



NetApp for GCP / GCVE

NetApp Solutions

NetApp
May 17, 2024

Table of Contents

- NetApp Capabilities for Google Cloud Platform GCVE 1
 - Configuring GCVE in GCP 1
 - NetApp Storage Options for GCVE 1
 - Solution Use Cases 1
 - Protecting Workloads on GCP / GCVE 1
 - Migrating Workloads on GCP / GCVE 9
 - Region Availability – Supplemental NFS datastore for Google Cloud Platform (GCP) 28
 - Security overview - NetApp Cloud Volumes Service (CVS) in Google Cloud 28

NetApp Capabilities for Google Cloud Platform GCVE

Learn more about the capabilities that NetApp brings to the Google Cloud Platform (GCP) Google Cloud VMware Engine (GCVE) - from NetApp as a guest connected storage device or a supplemental NFS datastore to migrating workflows, extending/bursting to the cloud, backup/restore and disaster recovery.

Jump to the section for the desired content by selecting from the following options:

- [Configuring GCVE in GCP](#)
- [NetApp Storage Options for GCVE](#)
- [NetApp / VMware Cloud Solutions](#)

Configuring GCVE in GCP

As with on-premises, planning a cloud based virtualization environment is critical for a successful production-ready environment for creating VMs and migration.

Unresolved directive in ehc/gcp-gcve.adoc - include::.../_include/ehc-config-vmware.adoc[tags=gcp-config;gcp;!ehc-gcp]

NetApp Storage Options for GCVE

NetApp storage can be utilized in several ways - either as guest connected or as a supplemental NFS datastore - within GCP GCVE.

Please visit [Supported NetApp Storage Options](#) for more information.

Unresolved directive in ehc/gcp-gcve.adoc - include::.../_include/ehc-datastore.adoc[tags=gcp-datastore;gcp;!ehc-gcp]

Solution Use Cases

With NetApp and VMware cloud solutions, many use cases are simple to deploy in Azure AVS. se cases are defined for each of the VMware defined cloud areas:

- Protect (includes both Disaster Recovery and Backup / Restore)
- Extend
- Migrate

[Browse the NetApp solutions for Google Cloud GCVE](#)

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

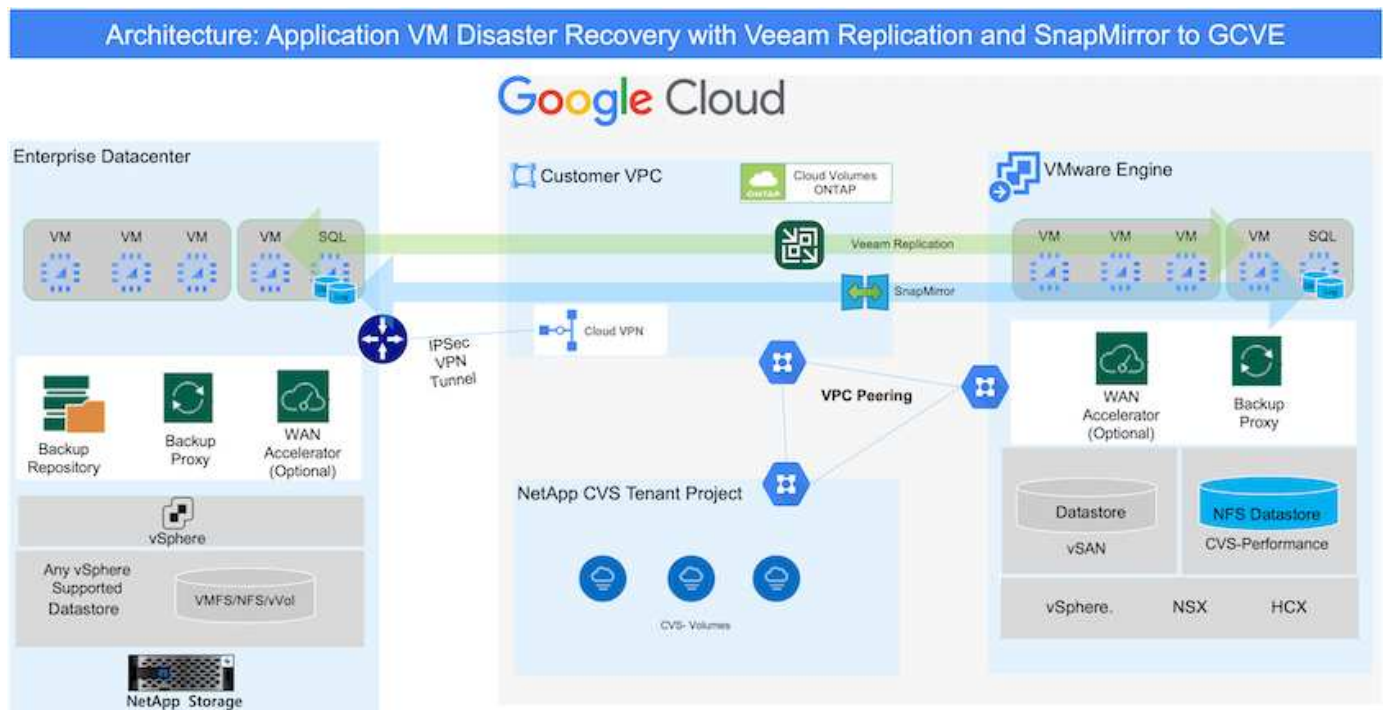
Overview

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 increase 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

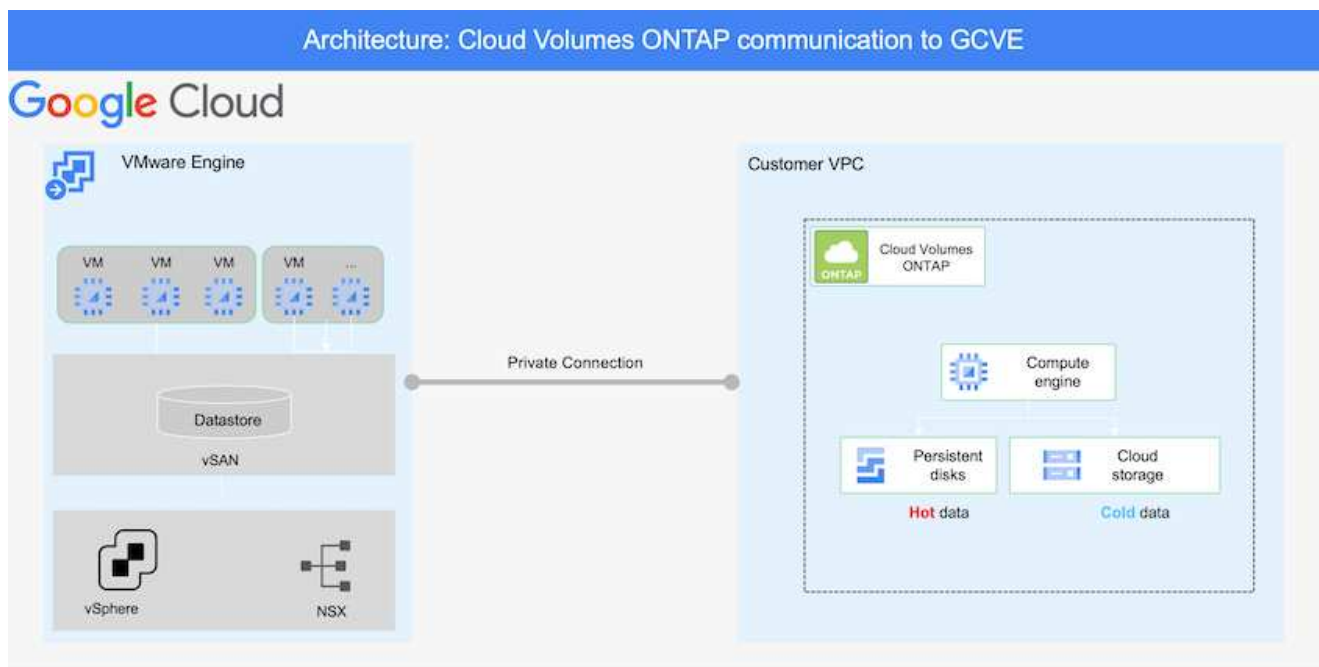
1. Make sure that application data is backed up using SnapCenter with the necessary RPO requirements.

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.

Migrating Workloads on GCP / GCVE

Migrate workloads to NetApp Cloud Volume Service datastore on Google Cloud VMware Engine using VMware HCX - Quickstart guide

One of the most common use cases for the Google Cloud VMware Engine and Cloud Volume Service datastore is the migration of VMware workloads. VMware HCX is a preferred option and provides various migration mechanisms to move on-premises virtual machines (VMs) and its data to Cloud Volume Service NFS datastores.

Author(s): NetApp Solutions Engineering

Overview: Migrating virtual machines with VMware HCX, NetApp Cloud Volume Service datastores, and Google Cloud VMware Engine (GCVE)

VMware HCX is primarily a migration platform that is designed to simplify application migration, workload rebalancing, and even business continuity across clouds. It is included as part of Google Cloud VMware Engine Private Cloud and offers many ways to migrate workloads and can be used for disaster recovery (DR) operations.

This document provides step-by-step guidance for provisioning Cloud Volume Service datastore followed by downloading, deploying, and configuring VMware HCX, including all its main components in on-premises and the Google Cloud VMware Engine side including Interconnect, Network Extension, and WAN optimization for enabling various VM migration mechanisms.



VMware HCX works with any datastore type as the migration is at the VM level. Hence this document is applicable to existing NetApp customers and non-NetApp customers who are planning to deploy Cloud Volume Service with Google Cloud VMware Engine for a cost-effective VMware cloud deployment.

High-level steps

This list provides the high-level steps necessary to pair & Migrate the VMs to HCX Cloud Manager on the Google Cloud VMware Engine side from HCX Connector on-premises:

1. Prepare HCX through the Google VMware Engine portal.
2. Download and deploy the HCX Connector Open Virtualization Appliance (OVA) installer in the on-premises VMware vCenter Server.
3. Activate HCX with the license key.
4. Pair the on-premises VMware HCX Connector with Google Cloud VMware Engine HCX Cloud Manager.
5. Configure the network profile, compute profile, and service mesh.
6. (Optional) Perform network extension to avoid re-IP during migrations.
7. Validate the appliance status and ensure that migration is possible.
8. Migrate the VM workloads.

Prerequisites

Before you begin, make sure the following prerequisites are met. For more information, see this [link](#). After the prerequisites, including connectivity, are in place, download HCX license key from the Google Cloud VMware Engine portal. After the OVA installer is downloaded, proceed with the installation process as described below.

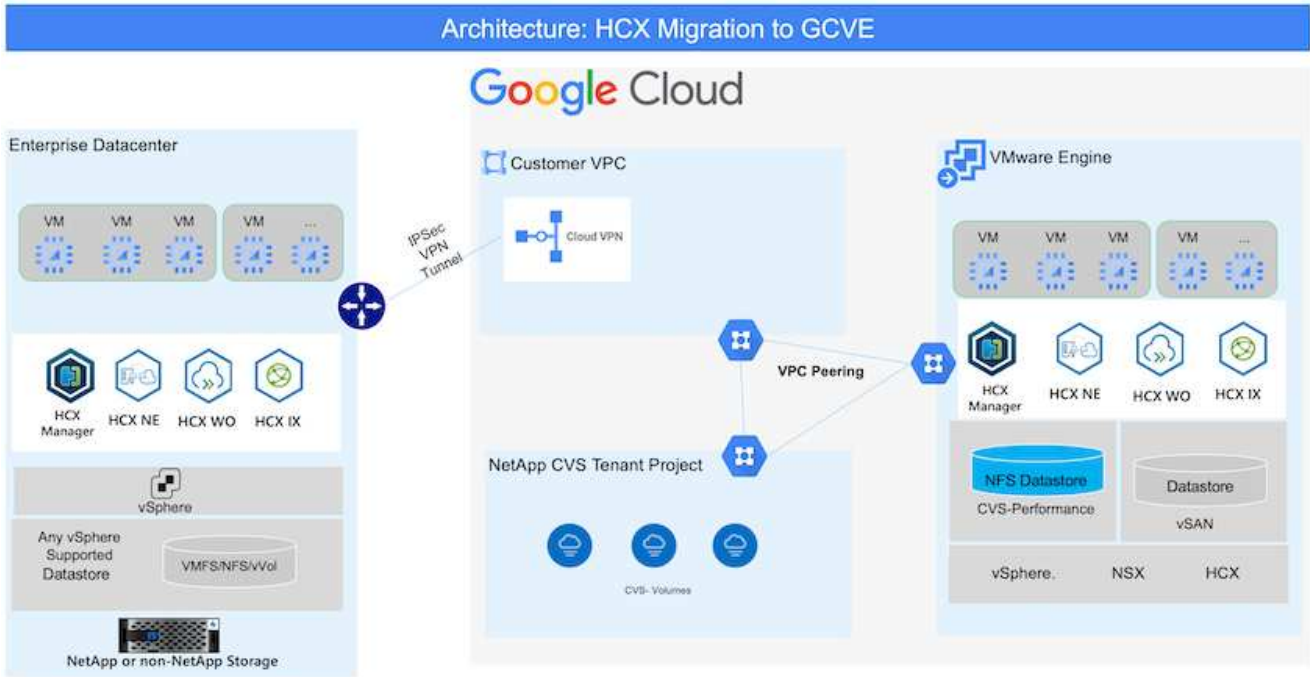


HCX advanced is the default option and VMware HCX Enterprise edition is also available through a support ticket and supported at no additional cost. Refer [this link](#)

- Use an existing Google Cloud VMware Engine software-defined data center (SDDC) or create a private cloud by using this [NetApp link](#) or this [Google link](#).
- Migration of VMs and associated data from the on-premises VMware vSphere- enabled data center requires network connectivity from the data center to the SDDC environment. Before migrating workloads, [set up a Cloud VPN or Cloud Interconnect connection](#) between the on-premises environment and the respective private cloud.
- The network path from on-premises VMware vCenter Server environment to the Google Cloud VMware Engine private cloud must support the migration of VMs by using vMotion.
- Make sure the required [firewall rules and ports](#) are allowed for vMotion traffic between the on-premises vCenter Server and SDDC vCenter.
- Cloud Volume Service NFS volume should be mounted as a datastore in Google Cloud VMware Engine. Follow the steps detailed in this [link](#) to attach Cloud Volume Service datastores to Google Cloud VMware Engines hosts.

High Level Architecture

For testing purposes, the lab environment from on-premises used for this validation was connected through a Cloud VPN, which allows on-premises connectivity to Google Cloud VPC.



For more detailed diagram on HCX, please refer [VMware link](#)

Solution Deployment

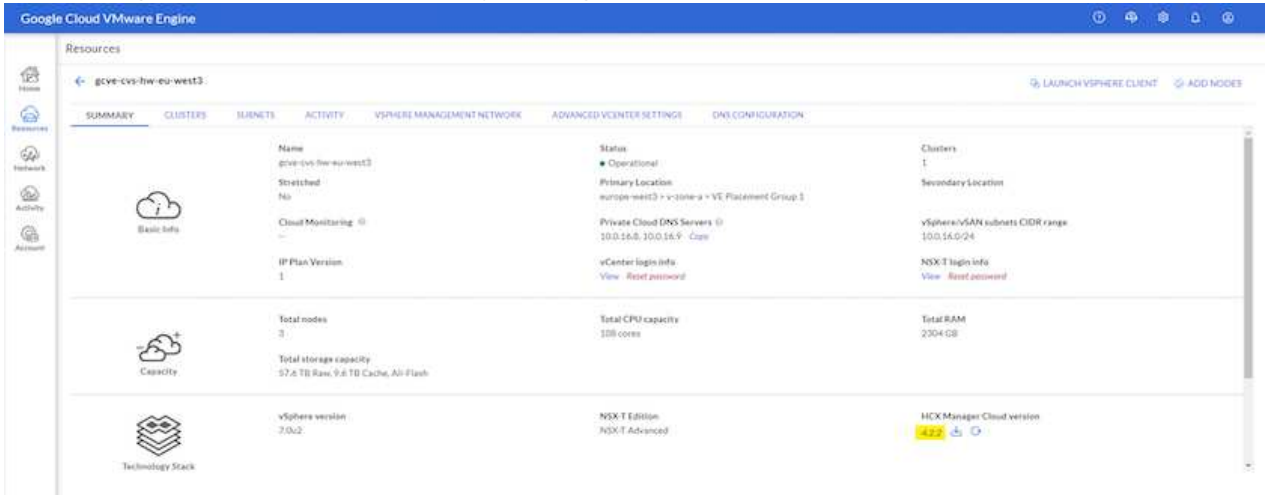
Follow the series of steps to complete the deployment of this solution:

Step 1: Prepare HCX through the Google VMware Engine Portal

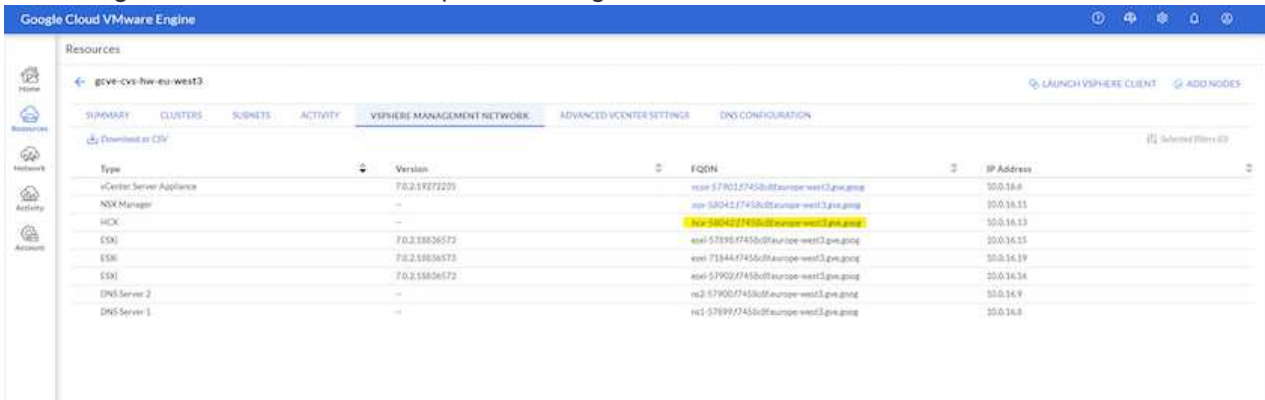
HCX Cloud Manager component automatically gets installed as you provision private cloud with VMware Engine. To prepare for site pairing, complete the following steps:

1. Log in to the Google VMware Engine Portal and sign-in to the HCX Cloud Manager.


You can login to HCX Console either by clicking on the HCX version link



or clicking on HCX FQDN under vSphere Management Network tab.



