Ready*.

1. To commit failback, navigate to **Replicas > Active**, select the VMs to be committed, right click and select **Commit failback**.



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

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

Conclusion

Google Cloud NetApp Volumes datastore functionality empowers Veeam and other validated third-party tools to deliver cost-effective disaster recovery (DR) solutions. By utilizing Pilot light clusters instead of large, dedicated clusters for VM replicas, organizations can significantly reduce expenses. This approach enables tailored DR strategies that leverage existing in-house backup solutions for cloud-based disaster recovery, eliminating the need for additional on-premises datacenters. In the event of a disaster, failover can be initiated with a single click or configured to occur automatically, ensuring business continuity with minimal downtime.

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

<https://netapp.hosted.panopto.com/Panopto/Pages/Embed.aspx?id=b2fb8597-c3fe-49e2-8a84-b1f10118db6d>

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.