2. In HCX Cloud Manager, go to **Administration > System Updates**.
3. Click **Request download link** and download the OVA file.



4. Update HCX Cloud Manager to the latest version available from the HCX Cloud Manager UI.

Step 2: Deploy the installer OVA in the on-premises vCenter Server

For the on-premises Connector to connect to the HCX Manager in Google Cloud VMware Engine, make sure the appropriate firewall ports are open in the on-premises environment.

To download and install HCX Connector in the on-premises vCenter Server, complete the following steps:

1. Have the ova downloaded from the HCX Console on Google Cloud VMware Engine as stated in previous step.
2. After the OVA is downloaded, deploy it on to the on-premises VMware vSphere environment by using the **Deploy OVF Template** option.

The screenshot shows the 'Deploy OVF Template' wizard in vSphere. The wizard is at step 1: 'Select an OVF template'. The left sidebar shows the progress: 1. Select an OVF template (selected), 2. Select a name and folder, 3. Select a compute resource, 4. Review details, 5. Select storage, 6. Ready to complete. The main area is titled 'Select an OVF template' and contains the following text: 'Select an OVF template from remote URL or local file system. Enter a URL to download and install the OVF package from the internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.' There are two radio buttons: 'URL' (unselected) and 'Local file' (selected). Below the 'Local file' option is an 'UPLOAD FILES' button and the filename 'VMware-HCX-Connector-4.5.2.0-20914338.ova'. At the bottom right, there are 'CANCEL' and 'NEXT' buttons. The 'NEXT' button is highlighted in blue.

3. Enter all the required information for the OVA deployment, click **Next**, and then click **Finish** to deploy the VMware HCX connector OVA.



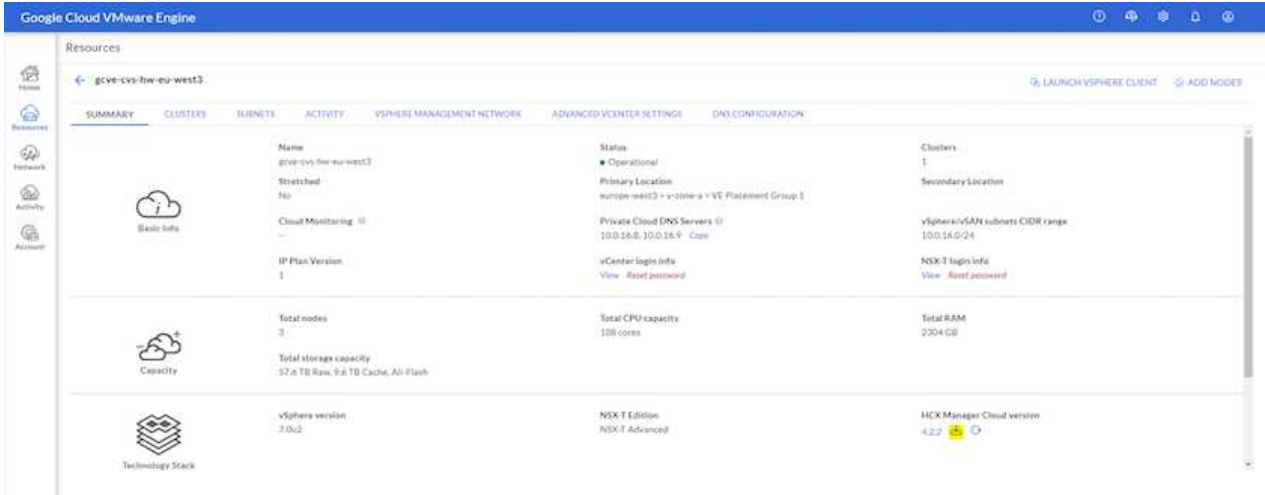
Power on the virtual appliance manually.

For step-by-step instructions, see the [VMware HCX User Guide](#).

Step 3: Activate HCX Connector with the license key

After you deploy the VMware HCX Connector OVA on-premises and start the appliance, complete the following steps to activate HCX Connector. Generate the license key from the Google Cloud VMware Engine portal and activate it in VMware HCX Manager.

1. From the VMware Engine portal, Click on Resources, select the private cloud, and **click on download icon under HCX Manager Cloud Version.**



Open Downloaded file and copy the License Key String.

2. Log into the on-premises VMware HCX Manager at <https://hcxmanagerIP:9443> using administrator credentials.



Use the hcxmanagerIP and password defined during the OVA deployment.

3. In the licensing, enter the key copied from step 3 and click **Activate**.



The on-premises HCX Connector should have internet access.

4. Under **Datacenter Location**, provide the nearest location for installing the VMware HCX Manager on-premises. Click **Continue**.

5. Under **System Name**, update the name and click **Continue**.

6. Click **Yes, Continue**.

7. Under **Connect your vCenter**, provide the fully qualified domain name (FQDN) or IP address of vCenter Server and the appropriate credentials and click **Continue**.



Use the FQDN to avoid connectivity issues later.

8. Under **Configure SSO/PSC**, provide the Platform Services Controller's(PSC) FQDN or IP address and click **Continue**.



For Embedded PSC, Enter the VMware vCenter Server FQDN or IP address.

9. Verify that the information entered is correct and click **Restart**.

10. After the services restart, vCenter Server is displayed as green on the page that appears. Both vCenter Server and SSO must have the appropriate configuration parameters, which should be the

same as the previous page.



This process should take approximately 10 to 20 minutes and for the plug-in to be added to the vCenter Server.


The screenshot displays the HCX Manager dashboard for an appliance named 'HCX-RTP'. The top navigation bar includes 'vm HCX Manager', 'Dashboard', 'Appliance Summary', 'Configuration', and 'Administration'. The top right corner shows the IP address '172.21.254.155', version '4.5.2.0', and user 'admin'. The main content area is divided into several sections:

- System Information:** IP Address: 172.21.254.155, Version: 4.5.2.0, Uptime: 13 days, 21 hours, 6 minutes, Current Time: Thursday, 16 February 2023 05:59:00 PM UTC.
- Resource Usage:** Three progress bars with warning icons: CPU (Used 552 MHz, Capacity 2095 MHz, 26% used), Memory (Used 9535 MB, Capacity 12008 MB, 79% used), and Storage (Used 7.7G, Capacity 84G, 9% used).
- Connected Instances:** Three panels for 'NSX', 'vCenter', and 'SSO'. Each panel shows a 'MANAGE' button and a status bar. The 'vCenter' and 'SSO' status bars show the URL 'https://a300-vcsa01.ehcdc.com' and a green dot, indicating they are connected. A red oval highlights the 'vCenter' and 'SSO' status bars.

Step 4: Pair on-premises VMware HCX Connector with Google Cloud VMware Engine HCX Cloud Manager

After HCX Connector is deployed and configured on on-premises vCenter, establish connection to Cloud Manager by adding the pairing. To configure the site pairing, complete the following steps:

1. To create a site pair between the on-premises vCenter environment and Google Cloud VMware Engine SDDC, log in to the on-premises vCenter Server and access the new HCX vSphere Web Client plug-in.



2. Under Infrastructure, click **Add a Site Pairing**.



Enter the Google Cloud VMware Engine HCX Cloud Manager URL or IP address and the credentials for user with Cloud-Owner-Role privileges for accessing the private cloud.

Connect to Remote Site



Remote HCX URL

https://hcx-58042.f7458c8f.europe-west3.g



Username

cloudowner@gve.local



Password

.....

CANCEL

CONNECT

3. Click **Connect**.





VMware HCX Connector must be able to route to HCX Cloud Manager IP over port 443.

4. After the pairing is created, the newly configured site pairing is available on the HCX Dashboard.

vSphere Client Administrator@EHCDC.COM

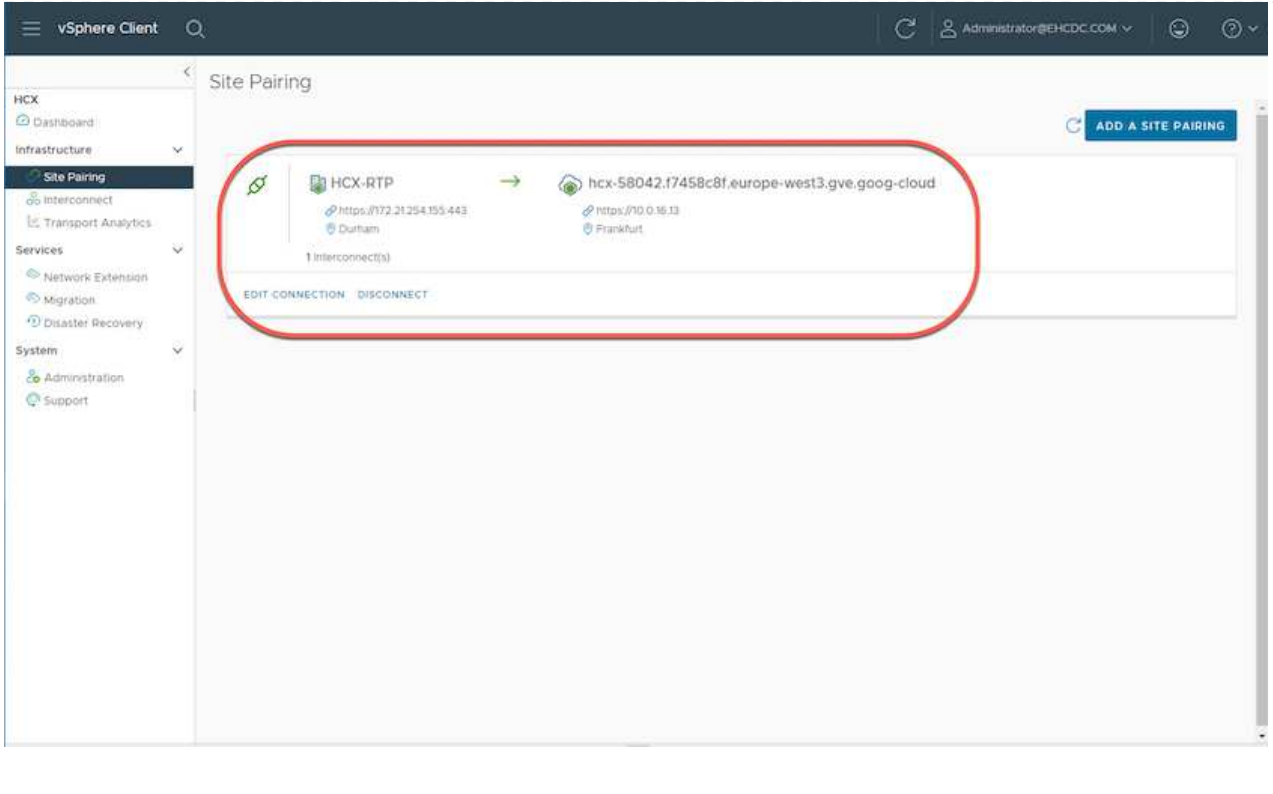
Site Pairing

ADD A SITE PAIRING

 HCX-RTP https://172.21254.155.443 Durham	→	 hcx-58042.f7458c8f.europe-west3.gve.goog-cloud https://10.0.16.13 Frankfurt
--	---	--

1 Interconnect(s)

EDIT CONNECTION DISCONNECT




Step 5: Configure the network profile, compute profile, and service mesh

The VMware HCX Interconnect service appliance provides replication and vMotion-based migration capabilities over the internet and private connections to the target site. The interconnect provides encryption, traffic engineering, and VM mobility. To create an Interconnect service appliance, complete the followings steps:

1. Under Infrastructure, select **Interconnect > Multi-Site Service Mesh > Compute Profiles > Create Compute Profile**.



The compute profiles define the deployment parameters including the appliances that are deployed and which portion of the VMware data center are accessible to HCX service.




2. After the compute profile is created, create the network profiles by selecting **Multi-Site Service Mesh > Network Profiles > Create Network Profile**.

The network profile defines a range of IP address and networks that are used by HCX for its virtual appliances.



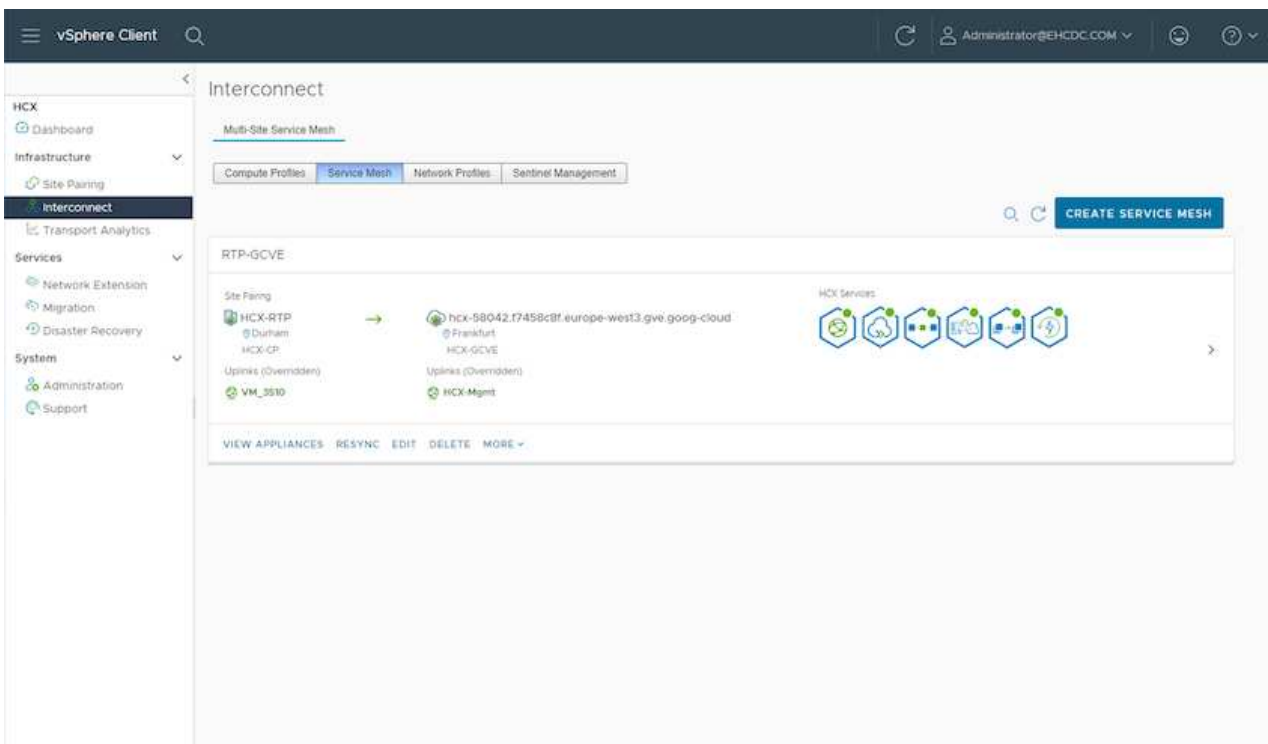
This step requires two or more IP addresses. These IP addresses are assigned from the management network to the Interconnect Appliances.



3. At this time, the compute and network profiles have been successfully created.
4. Create the Service Mesh by selecting the **Service Mesh** tab within the **Interconnect** option and select the on-premises and GCVE SDDC sites.
5. The Service Mesh specifies a local and remote compute and network profile pair.



As part of this process, the HCX appliances are deployed and automatically configured on both the source and target sites in order to create a secure transport fabric.



6. This is the final step of configuration. This should take close to 30 minutes to complete the deployment. After the service mesh is configured, the environment is ready with the IPsec tunnels successfully created to migrate the workload VMs.

The screenshot shows the vSphere Client interface with the 'Interconnect' section selected. The main area displays 'Appliances on HX-CX-RTP' and 'Appliances on hcx-SB042.1745@c8f.europe-west3.gcp.gcpcloud'. The appliances are listed in a table with columns for Name, Type, IP Address, Tunnel Status, and Current Version.

Appliance Name	Appliance Type	IP Address	Tunnel Status	Current Version
BTH-OCVE-0K-R M: 28A5749-4871-4887-4693-4329a3708670 Compute: A300-Outlet01 Storage: A300_MFS_2603	HX-WAN-01	172.21.254.81	Open	4.3.0
BTH-OCVE-0K-S M: 4791521-f644-47c6-4770-46208868906 Compute: A300-Outlet01 Storage: A300_MFS_2603 Network Container: HX-CX-S40300 Extended Networks: V9	HX-MET-EXT	172.21.254.82	Open	4.3.0
BTH-OCVE-WG-R M: 3224758-4779-4776-4996-4688843a9068 Compute: A300-Outlet01 Storage: A300_MFS_2603	HX-WAN-GPT			7.3.0

Appliance Name	Appliance Type	IP Address	Current Version
BTH-OCVE-0K-R1	HX-WAN-01	10.0.0.100	4.3.0
BTH-OCVE-WG-R1	HX-WAN-GPT		7.3.0

Step 6: Migrate workloads

Workloads can be migrated bidirectionally between on-premises and GCVE SDDCs using various VMware HCX migration technologies. VMs can be moved to and from VMware HCX-activated entities using multiple migration technologies such as HCX bulk migration, HCX vMotion, HCX Cold migration, HCX Replication Assisted vMotion (available with HCX Enterprise edition), and HCX OS Assisted Migration (available with the HCX Enterprise edition).

To learn more about various HCX migration mechanisms, see [VMware HCX Migration Types](#).

The HCX-IX appliance uses the Mobility Agent service to perform vMotion, Cold, and Replication Assisted vMotion (RAV) migrations.



The HCX-IX appliance adds the Mobility Agent service as a host object in the vCenter Server. The processor, memory, storage and networking resources displayed on this object do not represent actual consumption on the physical hypervisor hosting the IX appliance.

HCX vMotion

This section describes the HCX vMotion mechanism. This migration technology uses the VMware vMotion protocol to migrate a VM to GCVE. The vMotion migration option is used for migrating the VM state of a single VM at a time. There is no service interruption during this migration method.



Network Extension should be in place (for the port group in which the VM is attached) in order to migrate the VM without the need to make an IP address change.

1. From the on-premises vSphere client, go to Inventory, right-click on the VM to be migrated, and select HCX Actions > Migrate to HCX Target Site.

The screenshot shows the vSphere Client interface for a VM named 'Move2GCVE'. The 'Actions' menu is open, and the 'Migrate to HCX Target Site' option is highlighted with a red circle. The interface displays various VM details, including CPU usage, memory usage, and storage usage. The 'Migrate to HCX Target Site' option is located under the 'HCX Actions' sub-menu.

Task Name	Target
Power On virtual machine	vSAN
Initiate poweroff VM	SanApp ONTAP tools
Move into Replicate pool	NetApp SnapCenter
Reconfigure virtual machine	All Site Recovery actions

Instance Name	Quoted For	Start Time	Completion Time	Error
Powering on the new VM on...	System	03/16/2023, 2:30:50...	03/16/2023, 2:32:31 PM	4300-rcsadm@hcx.com
EMXC CONF Administrator	3 mi	03/16/2023, 2:30:50...	03/16/2023, 2:30:50...	4300-rcsadm@hcx.com
EMXC CONF Administrator	2 mi	03/16/2023, 2:30:33 P...	03/16/2023, 2:30:33 H...	4300-rcsadm@hcx.com
EMXC CONF Administrator	4 mi	03/16/2023, 2:30:19 PM	03/16/2023, 2:30:30...	4300-rcsadm@hcx.com

2. In the Migrate Virtual Machine wizard, select the Remote Site Connection (target GCVE).

HCX: Migrate Virtual Machine

Remote Site Connection:

Source: HCX-RTP / VC: a300-vcsa01.ehcdc.com
Destination: hcx-58042.f7458c8f.europe-west3.gve.goog-cloud / VC: vc5a-57901.f7458c8f.europe-west3.gve.goog
https://10.0.16.13

Transfer and Placement:

(Mandatory: Compute Container) (Mandatory: Storage) (Migration Profile)
(Specify Destination Folder) Same format as source (Optional: Switchover Schedule)

Switchover:

Extended Options:
Edit Extended Options

VM for Migration	Disk / Memory / vCPU	Migration Info
> Move2GCVE	2 GB / 2 GB / 1 vCPU	(Migration profile is not specified)

GO VALIDATE CLOSE

3. Update the mandatory fields (Cluster, Storage, and Destination Network), Click Validate.

HCX: Migrate Virtual Machine

Remote Site Connection:

Source: HCX-RTP / VC: a300-vcsa01.ehcdc.com
Destination: hcx-58042.f7458c8f.europe-west3.gve.goog-cloud / VC: vc5a-57901.f7458c8f.europe-west3.gve.goog
https://10.0.16.13

Transfer and Placement:

Workload gcp-ve-4 (807.6 GB/1TB) vMotion
(Specify Destination Folder) Same format as source (Optional: Switchover Schedule)

Switchover:

Extended Options:
Edit Extended Options Retain MAC

VM for Migration	Disk / Memory / vCPU	Migration Info
> Move2GCVE	2 GB / 2 GB / 1 vCPU	vMotion

Workload gcp-ve-4 (807.6 GB/1TB)
(Specify Destination Folder) Same format as source
Force Power-off VM
Enable Seed Checkpoint
Edit Extended Options Retain MAC
Network adapter 1 (VM_3509) → L2E_VM_3509-3509-a0041a8d


GO VALIDATE CLOSE

4. After the validation checks are complete, click Go to initiate the migration.



The vMotion transfer captures the VM active memory, its execution state, its IP address, and its MAC address. For more information about the requirements and limitations of HCX vMotion, see [Understanding VMware HCX vMotion and Cold Migration](#).

5. You can monitor the progress and completion of the vMotion from the HCX > Migration dashboard.



The target CVS NFS datastore should have sufficient space to handle the migration.

Conclusion

Whether you're targeting all-cloud or hybrid cloud and data residing on any type/vendor storage in on-premises, Cloud Volume Service and HCX provide excellent options to deploy and migrate the application workloads while reducing the TCO by making the data requirements seamless to the application layer. Whatever the use case, choose Google Cloud VMware Engine along with Cloud Volume Service for rapid realization of cloud benefits, consistent infrastructure, and operations across on-premises and multiple clouds, bidirectional portability of workloads, and enterprise-grade capacity and performance. It is the same familiar process and procedures used to connect the storage and migrate VMs using VMware vSphere Replication, VMware vMotion, or even network file copy (NFC).

Takeaways

The key points of this document include:

- You can now use Cloud Volume Service as a datastore on Google Cloud VMware Engine SDDC.
- You can easily migrate data from on-premises to Cloud Volume Service datastore.
- You can easily grow and shrink the Cloud Volume Service datastore to meet the capacity and performance requirements during migration activity.

Videos from Google and VMware for reference

From Google

- [Deploy HCX Connector with GCVE](#)
- [Configure HCX ServiceMesh with GCVE](#)
- [Migrate VM with HCX to GCVE](#)

From VMware

- [HCX Connector deployment for GCVE](#)
- [HCX ServiceMesh configuration for GCVE](#)
- [HCX Workload Migration to GCVE](#)

Where to find additional information

To learn more about the information described in this document, refer to the following website links:

- Google Cloud VMware Engine documentation

<https://cloud.google.com/vmware-engine/docs/overview>

- Cloud Volume Service documentation

<https://cloud.google.com/architecture/partners/netapp-cloud-volumes>

- VMware HCX User Guide

<https://docs.vmware.com/en/VMware-HCX/index.html>

VM Migration to NetApp Cloud Volume Service NFS Datastore on Google Cloud VMware Engine using Veeam Replication feature


Customers who currently use Veeam for their data protection requirements continue using that solution to migrate the workloads to GCVE and enjoy the benefits of NetApp Cloud Volume Service NFS Datastores.

Overview

Authors: Suresh Thoppay, NetApp

VM Workloads running on VMware vSphere can be migrated to Google Cloud VMware Engine (GCVE) utilizing Veeam Replication feature.

This document provides a step-by-step approach for setting up and performing VM migration that uses NetApp Cloud Volume Service, Veeam, and the Google Cloud VMware Engine (GCVE).



Assumptions

This document assumes you have either Google Cloud VPN or Cloud Interconnect or other networking option in place to establish network connectivity from existing vSphere servers to Google Cloud VMware Engine.



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 Migration Solution

Solution Deployment Overview

1. Make sure NFS datastore from NetApp Cloud Volume Service is mounted on GCVE vCenter.
2. Ensure Veeam Backup Recovery is deployed on existing VMware vSphere environment
3. Create Replication Job to start replicating virtual machines to Google Cloud VMware Engine instance.
4. Perform Failover of Veeam Replication Job.
5. Perform Permanent Failover on Veeam.

Deployment Details

Make sure NFS datastore from NetApp Cloud Volume Service is mounted on GCVE vCenter

Login to GCVE vCenter and ensure NFS datastore with sufficient space is available. If not, Please refer [Mount NetApp CVS as NFS datastore on GCVE](#)

Ensure Veeam Backup Recovery is deployed on existing VMware vSphere environment

Please refer [Veeam Replication Components](#) documentation to install required components.

Create Replication Job to start replicating virtual machines to Google Cloud VMware Engine instance.

Both on-premises vCenter and GCVE vCenter needs to be registered with Veeam. [Setup vSphere VM Replication Job](#)

Here is a video explaining how to [Configure Replication Job](#).



Replica VM can have different IP from the source VM and can also be connected to different port group. For more details, check the video above.

Perform Failover of Veeam Replication Job

To Migrate VMs, perform [Perform Failover](#)

Perform Permanent Failover on Veeam.

To treat GCVE as your new source environment, perform [Permanent Failover](#)

Benefits of this solution

- Existing Veeam backup infrastructure can be utilized for migration.
- Veeam Replication allows changing VM IP addresses on target site.
- Has ability to remap existing data replicated outside of Veeam (like replicated data from BlueXP)
- Has ability to specify different network portgroup on target site.
- Can specify the order of VMs to power on.
- Utilizes VMware Change Block Tracking to minimize the amount of data to send across WAN.
- Capability to execute pre and post scripts for replication.
- Capability to execute pre and post scripts for snapshots.

Region Availability – Supplemental NFS datastore for Google Cloud Platform (GCP)

Learn more about the the Global Region support for GCP, GCVE and CVS.



NFS datastore will be available in regions where both services (GCVE and CVS Performance) are available.

Unresolved directive in ehc/gcp-regions.adoc - include::.../../_include/gcp-region-support.adoc[]

Security overview - NetApp Cloud Volumes Service (CVS) in Google Cloud

TR-4918: Security overview - NetApp Cloud Volumes Service in Google Cloud

Oliver Krause, Justin Parisi, NetApp

Security, particularly in the cloud where infrastructure is outside of the control of storage

administrators, is paramount to trusting your data to service offerings provided by cloud providers. This document is an overview of the security offerings that NetApp [Cloud Volumes Service provides in Google Cloud](#).

Intended audience

This document's intended audience includes, but is not limited to, the following roles:

- Cloud providers
- Storage administrators
- Storage architects
- Field resources
- Business decision makers

If you have questions about the content of this technical report, see the section [“Contact us.”](#)

Abbreviation	Definition
CVS-SW	Cloud Volumes Service, Service Type CVS
CVS-Performance	Cloud Volume Service, Service Type CVS-Performance
PSA	

How Cloud Volumes Service in Google Cloud secures your data

Cloud Volumes Service in Google Cloud provides a multitude of ways to natively secure your data.

Secure architecture and tenancy model

Cloud Volumes Service provides a secure architecture in Google Cloud by segmenting the service management (control plane) and the data access (data plane) across different endpoints so that neither can impact the other (see the section [“Cloud Volumes Service architecture”](#)). It uses Google's [private services access](#) (PSA) framework to provide the service. This framework distinguishes between the service producer, which is provided and operated by NetApp, and the service consumer, which is a Virtual Private Cloud (VPC) in a customer project, hosting the clients that want to access Cloud Volumes Service file shares.

In this architecture, tenants (see the section [“Tenancy model”](#)) are defined as Google Cloud projects that are completely isolated from each other unless explicitly connected by the user. Tenants allow complete isolation of data volumes, external name services, and other essential pieces of the solution from other tenants using the Cloud Volumes Service volume platform. Because the Cloud Volumes Service platform is connected through VPC peering, that isolation applies to it also. You can enable sharing of Cloud Volumes Service volumes between multiple projects by using a shared-VPC (see the section [“Shared VPCs”](#)). You can apply access controls to SMB shares and NFS exports to limit who or what can view or modify datasets.

Strong identity management for the control plane

In the control plane where Cloud Volumes Service configuration takes place, identity management is managed by using [Identity Access Management \(IAM\)](#). IAM is a standard service that enables you to control authentication (logins) and authorization (permissions) to Google Cloud project instances. All configuration is

performed with Cloud Volumes Service APIs over a secure HTTPS transport using TLS 1.2 encryption, and authentication is performed by using JWT tokens for added security. The Google console UI for Cloud Volumes Service translates user input into Cloud Volumes Service API calls.

Security hardening - Limiting attack surfaces

Part of effective security is limiting the number of attack surfaces available in a service. Attack surfaces can include a variety of things, including data at-rest, in-flight transfers, logins, and the datasets themselves.

A managed service removes some of the attack surfaces inherently in its design. Infrastructure management, as described in the section [“Service operation,”](#) is handled by a dedicated team and is automated to reduce the number of times a human actually touches configurations, which helps reduce the number of intentional and unintentional errors. Networking is fenced off so that only necessary services can access one another. Encryption is baked into the data storage and only the data plane needs security attention from Cloud Volumes Service administrators. By hiding most of the management behind an API interface, security is achieved by limiting the attack surfaces.

Zero Trust model

Historically, IT security philosophy has been to trust but verify, and manifested as relying solely on external mechanisms (such as firewalls and intrusion detection systems) to mitigate threats. However, attacks and breaches evolved to bypass the verification in environments through phishing, social engineering, insider threats and other methods that provide the verification to enter networks and wreak havoc.

Zero Trust has become a new methodology in security, with the current mantra being “trust nothing while still verifying everything.” Therefore, nothing is allowed access by default. This mantra is enforced in a variety of ways, including standard firewalls and intrusion detection systems (IDS) and also with the following methods:

- Strong authentication methods (such as AES-encrypted Kerberos or JWT tokens)
- Single strong sources of identities (such as Windows Active Directory, Lightweight Directory Access Protocol (LDAP), and Google IAM)
- Network segmentation and secure multitenancy (only tenants are allowed access by default)
- Granular access controls with Least Privileged Access policies
- Small exclusive lists of dedicated, trusted administrators with digital audit and paper trails

Cloud Volumes Service running in Google Cloud adheres to the Zero Trust model by implementing the “trust nothing, verify everything” stance.

Encryption

Encrypt data at-rest (see the section [“Data encryption at rest”](#)) by using XTS-AES-256 ciphers with NetApp Volume Encryption (NVE) and in-flight with [“SMB encryption”](#) or NFS Kerberos 5p support. Rest easy knowing cross-region replication transfers are protected by TLS 1.2 encryption (see the section [“Cross-region replication”](#)). In addition, Google networking also provides encrypted communications (see the section [“Data encryption in transit”](#)) for an added layer of protection against attacks. For more information about transport encryption, see the section [“Google Cloud network”](#).

Data protection and backups

Security isn’t just about the prevention of attacks. It is also about how we recover from attacks if or when they occur. This strategy includes data protection and backups. Cloud Volumes Service provides methods to replicate to other regions in case of outages (see the section [“Cross-region replication”](#)) or if a dataset is affected by a ransomware attack. It can also perform asynchronous backups of data to locations outside of the

Cloud Volumes Service instance by using [Cloud Volumes Service backup](#). With regular backups, mitigation of security events can take less time and save money and angst for administrators.

Fast ransomware mitigation with industry leading Snapshot copies

In addition to data protection and backups, Cloud Volumes Service provides support for immutable Snapshot copies (see the section [“Immutable Snapshot copies”](#)) of volumes that allow recovery from ransomware attacks (see the section [“Service operation”](#)) within seconds of discovering the issue and with minimal disruption. Recovery time and effects depend on the Snapshot schedule, but you can create Snapshot copies that provide as little as one-hour deltas in ransomware attacks. Snapshot copies have a negligible effect on performance and capacity usage and are a low-risk, high-reward approach to protecting your datasets.

Security considerations and attack surfaces

The first step in understanding how to secure your data is identifying the risks and potential attack surfaces.

These include (but are not limited to) the following:

- Administration and logins
- Data at rest
- Data in flight
- Network and firewalls
- Ransomware, malware, and viruses

Understanding attack surfaces can help you to better secure your environments. Cloud Volumes Service in Google Cloud already considers many of these topics and implements security functionality by default, without any administrative interaction.

Ensuring secure logins

When securing your critical infrastructure components, it is imperative to make sure that only approved users can log in and manage your environments. If bad actors breach your administrative credentials, then they have the keys to the castle and can do anything they want—change configurations, delete volumes and backups, create backdoors, or disable Snapshot schedules.

Cloud Volumes Service for Google Cloud provides protection against unauthorized administrative logins through the obfuscation of storage as a service (StaaS). Cloud Volumes Service is completely maintained by the cloud provider with no availability to login externally. All setup and configuration operations are fully automated, so a human administrator never has to interact with the systems except in very rare circumstances.

If login is required, Cloud Volumes Service in Google Cloud secures logins by maintaining a very short list of trusted administrators that have access to log in to the systems. This gatekeeping helps reduce the number of potential bad actors with access. Additionally, the Google Cloud networking hides the systems behind layers of network security and exposes only what is needed to the outside world. For information about the Google Cloud, Cloud Volumes Service architecture, see the section [“Cloud Volumes Service architecture.”](#)

Cluster administration and upgrades

Two areas with potential security risks include cluster administration (what happens if a bad actor has admin access) and upgrades (what happens if a software image is compromised).

Storage administration protection

Storage provided as a service removes the added risk of exposure to administrators by removing that access to end users outside of the cloud data center. Instead, the only configuration done is for the data access plane by customers. Each tenant manages their own volumes, and no tenant can reach other Cloud Volumes Service instances. The service is managed by automation, with a very small list of trusted administrators given access to the systems through the processes covered in the section [“Service operation.”](#)

The CVS-Performance service type offers cross-region replication as an option to provide data protection to a different region in the event of a region failure. In those cases, Cloud Volumes Service can be failed over to the unaffected region to maintain data access.

Service upgrades

Updates help protect vulnerable systems. Each update provides security enhancements and bug fixes that minimize attack surfaces. Software updates are downloaded from centralized repositories and are validated before the updates are allowed to verify that official images are used and that the upgrades are not compromised by bad actors.

With Cloud Volumes Service, updates are handled by the cloud provider teams, which removes risk exposure for administrator teams by providing experts well versed in configuration and upgrades that have automated and fully tested the process. Upgrades are nondisruptive, and Cloud Volumes Service maintains the latest updates for best overall results.

For information about the administrator team that performs these service upgrades, see the section [“Service operation.”](#)

Securing data at-rest

Data-at-rest encryption is important to protect sensitive data in the event of a disk that is stolen, returned, or repurposed. Data in Cloud Volumes Service is protected at rest by using software-based encryption.

- Google-generated keys are used for CVS-SW.
- For CVS-Performance, the per-volume keys are stored in a key manager built into Cloud Volumes Service, which uses NetApp ONTAP CryptoMod to generate AES-256 encryption keys. CryptoMod is listed on the CMVP FIPS 140-2 validated modules list. See [FIPS 140-2 Cert #4144](#).

Starting in November 2021, preview Customer-managed Encryption (CMEK) functionality was made available for CVS-Performance. This functionality allows you to encrypt the per-volume keys with per-project, per-region master-keys that are hosted in Google Key Management Service (KMS). KMS enables you to attach external key managers.

For details about how to configure KMS for CVS-Performance, [see the Cloud Volumes Service documentation](#).

For more information about architecture, see the section [“Cloud Volumes Service architecture.”](#)

Securing data in-flight

In addition to securing data at rest, you must also be able to secure data when it is in flight between the Cloud Volumes Service instance and a client or replication target. Cloud Volumes Service provides encryption for in-flight data over NAS protocols by using encryption methods such as SMB encryption using Kerberos, the signing/sealing of packets, and NFS Kerberos 5p for end-to-end encryption of data transfers.

Replication of Cloud Volumes Service volumes uses TLS 1.2, which takes advantage of AES-GCM encryption methods.

Most insecure in-flight protocols such as telnet, NDMP, and so on are disabled by default. DNS, however, is not encrypted by Cloud Volumes Service (no DNS Sec support) and should be encrypted by using external network encryption when possible. See the section [“Data encryption in transit”](#) for more information about securing data in-flight.

For information about NAS protocol encryption, see the section [“NAS protocols.”](#)

Users and groups for NAS permissions

Part of securing your data in the cloud involves proper user and group authentication, where the users accessing the data are verified as real users in the environment and the groups contain valid users. These users and groups provide initial share and export access, as well as permission validation for files and folders in the storage system.

Cloud Volumes Service uses standard Active Directory-based Windows user and group authentication for SMB shares and Windows-style permissions. The service can also leverage UNIX identity providers such as LDAP for UNIX users and groups for NFS exports, NFSv4 ID validation, Kerberos authentication, and NFSv4 ACLs.



Currently only Active Directory LDAP is supported with Cloud Volumes Service for LDAP functionality.

Detection, prevention and mitigation of ransomware, malware, and viruses

Ransomware, malware, and viruses are a persistent threat to administrators, and detection, prevention, and mitigation of those threats are always top of mind for enterprise organizations. A single ransomware event on a critical dataset can potentially cost millions of dollars, so it is beneficial to do what you can to minimize the risk.

Although Cloud Volumes Service currently doesn't include native detection or prevention measures, such as antivirus protection or [automatic ransomware detection](#), there are ways to quickly recover from a ransomware event by enabling regular Snapshot schedules. Snapshot copies are immutable and read only pointers to changed blocks in the file system, are near instantaneous, have minimal impact on performance, and only use up space when data is changed or deleted. You can set schedules for Snapshot copies to match your desired acceptable recovery point objective (RPO)/recovery time objective (RTO) and can keep up to 1,024 Snapshot copies per volume.

Snapshot support is included at no additional cost (beyond data storage charges for changed blocks/data retained by Snapshot copies) with Cloud Volumes Service and, in the event of a ransomware attack, can be used to roll back to a Snapshot copy before the attack occurred. Snapshot restores take just seconds to complete, and you then can get back to serving data as normal. For more information, see [The NetApp Solution for Ransomware](#).

Preventing ransomware from affecting your business requires a multilayered approach that includes one or more of the following:

- Endpoint protection
- Protection against external threats through network firewalls
- Detection of data anomalies
- Multiple backups (onsite and offsite) of critical datasets
- Regular restore tests of backups
- Immutable read-only NetApp Snapshot copies
- Multifactor authentication for critical infrastructure

- Security audits of system logins

This list is far from exhaustive but is a good blueprint to follow when dealing with the potential of ransomware attacks. Cloud Volumes Service in Google Cloud provides several ways to protect against ransomware events and reduce their effects.

Immutable Snapshot copies

Cloud Volumes Service natively provides immutable read-only Snapshot copies that are taken on a customizable schedule for quick point-in-time recovery in the event of data deletion or if an entire volume has been victimized by a ransomware attack. Snapshot restores to previous good Snapshot copies are fast and minimize data loss based on the retention period of your Snapshot schedules and RTO/RPO. The performance effect with Snapshot technology is negligible.

Because Snapshot copies in Cloud Volumes Service are read-only, they cannot be infected by ransomware unless the ransomware has proliferated into the dataset unnoticed and Snapshot copies have been taken of the data infected by ransomware. This is why you must also consider ransomware detection based on data anomalies. Cloud Volumes Service does not currently provide detection natively, but you can use external monitoring software.

Backups and restores

Cloud Volumes Service provides standard NAS client backup capabilities (such as backups over NFS or SMB).

- CVS-Performance offers cross-region volume replication to other CVS-Performance volumes. For more information, see [volume replication](#) in the Cloud Volumes Service documentation.
- CVS-SW offers service-native volume backup/restore capabilities. For more information, see [cloud backup](#) in the Cloud Volumes Service documentation.

Volume replication provides an exact copy of the source volume for fast failover in the case of a disaster, including ransomware events.

Cross-region replication

CVS-Performance enables you to securely replicate volumes across Google Cloud regions for data protection and archive use cases by using TLS1.2 AES 256 GCM encryption on a NetApp-controlled backend service network using specific interfaces used for replication running on Google's network. A primary (source) volume contains the active production data and replicates to a secondary (destination) volume to provide an exact replica of the primary dataset.

Initial replication transfers all blocks, but updates only transmit the changed blocks in a primary volume. For instance, if a 1TB database that resides on a primary volume is replicated to the secondary volume, then 1TB of space is transferred on the initial replication. If that database has a few hundred rows (hypothetically, a few MB) that change between the initialization and the next update, only the blocks with the changed rows are replicated to the secondary (a few MB). This helps to make sure that the transfer times remain low and keeps replication charges down.

All permissions on files and folders are replicated to the secondary volume, but share access permissions (such as export policies and rules or SMB shares and share ACLs) must be handled separately. In the case of a site failover, the destination site should leverage the same name services and Active Directory domain connections to provide consistent handling of user and group identities and permissions. You can use a secondary volume as a failover target in the event of a disaster by breaking the replication relationship, which converts the secondary volume to read-write.

Volume replicas are read-only, which provides an immutable copy of data offsite for quick recovery of data in

instances where a virus has infected data or ransomware has encrypted the primary dataset. Read-only data won't be encrypted, but, if the primary volume is affected and replication occurs, the infected blocks also replicate. You can use older, non-affected Snapshot copies to recover, but SLAs might fall out of range of the promised RTO/RPO depending on how quickly an attack is detected.

In addition, you can prevent malicious administrative actions, such as volume deletions, Snapshot deletions, or Snapshot schedule changes, with cross-region replication (CRR) management in Google Cloud. This is done by creating custom roles that separate volume administrators, who can delete source volumes but not break mirrors and therefore cannot delete destination volumes, from CRR administrators, who cannot perform any volume operations. See [Security Considerations](#) in the Cloud Volumes Service documentation for permissions allowed by each administrator group.

Cloud Volumes Service backup

Although Cloud Volumes Service provides high data durability, external events can cause data loss. In the event of a security event such as a virus or ransomware, backups and restores become critical for resumption of data access in a timely manner. An administrator might accidentally delete a Cloud Volumes Service volume. Or users simply want to retain backup versions of their data for many months and keeping the extra Snapshot copy space inside the volume becomes a cost challenge. Although Snapshot copies should be the preferred way to keep backup versions for the last few weeks to restore lost data from them, they are sitting inside the volume and are lost if the volume goes away.

For all these reasons, NetApp Cloud Volumes Service offers backup services through [Cloud Volumes Service backup](#).

Cloud Volumes Service backup generates a copy of the volume on Google Cloud Storage (GCS). It only backs up the actual data stored within the volume, not the free space. It works as incremental forever, meaning it transfers the volume content once and from there on continues backing up changed data only. Compared to classical backup concepts with multiple full backups, it saves large amounts of backup storage, reducing cost. Because the monthly price of backup space is lower compared to a volume, it is an ideal place to keep backup versions longer.

Users can use a Cloud Volumes Service backup to restore any backup version to the same or a different volume within the same region. If the source volume is deleted, the backup data is retained and needs to be managed (for example, deleted) independently.

Cloud Volumes Service backup is built into Cloud Volumes Service as option. Users can decide which volumes to protect by activating Cloud Volumes Service backup on a per-volume basis. See the [Cloud Volumes Service backup documentation](#) for information about backups, the [number of maximum backup versions supported](#), scheduling, and [pricing](#).

All backup data of a project is stored within a GCS bucket, which is managed by the service and not visible to the user. Each project uses a different bucket. Currently, the buckets are in same region as the Cloud Volumes Service volumes, but more options are being discussed. Consult the documentation for the latest status.

Data transport from a Cloud Volumes Service bucket to GCS uses service-internal Google networks with HTTPS and TLS1.2. Data is encrypted at-rest with Google-managed keys.

To manage Cloud Volumes Service backup (creating, deleting, and restoring backups), a user must have the [roles/netappcloudvolumes.admin](#) role.

Architecture

Overview

Part of trusting a cloud solution is understanding the architecture and how it is secured. This section calls out different aspects of the Cloud Volumes Service architecture in Google to help alleviate potential concerns about how data is secured, as well as call out areas where additional configuration steps might be required to obtain the most secure deployment.

The general architecture of Cloud Volumes Service can be broken down into two main components: the control plane and the data plane.

Control plane

The control plane in Cloud Volumes Service is the backend infrastructure managed by Cloud Volumes Service administrators and NetApp native automation software. This plane is completely transparent to end users and includes networking, storage hardware, software updates, and so on to help deliver value to a cloud-resident solution such as Cloud Volumes Service.

Data plane

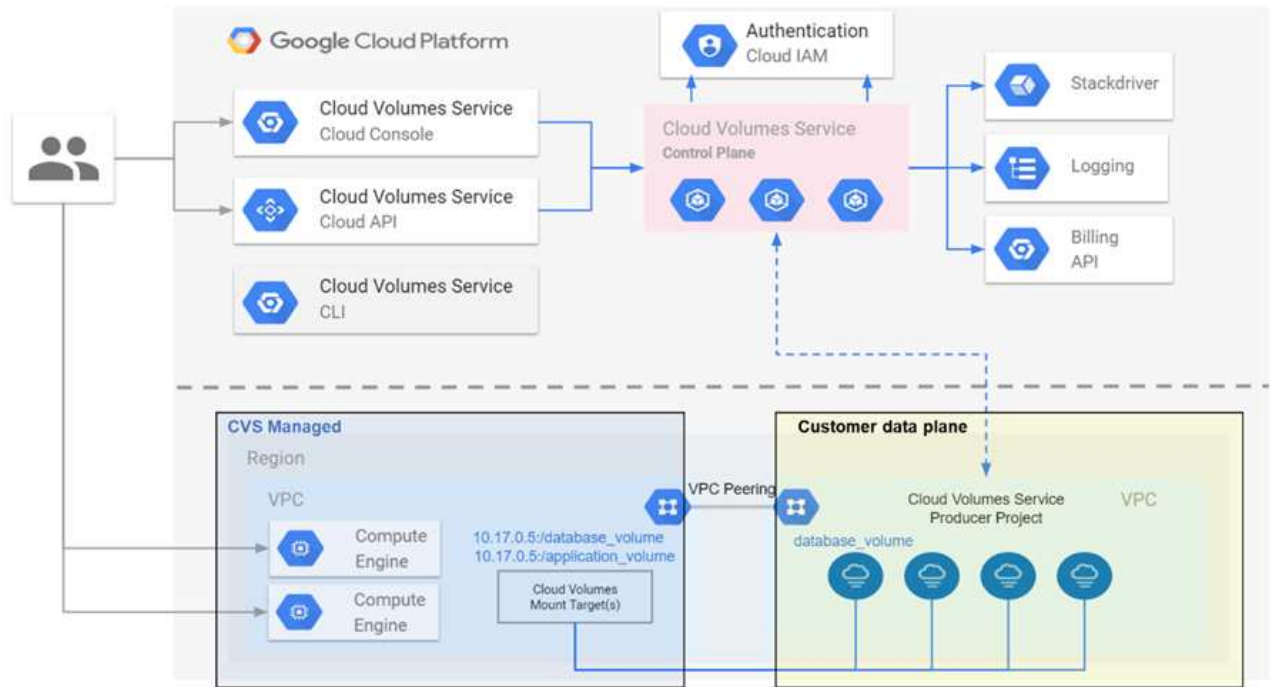
The data plane in Cloud Volumes Service includes the actual data volumes and the overall Cloud Volumes Service configuration (such as access control, Kerberos authentication, and so on). The data plane is entirely under the control of the end users and the consumers of the Cloud Volumes Service platform.

There are distinct differences in how each plane is secured and managed. The following sections cover these differences, starting with a Cloud Volumes Service architecture overview.

Cloud Volumes Service architecture

In a manner similar to other Google Cloud native services such as CloudSQL, Google Cloud VMware Engine (GCVE), and FileStore, Cloud Volumes Service uses [Google PSA](#) to deliver the service. In PSA, services are built inside a service producer project, which uses [VPC network peering](#) to connect to the service consumer. The service producer is provided and operated by NetApp, and the service consumer is a VPC in a customer project, hosting the clients that want to access Cloud Volumes Service file shares.

The following figure, referenced from the [architecture section](#) of the Cloud Volumes Service documentation, shows a high-level view.



The part above the dotted line shows the control plane of the service, which controls the volume lifecycle. The part below the dotted line shows the data plane. The left blue box depicts the user VPC (service consumer), the right blue box is the service producer provided by NetApp. Both are connected through VPC peering.

Tenancy model

In Cloud Volumes Service, individual projects are considered unique tenants. This means that manipulation of volumes, Snapshot copies, and so on are performed on a per-project basis. In other words, all volumes are owned by the project that they were created in and only that project can manage and access the data inside of them by default. This is considered the control plane view of the service.

Shared VPCs

On the data plane view, Cloud Volumes Service can connect to a shared VPC. You can create volumes in the hosting project or in one of the service projects connected to the shared VPC. All projects (host or service) connected to that shared VPC are able to reach the volumes at the network layer (TCP/IP). Because all clients with network connectivity on the shared-VPC can potentially access the data through NAS protocols, access control on the individual volume (such as user/group access control lists (ACLs) and hostnames/IP addresses for NFS exports) must be used to control who can access the data.

You can connect Cloud Volumes Service to up to five VPCs per customer project. On the control plane, the project enables you to manage all created volumes, no matter which VPC they are connected to. On the data plane, VPCs are isolated from one another, and each volume can only be connected to one VPC.

Access to the individual volumes is controlled by protocol specific (NFS/SMB) access control mechanisms.

In other words, on the network layer, all projects connected to the shared VPC are able to see the volume, while, on the management side, the control plane only allows the owner project to see the volume.

VPC Service Controls

VPC Service Controls establish an access control perimeter around Google Cloud services that are attached to

the internet and are accessible worldwide. These services provide access control through user identities but cannot restrict which network location requests originate from. VPC Service Controls close that gap by introducing the capabilities to restrict access to defined networks.

The Cloud Volumes Service data plane is not connected to the external internet but to private VPCs with well-defined network boundaries (perimeters). Within that network, each volume uses protocol-specific access control. Any external network connectivity is explicitly created by Google Cloud project administrators. The control plane, however, does not provide the same protections as the data plane and can be accessed by anyone from anywhere with valid credentials ([JWT tokens](#)).

In short, the Cloud Volumes Service data plane provides the capability of network access control, without the requirement to support VPC Service Controls and does not explicitly use VPC Service Controls.

Packet sniffing/trace considerations

Packet captures can be useful for troubleshooting network issues or other problems (such as NAS permissions, LDAP connectivity, and so on), but can also be used maliciously to gain information about network IP addresses, MAC addresses, user and group names, and what level of security is being used on endpoints. Because of the way Google Cloud networking, VPCs, and firewall rules are configured, unwanted access to network packets should be difficult to obtain without user login credentials or [JWT tokens](#) into the cloud instances. Packet captures are only possible on endpoints (such as virtual machines (VMs)) and only possible on endpoints internal to the VPC unless a shared VPC and/or external network tunnel/IP forwarding is in use to explicitly allow external traffic to endpoints. There is no way to sniff traffic outside of the clients.

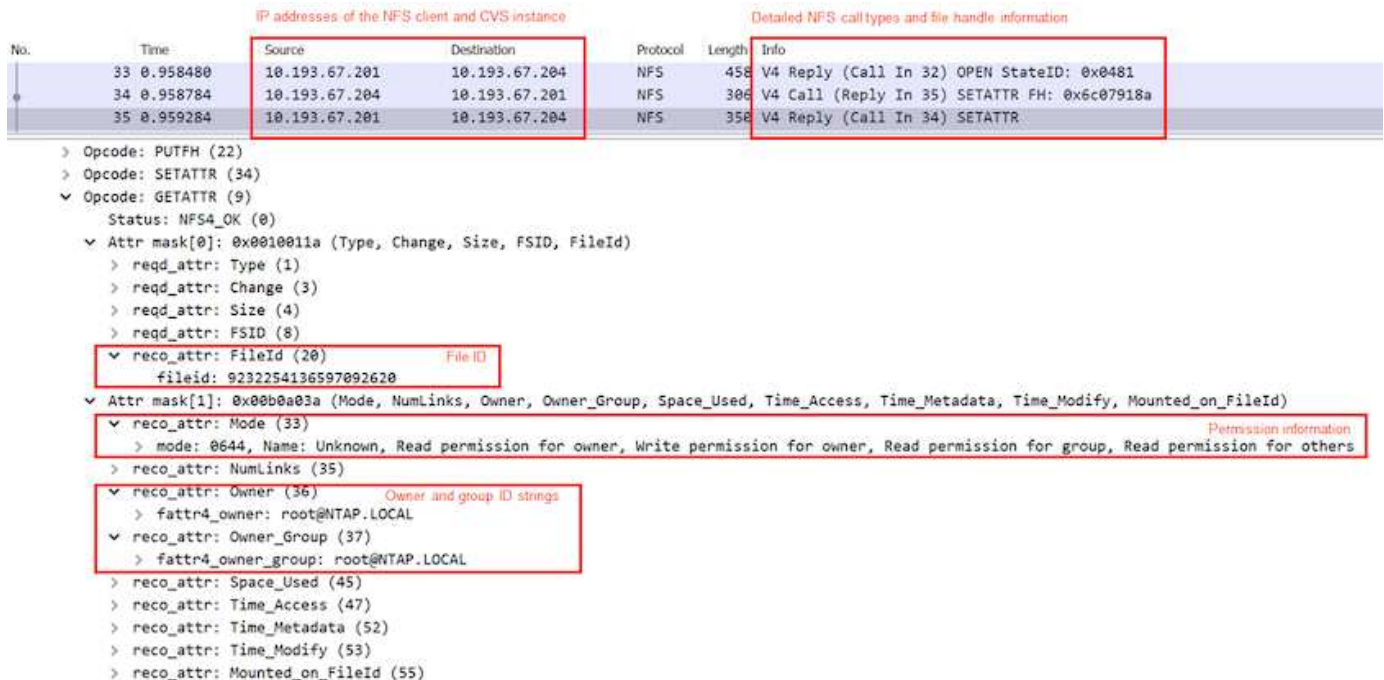

When shared VPCs are used, in-flight encryption with NFS Kerberos and/or [SMB encryption](#) can mask much of the information gleaned from traces. However, some traffic is still sent in plaintext, such as [DNS](#) and [LDAP queries](#). The following figure shows a packet capture from a plaintext LDAP query originating from Cloud Volumes Service and the potential identifying information that is exposed. LDAP queries in Cloud Volumes Service currently do not support encryption or LDAP over SSL. CVS-Performance support LDAP signing, if requested by Active Directory. CVS-SW does not support LDAP signing.





unixUserPassword is queried by LDAP and is not sent in plaintext but instead in a salted hash. By default, Windows LDAP does not populate the unixUserPassword fields. This field is only required if you need to leverage Windows LDAP for interactive logins through LDAP to clients. Cloud Volumes Service does not support interactive LDAP logins to the instances.

The following figure shows a packet capture from an NFS Kerberos conversation next to a capture of NFS over AUTH_SYS. Note how the information available in a trace differs between the two and how enabling in-flight encryption offers greater overall security for NAS traffic.



VM network interfaces

One trick attackers might attempt is to add a new network interface card (NIC) to a VM in [promiscuous mode](#) (port mirroring) or enable promiscuous mode on an existing NIC in order to sniff all traffic. In Google Cloud, adding a new NIC requires a VM to be shut down entirely, which creates alerts, so attackers cannot do this

unnoticed.

In addition, NICs cannot be set to promiscuous mode at all and will trigger alerts in Google Cloud.

Control plane architecture

All management actions to Cloud Volumes Service are done through API. Cloud Volumes Service management integrated into the GCP Cloud Console also uses the Cloud Volumes Service API.

Identity and Access Management

Identity and Access Management ([IAM](#)) is a standard service that enables you to control authentication (logins) and authorization (permissions) to Google Cloud project instances. Google IAM provides a full audit trail of permissions authorization and removal. Currently Cloud Volumes Service does not provide control plane auditing.

Authorization/permission overview

IAM offers built-in, granular permissions for Cloud Volumes Service. You can find a [complete list of granular permissions here](#).

IAM also offers two predefined roles called `netappcloudvolumes.admin` and `netappcloudvolumes.viewer`. These roles can be assigned to specific users or service accounts.

Assign appropriate roles and permission to allow IAM users to manage Cloud Volumes Service.

Examples for using granular permissions include the following:

- Build a custom role with only `get/list/create/update` permissions so that users cannot delete volumes.
- Use a custom role with only `snapshot.*` permissions to create a service account that is used to build application- consistent Snapshot integration.
- Build a custom role to delegate `volumereplication.*` to specific users.

Service accounts

To make Cloud Volumes Service API calls through scripts or [Terraform](#), you must create a service account with the `roles/netappcloudvolumes.admin` role. You can use this service account to generate the JWT tokens required to authenticate Cloud Volumes Service API requests in two different ways:

- Generate a JSON key and use Google APIs to derive a JWT token from it. This is the simplest approach, but it involves manual secrets (the JSON key) management.
- Use [Service account impersonation](#) with `roles/iam.serviceAccountTokenCreator`. The code (script, Terraform, and so on.) runs with [Application Default Credentials](#) and impersonates the service account to gain its permissions. This approach reflects Google security best practices.

See [Creating your service account and private key](#) in the Google cloud documentation for more information.

Cloud Volumes Service API

Cloud Volumes Service API uses a REST-based API by using HTTPS (TLSv1.2) as the underlying network transport. You can find the latest API definition [here](#) and information about how to use the API at [Cloud Volumes APIs in the Google cloud documentation](#).

The API endpoint is operated and secured by NetApp using standard HTTPS (TLSv1.2) functionality.

JWT tokens

Authentication to the API is performed with JWT bearer tokens ([RFC-7519](#)). Valid JWT tokens must be obtained by using Google Cloud IAM authentication. This must be done by fetching a token from IAM by providing a service account JSON key.

Audit logging

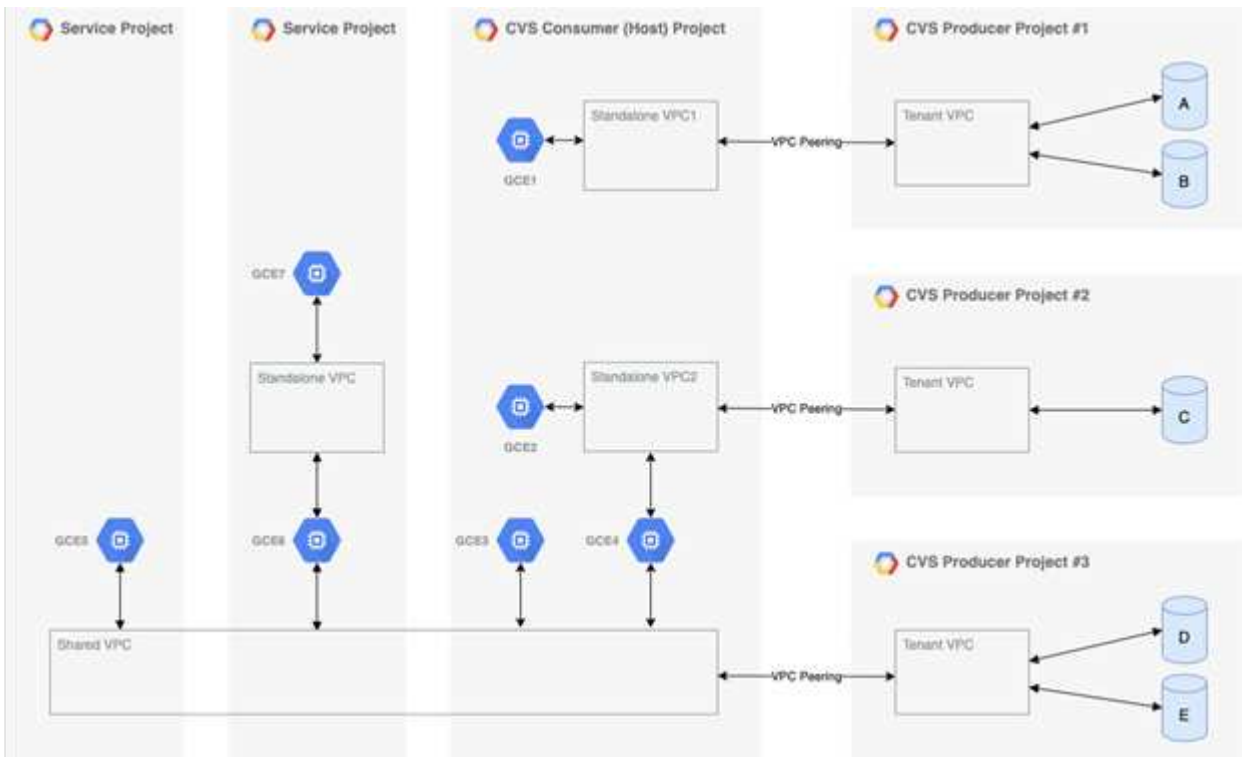
Currently, no user-accessible control plane audit logs are available.

Data plane architecture

Cloud Volumes Service for Google Cloud leverages the Google Cloud [private services access](#) framework. In this framework, users can connect to the Cloud Volumes Service. This framework uses Service Networking and VPC peering constructs like other Google Cloud services, ensuring complete isolation between tenants.

For an architecture overview of Cloud Volumes Service for Google Cloud, see [Architecture for Cloud Volumes Service](#).

User VPCs (standalone or shared) are peered to VPCs within Cloud Volumes Service managed tenant projects, which hosts the volumes.



The preceding figure shows a project (the CVS consumer project in the middle) with three VPC networks connected to Cloud Volumes Service and multiple Compute Engine VMs (GCE1-7) sharing volumes:

- VPC1 allows GCE1 to access volumes A and B.
- VPC2 allows GCE2 and GCE4 to access volume C.

- The third VPC network is a shared VPC, shared with two service projects. It allows GCE3, GCE4, GCE5, and GCE6 to access volumes D and E. Shared VPC networks are only supported for volumes of the CVS-Performance service type.



GCE7 cannot access any volume.

Data can be encrypted both in-transit (using Kerberos and/or SMB encryption) and at-rest in Cloud Volumes Service.

Data encryption in transit

Data in transit can be encrypted at the NAS protocol layer, and the Google Cloud network itself is encrypted, as described in the following sections.

Google Cloud network

Google Cloud encrypts traffic on the network level as described in [Encryption in transit](#) in the Google documentation. As mentioned in the section “Cloud Volumes Services architecture,” Cloud Volumes Service is delivered out of a NetApp-controlled PSA producer project.

In case of CVS-SW, the producer tenant runs Google VMs to provide the service. Traffic between user VMs and Cloud Volumes Service VMs is encrypted automatically by Google.

Although the data path for CVS-Performance isn't fully encrypted on the network layer, NetApp and Google use a combination of [IEEE 802.1AE encryption \(MACSec\)](#), [encapsulation](#) (data encryption), and physically restricted networks to protect data in transit between the Cloud Volumes Service CVS-Performance service type and Google Cloud.

NAS protocols

NFS and SMB NAS protocols provide optional transport encryption at the protocol layer.

SMB encryption

[SMB encryption](#) provides end-to-end encryption of SMB data and protects data from eavesdropping occurrences on untrusted networks. You can enable encryption for both the client/server data connection (only available to SMB3.x capable clients) and the server/domain controller authentication.

When SMB encryption is enabled, clients that do not support encryption cannot access the share.

Cloud Volumes Service supports RC4-HMAC, AES-128-CTS-HMAC-SHA1, and AES-256-CTS-HMAC-SHA1 security ciphers for SMB encryption. SMB negotiates to the highest supported encryption type by the server.

NFSv4.1 Kerberos

For NFSv4.1, CVS-Performance offers Kerberos authentication as described in [RFC7530](#). You can enable Kerberos on a per-volume basis.

The current strongest available encryption type for Kerberos is AES-256-CTS-HMAC-SHA1. NetApp Cloud Volumes Service supports AES-256-CTS-HMAC-SHA1, AES-128-CTS-HMAC-SHA1, DES3, and DES for NFS. It also supports ARCFOUR-HMAC (RC4) for CIFS/SMB traffic, but not for NFS.

Kerberos provides three different security levels for NFS mounts that offer choices for how strong the Kerberos security should be.

As per RedHat's [Common Mount Options](#) documentation:

```
sec=krb5 uses Kerberos V5 instead of local UNIX UIDs and GIDs to
authenticate users.
sec=krb5i uses Kerberos V5 for user authentication and performs integrity
checking of NFS operations using secure checksums to prevent data
tampering.
sec=krb5p uses Kerberos V5 for user authentication, integrity checking,
and encrypts NFS traffic to prevent traffic sniffing. This is the most
secure setting, but it also involves the most performance overhead.
```

As a general rule, the more the Kerberos security level has to do, the worse the performance is, as the client and server spend time encrypting and decrypting NFS operations for each packet sent. Many clients and NFS servers provide support for AES-NI offloading to the CPUs for a better overall experience, but the performance impact of Kerberos 5p (full end-to-end encryption) is significantly greater than the impact of Kerberos 5 (user authentication).

The following table shows differences in what each level does for security and performance.

Security level	Security	Performance
NFSv3—sys	<ul style="list-style-type: none">• Least secure; plain text with numeric user IDs/group IDs• Able to view UID, GID, client IP addresses, export paths, file names, permissions in packet captures	<ul style="list-style-type: none">• Best for most cases
NFSv4.x—sys	<ul style="list-style-type: none">• More secure than NFSv3 (client IDs, name string/domain string matching) but still plain text• Able to view UID, GID, client IP addresses, name strings, domain IDs, export paths, file names, permissions in packet captures	<ul style="list-style-type: none">• Good for sequential workloads (such as VMs, databases, large files)• Bad with high file count/high metadata (30-50% worse)

Security level	Security	Performance
NFS—krb5	<ul style="list-style-type: none"> • Kerberos encryption for credentials in every NFS packet—wraps UID/GID of users/groups in RPC calls in GSS wrapper • User requesting access to mount needs a valid Kerberos ticket (either through username/password or manual key tab exchange); ticket expires after a specified time period and user must reauthenticate for access • No encryption for NFS operations or ancillary protocols like mount/portmapper/nlm (can see export paths, IP addresses, file handles, permissions, file names, atime/mtime in packet captures) 	<ul style="list-style-type: none"> • Best in most cases for Kerberos; worse than AUTH_SYS
NFS—krb5i	<ul style="list-style-type: none"> • Kerberos encryption for credentials in every NFS packet—wraps UID/GID of users/groups in RPC calls in GSS wrapper • User requesting access to mount needs a valid Kerberos ticket (either via username/password or manual key tab exchange); ticket expires after a specified time period and user must reauthenticate for access • No encryption for NFS operations or ancillary protocols like mount/portmapper/nlm (can see export paths, IP addresses, file handles, permissions, file names, atime/mtime in packet captures) • Kerberos GSS checksum is added to every packet to ensure nothing intercepts the packets. If checksums match, conversation is allowed. 	<ul style="list-style-type: none"> • Better than krb5p because the NFS payload is not encrypted; only added overhead compared to krb5 is the integrity checksum. Performance of krb5i won't be much worse than krb5 but will see some degradation.

Security level	Security	Performance
NFS – krb5p	<ul style="list-style-type: none"> • Kerberos encryption for credentials in every NFS packet—wraps UID/GID of users/groups in RPC calls in GSS wrapper • User requesting access to mount needs a valid Kerberos ticket (either via username/password or manual keytab exchange); ticket expires after specified time period and user must reauthenticate for access • All of the NFS packet payloads are encrypted with the GSS wrapper (cannot see file handles, permissions, file names, atime/mtime in packet captures). • Includes integrity check. • NFS operation type is visible (FSINFO, ACCESS, GETATTR, and so on). • Ancillary protocols (mount, portmap, nlm, and so on) are not encrypted - (can see export paths, IP addresses) 	<ul style="list-style-type: none"> • Worst performance of the security levels; krb5p has to encrypt/decrypt more. • Better performance than krb5p with NFSv4.x for high file count workloads.

In Cloud Volumes Service, a configured Active Directory server is used as Kerberos server and LDAP server (to lookup user identities from an RFC2307 compatible schema). No other Kerberos or LDAP servers are supported. NetApp highly recommends that you use LDAP for identity management in Cloud Volumes Service. For information on how NFS Kerberos is shown in packet captures, see the section [“Packet sniffing/trace considerations.”](#)

Data encryption at rest

All volumes in Cloud Volumes Service are encrypted-at-rest using AES-256 encryption, which means all user data written to media is encrypted and can only be decrypted with a per-volume key.

- For CVS-SW, Google-generated keys are used.
- For CVS-Performance, the per-volume keys are stored in a key manager built into the Cloud Volumes Service.

Starting in November 2021, preview customer-managed encryption keys (CMEK) functionality was made available. This enables you to encrypt the per-volume keys with a per-project, per-region master key that is hosted in [Google Key Management Service \(KMS\)](#). KMS enables you to attach external key managers.

For information about configuring KMS for CVS-Performance, see [Setting up customer-managed encryption keys](#).

Firewall

Cloud Volumes Service exposes multiple TCP ports to serve NFS and SMB shares:

- [Ports required for NFS access](#)
- [Ports required for SMB access](#)

Additionally, SMB, NFS with LDAP including Kerberos, and dual-protocol configurations require access to a Windows Active Directory domain. Active Directory connections must be [configured](#) on a per-region basis. Active Directory Domain controllers (DC) are identified by using [DNS-based DC discovery](#) using the specified DNS servers. Any of the DCs returned are used. The list of eligible DCs can be limited by specifying an Active Directory site.

Cloud Volumes Service reaches out with IP addresses from the CIDR range allocated with the `gcloud compute address` command while [on-boarding the Cloud Volumes Service](#). You can use this CIDR as source addresses to configure inbound firewalls to your Active Directory domain controllers.

Active Directory Domain Controllers must [expose ports to the Cloud Volumes Service CIDRs as mentioned here](#).

NAS protocols

NAS protocols overview

NAS protocols include NFS (v3 and v4.1) and SMB/CIFS (2.x and 3.x). These protocols are how CVS allows shared access to data across multiple NAS clients. In addition, Cloud Volumes Service can provide access to NFS and SMB/CIFS clients simultaneously (dual-protocol) while honoring all of the identity and permission settings on files and folders in the NAS shares. To maintain the highest possible data transfer security, Cloud Volumes Service supports protocol encryption in flight using SMB encryption and NFS Kerberos 5p.



Dual-protocol is available with CVS-Performance only.

Basics of NAS protocols

NAS protocols are ways for multiple clients on a network to access the same data on a storage system, such as Cloud Volumes Service on GCP. NFS and SMB are the defined NAS protocols and operate on a client/server basis where Cloud Volumes Service acts as the server. Clients send access, read, and write requests to the server, and the server is responsible for coordinating the locking mechanisms for files, storing permissions and handling identity and authentication requests.

For example, the following general process is followed if a NAS client wants to create a new file in a folder.

1. The client asks the server for information about the directory (permissions, owner, group, file ID, available space, and so on); the server responds with the information if the requesting client and user have the

necessary permissions on the parent folder.

2. If the permissions on the directory allow access, the client then asks the server if the file name being created already exists in the file system. If the file name is already in use, creation fails. If the file name does not exist, the server lets the client know it can proceed.
3. The client issues a call to the server to create the file with the directory handle and file name and sets the access and modified times. The server issues a unique file ID to the file to make sure that no other files are created with the same file ID.
4. The client sends a call to check file attributes before the WRITE operation. If permissions allow it, the client then writes the new file. If locking is used by the protocol/application, the client asks the server for a lock to prevent other clients from accessing the file while locked to prevent data corruption.

NFS

NFS is a distributed file system protocol that is an open IETF standard defined in Request for Comments (RFC) that allows anyone to implement the protocol.

Volumes in Cloud Volumes Service are shared out to NFS clients by exporting a path that is accessible to a client or set of clients. Permissions to mount these exports are defined by export policies and rules, which are configurable by Cloud Volumes Service administrators.

The NetApp NFS implementation is considered a gold standard for the protocol and is used in countless enterprise NAS environments. The following sections cover NFS and specific security features available in Cloud Volumes Service and how they are implemented.

Default local UNIX users and groups

Cloud Volumes Service contains several default UNIX users and groups for various basic functionalities. These users and groups cannot currently be modified or deleted. New local users and groups cannot currently be added to Cloud Volumes Service. UNIX users and groups outside of the default users and groups need to be provided by an external LDAP name service.

The following table shows the default users and groups and their corresponding numeric IDs. NetApp recommends not creating new users or groups in LDAP or on the local clients that re-use these numeric IDs.

Default users: numeric IDs	Default groups: numeric IDs
<ul style="list-style-type: none">• root:0• pcuser:65534• nobody:65535	<ul style="list-style-type: none">• root:0• daemon:1• pcuser:65534• nobody:65535



When using NFSv4.1, the root user might display as nobody when running directory listing commands on NFS clients. This is due to the client's ID domain mapping configuration. See the section called [NFSv4.1 and the nobody user/group](#) for details on this issue and how to resolve it.

The root user

In Linux, the root account has access to all commands, files, and folders in a Linux-based file system. Because of the power of this account, security best practices often require the root user to be disabled or restricted in some fashion. In NFS exports, the power a root user has over the files and folders can be controlled in Cloud

Volumes Service through export policies and rules and a concept known as root squash.

Root squashing ensures that the root user accessing an NFS mount is squashed to the anonymous numeric user 65534 (see the section “[The anonymous user](#)”) and is currently only available when using CVS-Performance by selecting Off for root access during export policy rule creation. If the root user is squashed to the anonymous user, it no longer has access to run `chown` or [setuid/setgid commands \(the sticky bit\)](#) on files or folders in the NFS mount, and files or folders created by the root user show the anon UID as the owner/group. In addition, NFSv4 ACLs cannot be modified by the root user. However, the root user still has access to `chmod` and deleted files that it does not have explicit permissions for. If you want to limit access to a root user’s file and folder permissions, consider using a volume with NTFS ACLs, creating a Windows user named `root`, and applying the desired permissions to the files or folders.

The anonymous user

The anonymous (anon) user ID specifies a UNIX user ID or username that is mapped to client requests that arrive without valid NFS credentials. This can include the root user when root squashing is used. The anon user in Cloud Volumes Service is 65534.

This UID is normally associated with the username `nobody` or `nfsnobody` in Linux environments. Cloud Volumes Service also uses 65534 as the local UNIX user `pcuser` (see the section “[Default local UNIX users and groups](#)”), which is also the default fallback user for Windows to UNIX name mappings when no valid matching UNIX user can be found in LDAP.

Because of the differences in usernames across Linux and Cloud Volumes Service for UID 65534, the name string for users mapped to 65534 might not match when using NFSv4.1. As a result, you might see `nobody` as the user on some files and folders. See the section “[NFSv4.1 and the nobody user/group](#)” for information about this issue and how to resolve it.

Access control/exports

Initial export/share access for NFS mounts is controlled through host-based export policy rules contained within an export policy. A host IP, host name, subnet, netgroup, or domain is defined to allow access to mount the NFS share and the level of access allowed to the host. Export policy rule configuration options depend on the Cloud Volumes Service level.

For CVS-SW, the following options are available for export-policy configuration:

- **Client match.** Comma-separated list of IP addresses, comma-separated list of hostnames, subnets, netgroups, domain names.
- **RO/RW access rules.** Select read/write or read only to control level of access to export. CVS-Performance provides the following options:
 - **Client match.** Comma-separated list of IP addresses, comma-separated list of hostnames, subnets, netgroups, domain names.
 - **RO/RW access rules.** Select read/write or read only to control level of access to export.
 - **Root access (on/off).** Configures root squash (see the section “[The root user](#)” for details).
 - **Protocol type.** This limits access to the NFS mount to a specific protocol version. When specifying both NFSv3 and NFSv4.1 for the volume, either leave both blank or check both boxes.
- **Kerberos security level (when Enable Kerberos is selected).** Provides the options of `krb5`, `krb5i`, and/or `krb5p` for read-only or read-write access.

Change ownership (chown) and change group (chgrp)

NFS on Cloud Volumes Service only allows the root user to run `chown/chgrp` on files and folders. Other users see an `Operation not permitted` error— even on files they own. If you use `root squash` (as covered in the section “[The root user](#)”), the root is squashed to a nonroot user and is not allowed access to `chown` and `chgrp`. There are currently no workarounds in Cloud Volumes Service to allow `chown` and `chgrp` for non-root users. If ownership changes are required, consider using dual protocol volumes and set the security style to NTFS to control permissions from the Windows side.

Permission management

Cloud Volumes Service supports both mode bits (such as 644, 777, and so on for `rwx`) and NFSv4.1 ACLs to control permissions on NFS clients for volumes that use the UNIX security style. Standard permission management is used for these (such as `chmod`, `chown`, or `nfs4_setfacl`) and work with any Linux client that supports them.

Additionally, when using dual protocol volumes set to NTFS, NFS clients can leverage Cloud Volumes Service name mapping to Windows users, which then are used to resolve the NTFS permissions. This requires an LDAP connection to Cloud Volumes Service to provide numeric-ID-to-username translations because Cloud Volumes Service requires a valid UNIX username to map properly to a Windows username.

Providing granular ACLs for NFSv3

Mode bit permissions cover only owner, group, and everyone else in the semantics—meaning that there are no granular user access controls in place for basic NFSv3. Cloud Volumes Service does not support POSIX ACLs, nor extended attributes (such as `chattr`), so granular ACLs are only possible in the following scenarios with NFSv3:

- NTFS security style volumes (CIFS server required) with valid UNIX to Windows user mappings.
- NFSv4.1 ACLs applied using an admin client mounting NFSv4.1 to apply ACLs.

Both methods require an LDAP connection for UNIX identity management and a valid UNIX user and group information populated (see the section “[LDAP](#)”) and are only available with CVS-Performance instances. To use NTFS security style volumes with NFS, you must use dual-protocol (SMB and NFSv3) or dual-protocol (SMB and NFSv4.1), even if no SMB connections are made. To use NFSv4.1 ACLs with NFSv3 mounts, you must select `Both (NFSv3/NFSv4.1)` as the protocol type.

Regular UNIX mode bits don't provide the same level of granularity in permissions that NTFS or NFSv4.x ACLs provide. The following table compares the permission granularity between NFSv3 mode bits and NFSv4.1 ACLs. For information about NFSv4.1 ACLs, see [nfs4_acl - NFSv4 Access Control Lists](#).

NFSv3 mode bits	NFSv4.1 ACLs
<ul style="list-style-type: none"> • Set user ID on execution • Set group ID on execution • Save swapped text (not defined in POSIX) • Read permission for owner • Write permission for owner • Execute permission for owner on a file; or look up (search) permission for owner in directory • Read permission for group • Write permission for group • Execute permission for group on a file; or look up (search) permission for group in directory • Read permission for others • Write permission for others • Execute permission for others on a file; or look up (search) permission for others in directory 	<p>Access control entry (ACE) types (Allow/Deny/Audit)</p> <ul style="list-style-type: none"> * Inheritance flags * directory-inherit * file-inherit * no-propagate-inherit * inherit-only <p>Permissions</p> <ul style="list-style-type: none"> * read-data (files) / list-directory (directories) * write-data (files) / create-file (directories) * append-data (files) / create-subdirectory (directories) * execute (files) / change-directory (directories) * delete * delete-child * read-attributes * write-attributes * read-named-attributes * write-named-attributes * read-ACL * write-ACL * write-owner * Synchronize

Finally, NFS group membership (in both NFSv3 and NFSv4.x) is limited to a default maximum of 16 for AUTH_SYS as per the RPC packet limits. NFS Kerberos provides up to 32 groups and NFSv4 ACLs remove the limitation by way of granular user and group ACLs (up to 1024 entries per ACE).

Additionally, Cloud Volumes Service provides extended group support to extend the maximum supported groups up to 32. This requires an LDAP connection to an LDAP server that contains valid UNIX user and group identities. For more information about configuring this, see [Creating and managing NFS volumes](#) in the Google documentation.

NFSv3 user and group IDs

NFSv3 user and group IDs come across the wire as numeric IDs rather than names. Cloud Volumes Service does no username resolution for these numeric IDs with NFSv3, with UNIX security style volumes using just mode bits. When NFSv4.1 ACLs are present, a numeric ID lookup and/or name string lookup is needed to resolve the ACL properly—even when using NFSv3. With NTFS security style volumes, Cloud Volumes Service must resolve a numeric ID to a valid UNIX user and then map to a valid Windows user to negotiate access rights.

Security limitations of NFSv3 user and group IDs

With NFSv3, the client and server never have to confirm that the user attempting a read or write with a numeric ID is a valid user; it is just implicitly trusted. This opens the file system up to potential breaches simply by spoofing any numeric ID. To prevent security holes like this, there are a few options available to Cloud Volumes Service.

- Implementing Kerberos for NFS forces users to authenticate with a username and password or keytab file to get a Kerberos ticket to allow access into a mount. Kerberos is available with CVS-Performance instances and only with NFSv4.1.

- Limiting the list of hosts in your export policy rules limits which NFSv3 clients have access to the Cloud Volumes Service volume.
- Using dual-protocol volumes and applying NTFS ACLs to the volume forces NFSv3 clients to resolve numeric IDs to valid UNIX usernames to authenticate properly to access mounts. This requires enabling LDAP and configuring UNIX user and group identities.
- Squashing the root user limits the damage a root user can do to an NFS mount but does not completely remove risk. For more information, see the section “[The root user.](#)”

Ultimately, NFS security is limited to what the protocol version you are using offers. NFSv3, while more performant in general than NFSv4.1, does not provide the same level of security.

NFSv4.1

NFSv4.1 provides greater security and reliability as compared to NFSv3, for the following reasons:

- Integrated locking through a lease-based mechanism
- Stateful sessions
- All NFS functionality over a single port (2049)
- TCP only
- ID domain mapping
- Kerberos integration (NFSv3 can use Kerberos, but only for NFS, not for ancillary protocols such as NLM)

NFSv4.1 dependencies

Because of the additional security features in NFSv4.1, there are some external dependencies involved that were not needed to use NFSv3 (similar to how SMB requires dependencies such as Active Directory).

NFSv4.1 ACLs

Cloud Volumes Service offers support for NFSv4.x ACLs, which deliver distinct advantages over normal POSIX-style permissions, such as the following:

- Granular control of user access to files and directories
 - Better NFS security
 - Improved interoperability with CIFS/SMB
 - Removal of the NFS limitation of 16 groups per user with AUTH_SYS security
 - ACLs bypass the need for group ID (GID) resolution, which effectively removes the GID limit
- NFSv4.1 ACLs are controlled from NFS clients—not from Cloud Volumes Service. To use NFSv4.1 ACLs, be sure your client’s software version supports them and the proper NFS utilities are installed.

Compatibility between NFSv4.1 ACLs and SMB clients

NFSv4 ACLs are different from Windows file-level ACLs (NTFS ACLs) but carry similar functionality. However, in multiprotocol NAS environments, if NFSv4.1 ACLs are present and you are using dual-protocol access (NFS and SMB on the same datasets), clients using SMB2.0 and later won’t be able to view or manage ACLs from Windows security tabs.

How NFSv4.1 ACLs work

For reference, the following terms are defined:

- **Access control list (ACL).** A list of permissions entries.
- **Access control entry (ACE).** A permission entry in the list.

When a client sets an NFSv4.1 ACL on a file during a SETATTR operation, Cloud Volumes Service sets that ACL on the object, replacing any existing ACL. If there is no ACL on a file, then the mode permissions on the file are calculated from OWNER@, GROUP@, and EVERYONE@. If there are any existing SUID/SGID/STICKY bits on the file, they are not affected.

When a client gets an NFSv4.1 ACL on a file during the course of a GETATTR operation, Cloud Volumes Service reads the NFSv4.1 ACL associated with the object, constructs a list of ACEs, and returns the list to the client. If the file has an NT ACL or mode bits, then an ACL is constructed from mode bits and is returned to the client.

Access is denied if a DENY ACE is present in the ACL; access is granted if an ALLOW ACE exists. However, access is also denied if neither of the ACEs is present in the ACL.

A security descriptor consists of a security ACL (SACL) and a discretionary ACL (DACL). When NFSv4.1 interoperates with CIFS/SMB, the DACL is one-to-one mapped with NFSv4 and CIFS. The DACL consists of the ALLOW and the DENY ACEs.

If a basic `chmod` is run on a file or folder with NFSv4.1 ACLs set, existing user and group ACLs are preserved, but the default OWNER@, GROUP@, EVERYONE@ ACLs are modified.

A client using NFSv4.1 ACLs can set and view ACLs for files and directories on the system. When a new file or subdirectory is created in a directory that has an ACL, that object inherits all ACEs in the ACL that have been tagged with the appropriate [inheritance flags](#).

If a file or directory has an NFSv4.1 ACL, that ACL is used to control access no matter which protocol is used to access the file or directory.

Files and directories inherit ACEs from NFSv4 ACLs on parent directories (possibly with appropriate modifications) as long as the ACEs have been tagged with the correct inheritance flags.

When a file or directory is created as the result of an NFSv4 request, the ACL on the resulting file or directory depends on whether the file creation request includes an ACL or only standard UNIX file access permissions. The ACL also depends on whether the parent directory has an ACL.

- If the request includes an ACL, that ACL is used.
- If the request includes only standard UNIX file access permissions and the parent directory does not have an ACL, the client file mode is used to set standard UNIX file access permissions.
- If the request includes only standard UNIX file access permissions and the parent directory has a noninheritable ACL, a default ACL based on the mode bits passed into the request is set on the new object.
- If the request includes only standard UNIX file access permissions but the parent directory has an ACL, the ACEs in the parent directory's ACL are inherited by the new file or directory as long as the ACEs have been tagged with the appropriate inheritance flags.

ACE permissions

NFSv4.1 ACLs permissions uses a series of upper- and lower-case letter values (such as `rxwtncy`) to control access. For more information about these letter values, see [HOW TO: Use NFSv4 ACL](#).

NFSv4.1 ACL behavior with umask and ACL inheritance

[NFSv4 ACLs provide the ability to offer ACL inheritance](#). ACL inheritance means that files or folders created beneath objects with NFSv4.1 ACLs set can inherit the ACLs based on the configuration of the [ACL inheritance flag](#).

[Umask](#) is used to control the permission level at which files and folders are created in a directory without administrator interaction. By default, Cloud Volumes Service allows umask to override inherited ACLs, which is expected behavior as per [RFC 5661](#).

ACL formatting

NFSv4.1 ACLs have specific formatting. The following example is an ACE set on a file:

```
A::ldapuser@domain.netapp.com:rwatTnNcCy
```

The preceding example follows the ACL format guidelines of:

```
type:flags:principal:permissions
```

A type of `A` means “allow.” The inherit flags are not set in this case, because the principal is not a group and does not include inheritance. Also, because the ACE is not an AUDIT entry, there is no need to set the audit flags. For more information about NFSv4.1 ACLs, see http://linux.die.net/man/5/nfs4_acl.

If the NFSv4.1 ACL is not set properly (or a name string cannot be resolved by the client and server), the ACL might not behave as expected, or the ACL change might fail to apply and throw an error.

Sample errors include:

```
Failed setattr operation: Invalid argument
Scanning ACE string 'A:: user@rwaDxtTnNcCy' failed.
```

Explicit DENY

NFSv4.1 permissions can include explicit DENY attributes for OWNER, GROUP, and EVERYONE. That is because NFSv4.1 ACLs are default-deny, which means that if an ACL is not explicitly granted by an ACE, then it is denied. Explicit DENY attributes override any ACCESS ACEs, explicit or not.

DENY ACEs are set with an attribute tag of `D`.

In the example below, `GROUP@` is allowed all read and execute permissions, but denied all write access.

```
sh-4.1$ nfs4_getfacl /mixed
A::ldapuser@domain.netapp.com:ratTnNcCy
A::OWNER@:rwaDxtTnNcCy
D::OWNER@:
A:g:GROUP@:rxtncy
D:g:GROUP@:waDTC
A::EVERYONE@:rxtncy
D::EVERYONE@:waDTC
```

DENY ACEs should be avoided whenever possible because they can be confusing and complicated; ALLOW ACLs that are not explicitly defined are implicitly denied. When DENY ACEs are set, users might be denied access when they expect to be granted access.

The preceding set of ACEs is equivalent to 755 in mode bits, which means:

- The owner has full rights.
- Groups have read only.
- Others have read only.

However, even if permissions are adjusted to the 775 equivalent, access can be denied because of the explicit DENY set on EVERYONE.

NFSv4.1 ID domain mapping dependencies

NFSv4.1 leverages ID domain mapping logic as a security layer to help verify that a user attempting access to an NFSv4.1 mount is indeed who they claim to be. In these cases, the username and group name coming from the NFSv4.1 client appends a name string and sends it to the Cloud Volumes Service instance. If that username/group name and ID string combination does not match, then the user and/or group is squashed to the default nobody user specified in the `/etc/idmapd.conf` file on the client.

This ID string is a requirement for proper permission adherence, especially when NFSv4.1 ACLs and/or Kerberos are in use. As a result, name service server dependencies such as LDAP servers are necessary to ensure consistency across clients and Cloud Volumes Service for proper user and group name identity resolution.

Cloud Volumes Service uses a static default ID domain name value of `defaultv4iddomain.com`. NFS clients default to the DNS domain name for its ID domain name settings, but you can manually adjust the ID domain name in `/etc/idmapd.conf`.

If LDAP is enabled in Cloud Volumes Service, then Cloud Volumes Service automates the NFS ID domain to change to what is configured for the search domain in DNS and clients won't need to be modified unless they use different DNS domain search names.

When Cloud Volumes Service can resolve a username or group name in local files or LDAP, the domain string is used and non-matching domain IDs squash to nobody. If Cloud Volumes Service cannot find a username or group name in local files or LDAP, the numeric ID value is used and the NFS client resolves the name properly (this is similar to NFSv3 behavior).

Without changing the client's NFSv4.1 ID domain to match what the Cloud Volumes Service volume is using, you see the following behavior:

- UNIX users and groups with local entries in Cloud Volumes Service (such as root, as defined in local UNIX users and groups) are squashed to the nobody value.
- UNIX users and groups with entries in LDAP (if Cloud Volumes Service is configured to use LDAP) squashes to nobody if DNS domains are different between NFS clients and Cloud Volumes Service.
- UNIX users and groups with no local entries or LDAP entries use the numeric ID value and resolve to the name specified on the NFS client. If no name exists on the client, only the numeric ID is shown.

The following shows the results of the preceding scenario:

```
# ls -la /mnt/home/prof1/nfs4/
total 8
drwxr-xr-x 2 nobody nobody 4096 Feb  3 12:07 .
drwxrwxrwx 7 root    root    4096 Feb  3 12:06 ..
-rw-r--r-- 1  9835   9835     0 Feb  3 12:07 client-user-no-name
-rw-r--r-- 1 nobody nobody     0 Feb  3 12:07 ldap-user-file
-rw-r--r-- 1 nobody nobody     0 Feb  3 12:06 root-user-file
```

When the client and server ID domains match, this is how the same file listing looks:

```
# ls -la
total 8
drwxr-xr-x 2 root    root    4096 Feb  3 12:07 .
drwxrwxrwx 7 root    root    4096 Feb  3 12:06 ..
-rw-r--r-- 1  9835   9835     0 Feb  3 12:07 client-user-no-name
-rw-r--r-- 1 apache apache-group 0 Feb  3 12:07 ldap-user-file
-rw-r--r-- 1 root    root    0 Feb  3 12:06 root-user-file
```

For more information about this issue and how to resolve it, see the section [“NFSv4.1 and the nobody user/group.”](#)

Kerberos dependencies

If you plan to use Kerberos with NFS, you must have the following with Cloud Volumes Service:

- Active Directory domain for Kerberos Distribution Center services (KDC)
- Active Directory domain with user and group attributes populated with UNIX information for LDAP functionality (NFS Kerberos in Cloud Volumes Service requires a user SPN to UNIX user mapping for proper functionality.)
- LDAP enabled on the Cloud Volumes Service instance
- Active Directory domain for DNS services

NFSv4.1 and the nobody user/group

One of the most common issues seen with an NFSv4.1 configuration is when a file or folder is shown in a listing using `ls` as being owned by the `user:group` combination of `nobody:nobody`.

For example:

```
sh-4.2$ ls -la | grep prof1-file
-rw-r--r-- 1 nobody nobody    0 Apr 24 13:25 prof1-file
```

And the numeric ID is 99.

```
sh-4.2$ ls -lan | grep prof1-file
-rw-r--r-- 1 99 99    0 Apr 24 13:25 prof1-file
```

In some instances, the file might show the correct owner but `nobody` as the group.

```
sh-4.2$ ls -la | grep newfile1
-rw-r--r-- 1 prof1 nobody    0 Oct  9 2019 newfile1
```

Who is `nobody`?

The `nobody` user in NFSv4.1 is different from the `nfsnobody` user. You can view how an NFS client sees each user by running the `id` command:

```
# id nobody
uid=99(nobody) gid=99(nobody) groups=99(nobody)
# id nfsnobody
uid=65534(nfsnobody) gid=65534(nfsnobody) groups=65534(nfsnobody)
```

With NFSv4.1, the `nobody` user is the default user defined by the `idmapd.conf` file and can be defined as any user you want to use.

```
# cat /etc/idmapd.conf | grep nobody
#Nobody-User = nobody
#Nobody-Group = nobody
```

Why does this happen?

Because security through name string mapping is a key tenet of NFSv4.1 operations, the default behavior when a name string does not match properly is to squash that user to one that won't normally have any access to files and folders owned by users and groups.

When you see `nobody` for the user and/or group in file listings, this generally means something in NFSv4.1 is misconfigured. Case sensitivity can come into play here.

For example, if `user1@CVSDemo.local` (uid 1234, gid 1234) is accessing an export, then Cloud Volumes Service must be able to find `user1@CVSDemo.local` (uid 1234, gid 1234). If the user in Cloud Volumes Service is `USER1@CVSDemo.local`, then it won't match (uppercase `USER1` versus lowercase `user1`). In

many cases, you can see the following in the messages file on the client:

```
May 19 13:14:29 centos7 nfsidmap[17481]: nss_getpwnam: name
'root@defaultv4iddomain.com' does not map into domain 'CVSDemo.LOCAL'
May 19 13:15:05 centos7 nfsidmap[17534]: nss_getpwnam: name 'nobody' does
not map into domain 'CVSDemo.LOCAL'
```

The client and server must both agree that a user is indeed who they are claiming to be, so you must check the following to ensure that the user that the client sees has the same information as the user that Cloud Volumes Service sees.

- **NFSv4.x ID domain.** Client: `idmapd.conf` file; Cloud Volumes Service uses `defaultv4iddomain.com` and cannot be changed manually. If using LDAP with NFSv4.1, Cloud Volumes Service changes the ID domain to what the DNS search domain is using, which is the same as the AD domain.
- **User name and numeric IDs.** This determines where the client is looking for user names and leverages the name service switch configuration—client: `nsswitch.conf` and/or local `passwd` and `group` files; Cloud Volumes Service does not allow modifications to this but automatically adds LDAP to the configuration when it is enabled.
- **Group name and numeric IDs.** This determines where the client is looking for group names and leverages the name service switch configuration—client: `nsswitch.conf` and/or local `passwd` and `group` files; Cloud Volumes Service does not allow modifications to this but automatically adds LDAP to the configuration when it is enabled.

In almost all cases, if you see `nobody` in user and group listings from clients, the issue is user or group name domain ID translation between Cloud Volumes Service and the NFS client. To avoid this scenario, use LDAP to resolve user and group information between clients and Cloud Volumes Service.

Viewing name ID strings for NFSv4.1 on clients

If you are using NFSv4.1, there is a name-string mapping that takes place during NFS operations, as previously described.

In addition to using `/var/log/messages` to find an issue with NFSv4 IDs, you can use the `nfsidmap -l` command on the NFS client to view which usernames have properly mapped to the NFSv4 domain.

For example, this is output of the command after a user that can be found by the client and Cloud Volumes Service accesses an NFSv4.x mount:

```
# nfsidmap -l
4 .id_resolver keys found:
gid:daemon@CVSDemo.LOCAL
uid:nfs4@CVSDemo.LOCAL
gid:root@CVSDemo.LOCAL
uid:root@CVSDemo.LOCAL
```

When a user that does not map properly into the NFSv4.1 ID domain (in this case, `netapp-user`) tries to access the same mount and touches a file, they are assigned `nobody:nobody`, as expected.

```

# su netapp-user
sh-4.2$ id
uid=482600012(netapp-user), 2000(secondary)
sh-4.2$ cd /mnt/nfs4/
sh-4.2$ touch newfile
sh-4.2$ ls -la
total 16
drwxrwxrwx  5 root  root  4096 Jan 14 17:13 .
drwxr-xr-x.  8 root  root    81 Jan 14 10:02 ..
-rw-r--r--  1 nobody nobody    0 Jan 14 17:13 newfile
drwxrwxrwx  2 root  root  4096 Jan 13 13:20 qtrees1
drwxrwxrwx  2 root  root  4096 Jan 13 13:13 qtrees2
drwxr-xr-x  2 nfs4  daemon 4096 Jan 11 14:30 testdir

```

The `nfsidmap -l` output shows the user `pcuser` in the display but not `netapp-user`; this is the anonymous user in our `export-policy` rule (65534).

```

# nfsidmap -l
6 .id_resolver keys found:
gid:pcuser@CVSDemo.LOCAL
uid:pcuser@CVSDemo.LOCAL
gid:daemon@CVSDemo.LOCAL
uid:nfs4@CVSDemo.LOCAL
gid:root@CVSDemo.LOCAL
uid:root@CVSDemo.LOCAL

```

SMB

SMB is a network file sharing protocol developed by Microsoft that provides centralized user/group authentication, permissions, locking, and file sharing to multiple SMB clients over an Ethernet network. Files and folders are presented to clients by way of shares, which can be configured with a variety of share properties and offers access control through share-level permissions. SMB can be presented to any client that offers support for the protocol, including Windows, Apple, and Linux clients.

Cloud Volumes Service provides support for the SMB 2.1 and 3.x versions of the protocol.

Access control/SMB shares

- When a Windows username requests access to the Cloud Volumes Service volume, Cloud Volumes Service looks for a UNIX username using the methods configured by Cloud Volumes Service administrators.
- If an external UNIX identity provider (LDAP) is configured and Windows/UNIX usernames are identical, then Windows usernames will map 1:1 to UNIX usernames without any additional configuration needed. When LDAP is enabled, Active Directory is used to host those UNIX attributes for user and group objects.

- If Windows names and UNIX names do not match identically, then LDAP must be configured to allow Cloud Volumes Service to use the LDAP name mapping configuration (see the section [“Using LDAP for asymmetric name mapping”](#)).
- If LDAP is not in use, then Windows SMB users map to a default local UNIX user named `pcuser` in Cloud Volumes Service. This means files written in Windows by users that map to the `pcuser` show UNIX ownership as `pcuser` in multiprotocol NAS environments. `pcuser` here is effectively the `nobody` user in Linux environments (UID 65534).

In deployments with SMB only, the `pcuser` mapping still occurs, but it won't matter, because Windows user and group ownership is correctly displayed and NFS access to the SMB-only volume is not allowed. In addition, SMB-only volumes do not support conversion to NFS or dual-protocol volumes after they are created.

Windows leverages Kerberos for username authentication with the Active Directory domain controllers, which requires a username/password exchange with the AD DCs, which is external to the Cloud Volumes Service instance. Kerberos authentication is used when the `\\SERVERNAME` UNC path is used by the SMB clients and the following is true:

- DNS A/AAAA entry exists for SERVERNAME
- A valid SPN for SMB/CIFS access exists for SERVERNAME

When a Cloud Volumes Service SMB volume is created, the machine account name is created as defined in the section [“How Cloud Volumes Service shows up in Active Directory.”](#) That machine account name also becomes the SMB share access path because Cloud Volumes Service leverages Dynamic DNS (DDNS) to create the necessary A/AAAA and PTR entries in DNS and the necessary SPN entries on the machine account principal.



For PTR entries to be created, the reverse lookup zone for the Cloud Volumes Service instance IP address must exist on the DNS server.

For example, this Cloud Volumes Service volume uses the following UNC share path: `\\cvs-east-433d.cvsdemo.local`.

In Active Directory, these are the Cloud Volumes Service-generated SPN entries:

```
PS C:\> setspn /L CVS-EAST-433D
Registered ServicePrincipalNames for CN=CVS-EAST-433D,CN=Computers,DC=cvsdemo,DC=local:
HOST/cvs-east-433d.cvsdemo.local
HOST/ CVS-EAST-433D
```

This is the DNS forward/reverse lookup result:


```
PS C:\> nslookup CVS-EAST-433D
Server:    activedirectory.region.lab.internal
Address:  10. xx.0. xx
Name:     CVS-EAST-433D.cvsdemo.local
Address:  10. xxx.0. x
PS C:\> nslookup 10. xxx.0. x
Server:    activedirectory.region.lab.internal
Address:  10.xx.0.xx
Name:     CVS-EAST-433D.CVSDEMO.LOCAL
Address:  10. xxx.0. x
```

Optionally, more access control can be applied by enabling/requiring SMB encryption for SMB shares in Cloud Volumes Service. If SMB encryption isn't supported by one of the endpoints, then access is not allowed.

Using SMB name aliases

In some cases, it might be a security concern for end users to know the machine account name in use for Cloud Volumes Service. In other cases, you might simply want to provide a simpler access path to your end users. In those cases, you can create SMB aliases.

If you want to create aliases for the SMB share path, you can leverage what is known as a CNAME record in DNS. For example, if you want to use the name `\\CIFs` to access shares instead of `\\cvs-east-433d.cvsdemo.local`, but you still want to use Kerberos authentication, a CNAME in DNS that points to the existing A/AAAA record and an additional SPN added to the existing machine account provides Kerberos access.



This is the resulting DNS forward lookup result after adding a CNAME:

```
PS C:\> nslookup cifs
Server: ok-activedirectory.us-east4-a.c.cv-solution-architect-
lab.internal
Address: 10. xx.0. xx
Name: CVS-EAST-433D.cvsdemo.local
Address: 10. xxx.0. x
Aliases: cifs.cvsdemo.local
```

This is the resulting SPN query after adding new SPNs:

```
PS C:\> setspn /L CVS-EAST-433D
Registered ServicePrincipalNames for CN=CVS-EAST-433D,CN=Computers,DC=cvsdemo,DC=local:
cifs/cifs.cvsdemo.local
cifs/cifs
HOST/cvs-east-433d.cvsdemo.local
HOST/CVS-EAST-433D
```

In a packet capture, we can see the Session Setup Request using the SPN tied to the CNAME.

431	4.156722	SMB2	308	Negotiate Protocol Response
432	4.156785	SMB2	232	Negotiate Protocol Request
434	4.158108	SMB2	374	Negotiate Protocol Response
435	4.160977	SMB2	1978	Session Setup Request
437	4.166224	SMB2	322	Session Setup Response
438	4.166891	SMB2	152	Tree Connect Request Tree: \\cifs\IPC\$
439	4.168063	SMB2	138	Tree Connect Response

```
realm: CVSDemo.LOCAL
  v sname
    name-type: kRB5-NT-SRV-INST (2)
    v sname-string: 2 items
      SNameString: cifs
      SNameString: cifs
  v enc-part
    etype: eTYPE-ARCFOUR-HMAC-MD5 (23)
```

SMB authentication dialects

Cloud Volumes Service supports the following [dialects](#) for SMB authentication:

- LM
- NTLM
- NTLMv2
- Kerberos

Kerberos authentication for SMB share access is the most secure level of authentication you can use. With AES and SMB encryption enabled, the security level is further increased.

Cloud Volumes Service also supports backward compatibility for LM and NTLM authentication. When Kerberos is misconfigured (such as when creating SMB aliases), share access falls back to weaker authentication methods (such as NTLMv2). Because these mechanisms are less secure, they are disabled in some Active Directory environments. If weaker authentication methods are disabled and Kerberos is not configured properly, share access fails because there is no valid authentication method to fall back to.

For information about configuring/viewing your supported authentication levels in Active Directory, see [Network security: LAN Manager authentication level](#).

Permission models

NTFS/File permissions

NTFS permissions are the permissions applied to files and folders in file systems adhering to NTFS logic. You can apply NTFS permissions in *Basic* or *Advanced* and can be set to *Allow* or *Deny* for access control.

Basic permissions include the following:

- Full Control
- Modify
- Read & Execute
- Read
- Write

When you set permissions for a user or group, referred to as an ACE, it resides in an ACL. NTFS permissions use the same read/write/execute basics as UNIX mode bits, but they can also extend to more granular and extended access controls (also known as Special Permissions), such as Take Ownership, Create Folders/Append Data, Write Attributes, and more.

Standard UNIX mode bits do not provide the same level of granularity as NTFS permissions (such as being able to set permissions for individual user and group objects in an ACL or setting extended attributes). However, NFSv4.1 ACLs do provide the same functionality as NTFS ACLs.

NTFS permissions are more specific than share permissions and can be used in conjunction with share permissions. With NTFS permission structures, the most restrictive applies. As such, explicit denials to a user or group overrides even Full Control when defining access rights.

NTFS permissions are controlled from Windows SMB clients.




Share permissions

Share permissions are more general than NTFS permissions (Read/Change/Full Control only) and control the initial entry into an SMB share—similar to how NFS export policy rules work.

Although NFS export policy rules control access through host-based information such as IP addresses or host names, SMB share permissions can control access by using user and group ACEs in a share ACL. You can set share ACLs either from the Windows client or from the Cloud Volumes Service management UI.

By default, share ACLs and initial volume ACLs include Everyone with Full Control. The file ACLs should be changed but share permissions are overruled by the file permissions on objects in the share.

For instance, if a user is only allowed Read access to the Cloud Volumes Service volume file ACL, they are denied access to create files and folders even though the share ACL is set to Everyone with Full Control, as shown in the following figure.



For best security results, do the following:

- Remove Everyone from the share and file ACLs and instead set share access for users or groups.
- Use groups for access control instead of individual users for ease of management and faster removal/addition of users to share ACLs through group management.
- Allow less restrictive, more general share access to the ACEs on the share permissions and lock down access to users and groups with file permissions for more granular access control.
- Avoid general use of explicit deny ACLs, because they override allow ACLs. Limit use of explicit deny ACLs for users or groups that need to be restricted from access to a file system quickly.
- Make sure that you pay attention to the [ACL inheritance](#) settings when modifying permissions; setting the inheritance flag at the top level of a directory or volume with high file counts means that each file below that

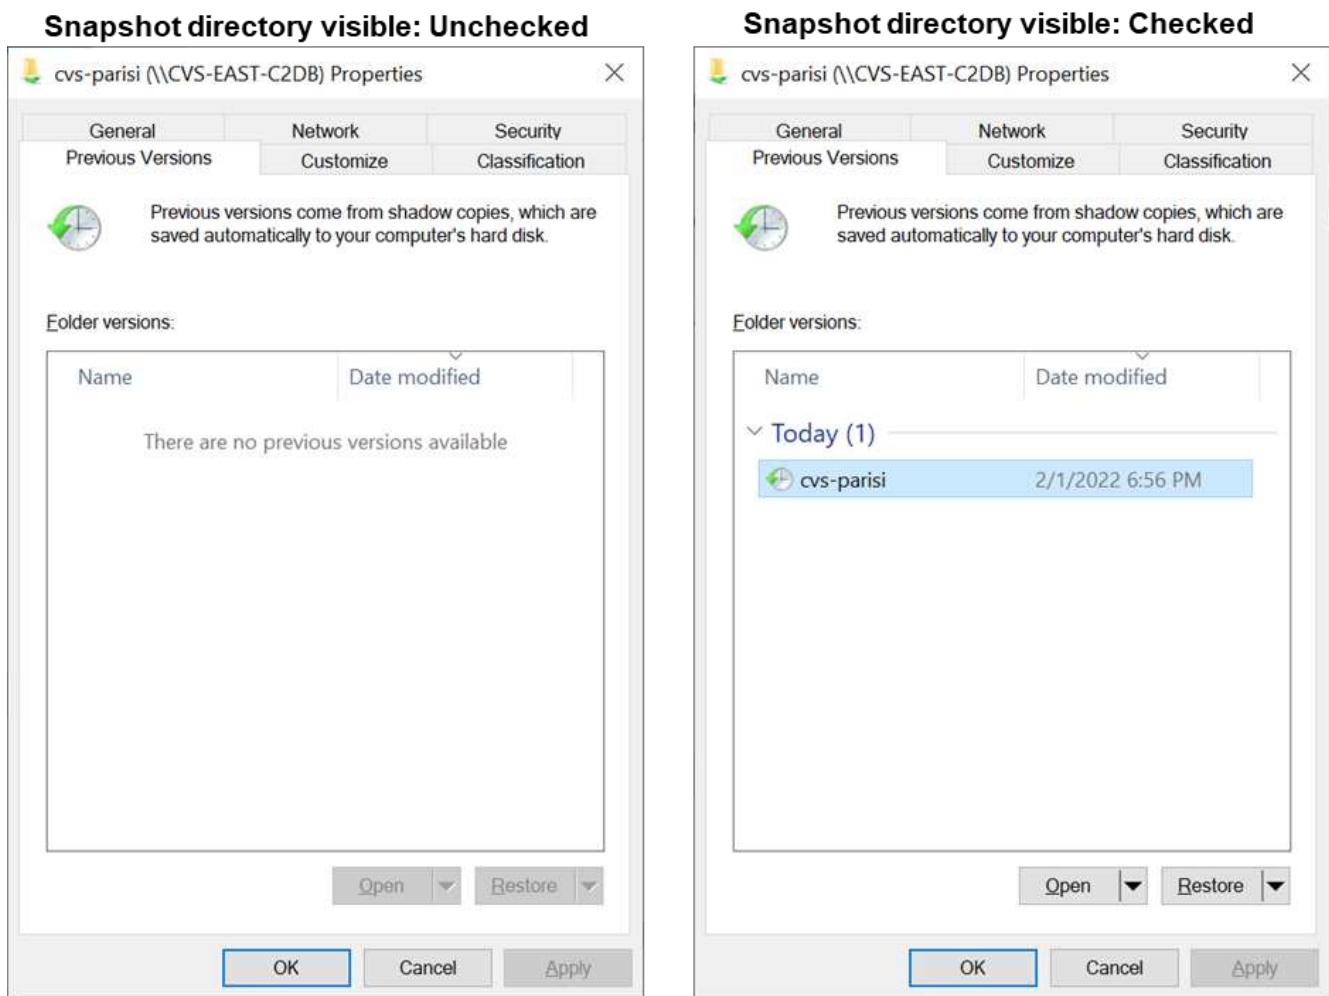
directory or volume has inherited permissions added to it, which can create unwanted behavior such as unintended access/denial and long churn of permission modification as each file is adjusted.

SMB share security features

When you first create a volume with SMB access in Cloud Volumes Service, you are presented with a series of choices for securing that volume.

Some of these choices depend on the Cloud Volumes Service level (Performance or Software) and choices include:

- **Make snapshot directory visible (available for both CVS-Performance and CVS-SW).** This option controls whether or not SMB clients can access the Snapshot directory in an SMB share (`\\server\share\~snapshot` and/or Previous Versions tab). The default setting is Not Checked, which means that the volume defaults to hiding and disallowing access to the `~snapshot` directory, and no Snapshot copies appear in the Previous Versions tab for the volume.



Hiding Snapshot copies from end users might be desired for security reasons, performance reasons (hiding these folders from AV scans) or preference. Cloud Volumes Service Snapshots are read-only, so even if these Snapshots are visible, end users cannot delete or modify files in the Snapshot directory. File permissions on the files or folders at the time the Snapshot copy was taken apply. If a file or folder's permissions change between Snapshot copies, then the changes also apply to the files or folders in the Snapshot directory. Users and groups can gain access to these files or folders based on permissions. While deletes or modifications of files in the Snapshot directory are not possible, it is possible to copy files or folders out of the Snapshot

directory.

- **Enable SMB encryption (available for both CVS-Performance and CVS-SW).** SMB encryption is disabled on the SMB share by default (unchecked). Checking the box enables SMB encryption, which means traffic between the SMB client and server is encrypted in-flight with the highest supported encryption levels negotiated. Cloud Volumes Service supports up to AES-256 encryption for SMB. Enabling SMB encryption does carry a performance penalty that might or might not be noticeable to your SMB clients—roughly in the 10-20% range. NetApp strongly encourages testing to see if that performance penalty is acceptable.
- **Hide SMB share (available for both CVS-Performance and CVS-SW).** Setting this option hides the SMB share path from normal browsing. This means that clients that do not know the share path cannot see the shares when accessing the default UNC path (such as \\CVS-SMB). When the checkbox is selected, only clients that explicitly know the SMB share path or have the share path defined by a Group Policy Object can access it (security through obfuscation).
- **Enable access-based enumeration (ABE) (CVS-SW only).** This is similar to hiding the SMB share, except the shares or files are only hidden from users or groups that do not have permissions to access the objects. For instance, if Windows user joe is not allowed at least Read access through the permissions, then the Windows user joe cannot see the SMB share or files at all. This is disabled by default, and you can enable it by selecting the checkbox. For more information on ABE, see the NetApp Knowledge Base article [How does Access Based Enumeration \(ABE\) work?](#)
- **Enable Continuously Available (CA) share support (CVS-Performance only).** [Continuously Available SMB shares](#) provide a way to minimize application disruptions during failover events by replicating lock states across nodes in the Cloud Volumes Service backend system. This is not a security feature, but it does offer better overall resiliency. Currently, only SQL Server and FSLogix applications are supported for this functionality.

Default hidden shares

When an SMB server is created in Cloud Volumes Service, there are [hidden administrative shares](#) (using the \$ naming convention) that are created in addition to the data volume SMB share. These include C\$ (namespace access) and IPC\$ (sharing named pipes for communication between programs, such as the remote procedure calls (RPC) used for Microsoft Management Console (MMC) access).

The IPC\$ share contains no share ACLs and cannot be modified—it is strictly used for RPC calls and [Windows disallows anonymous access to these shares by default](#).

The C\$ share allows BUILTIN/Administrators access by default, but Cloud Volumes Service automation removes the share ACL and does not allow access to anyone because access to the C\$ share allows visibility into all mounted volumes in the Cloud Volumes Service file systems. As a result, attempts to navigate to \\SERVER\C\$ fail.

Accounts with local/BUILTIN administrator/backup rights

Cloud Volumes Service SMB servers maintain similar functionality to regular Windows SMB servers in that there are local groups (such as BUILTIN\Administrators) that apply access rights to select domain users and groups.

When you specify a user to be added to Backup Users, the user is added to the BUILTIN\Backup Operators group in the Cloud Volumes Service instance that uses that Active Directory connection, which then gets the [SeBackupPrivilege and SeRestorePrivilege](#).

When you add a user to Security Privilege Users, the user is given the SeSecurityPrivilege, which is useful in some application use cases, such as [SQL Server on SMB shares](#).

Backup Users

Provide a comma separated list of domain users or a domain group name that require elevated privileges to access volumes created by Cloud Volumes Service.


Accountnames
administrator,cvs-svc

Security Privilege Users

Provide a list of comma separated domain user accounts that require elevated privileges to manage security log for the Active Directory associated with Cloud Volumes Service.

Accountnames
administrator,cvs-svc

You can view Cloud Volumes Service local group memberships through the MMC with the proper privileges. The following figure shows users that have been added by using the Cloud Volumes Service console.



The following table shows the list of default BUILTIN groups and what users/groups are added by default.

Local/BUILTIN group	Default members
BUILTIN\Administrators*	DOMAIN\Domain Admins
BUILTIN\Backup Operators*	None
BUILTIN\Guests	DOMAIN\Domain Guests
BUILTIN\Power Users	None
BUILTIN\Domain Users	DOMAIN\Domain Users

*Group membership controlled in Cloud Volumes Service Active Directory connection configuration.

You can view local users and groups (and group members) in the MMC window, but you cannot add or delete objects or change group memberships from this console. By default, only the Domain Admins group and Administrator are added to the BUILTIN\Administrators group in Cloud Volumes Service. Currently, you cannot modify this.

Computer Management (CVS-EAST-C2DB)		
Name	Full Name	Description
Administrator		Built-in administrator account

Computer Management (CVS-EAST-C2DB)		
Name	Full Name	Description
Administrators		Built-in Administrators group
Users		All users
Guests		Built-in Guests Group
Power Users		Restricted administrative privileges
Backup Operators		Backup Operators group

Administrators Properties

General

Administrators

Description:

Members:

Administrator

CVSDemo Domain Admins

Add...
Remove

Changes to a user's group membership are not effective until the next time the user logs on.

OK
Cancel
Apply
Help

MMC/Computer Management access


SMB access in Cloud Volumes Service provides connectivity to the Computer Management MMC, which allows you to view shares, manage share ACLs, and view/manage SMB sessions and open files.


To use the MMC to view SMB shares and sessions in Cloud Volumes Service, the user logged in currently must be a domain administrator. Other users are allowed access to view or manage the SMB server from MMC and receive a You Do Not Have Permissions dialog box when attempting to view shares or sessions on the Cloud Volumes Service SMB instance.

To connect to the SMB server, open Computer Management, right click Computer Management and then select Connect To Another Computer. This opens the Select Computer dialog box where you can enter the SMB server name (found in the Cloud Volumes Service volume information).

When you view SMB shares with the proper permissions, you see all available shares in the Cloud Volumes Service instance that share the Active Directory connection. To control this behavior, set the Hide SMB Shares option on the Cloud Volumes Service volume instance.

Remember, only one Active Directory connection is allowed per region.





The following table shows a list of supported/unsupported functionality for the MMC.

Supported functions	Unsupported functions
<ul style="list-style-type: none"> • View shares • View active SMB sessions • View open files • View local users and groups • View local group memberships • Enumerate the list of sessions, files, and tree connections in the system • Close open files in the system • Close open sessions • Create/manage shares 	<ul style="list-style-type: none"> • Creating new local users/groups • Managing/viewing existing local user/groups • View events or performance logs • Managing storage • Managing services and applications

SMB server security information

The SMB server in Cloud Volumes Service uses a series of options that define security policies for SMB connections, including things such as Kerberos clock skew, ticket age, encryption, and more.

The following table contains a list of those options, what they do, the default configurations, and if they can be modified with Cloud Volumes Service. Some options do not apply to Cloud Volumes Service.

Security option	What it does	Default value	Can change?
Maximum Kerberos Clock Skew (minutes)	Maximum time skew between Cloud Volumes Service and domain controllers. If the time skew exceeds 5 minutes, Kerberos authentication fails. This is set to the Active Directory default value.	5	No
Kerberos Ticket Lifetime (hours)	Maximum time a Kerberos ticket remains valid before requiring a renewal. If no renewal occurs before the 10 hours, you must obtain a new ticket. Cloud Volumes Service performs these renewals automatically. 10 hours is the Active Directory default value.	10	No
Maximum Kerberos Ticket Renewal (days)	Maximum number of days that a Kerberos ticket can be renewed before a new authorization request is needed. Cloud Volumes Service automatically renews tickets for SMB connections. Seven days is the Active Directory default value.	7	No
Kerberos KDC Connection Timeout (secs)	The number of seconds before a KDC connection times out.	3	No
Require Signing for Incoming SMB Traffic	Setting to require signing for SMB traffic. If set to true, clients that do not support signing fail connectivity.	False	
Require Password Complexity for Local User Accounts	Used for passwords on local SMB users. Cloud Volumes Service does not support local user creation, so this option does not apply to Cloud Volumes Service.	True	No

Security option	What it does	Default value	Can change?
Use start_tls for Active Directory LDAP Connections	Used to enable start TLS connections for Active Directory LDAP. Cloud Volumes Service does not currently support enabling this.	False	No
Is AES-128 and AES-256 Encryption for Kerberos Enabled	This controls whether AES encryption is used for Active Directory connections and is controlled with the Enable AES Encryption for Active Directory Authentication option when creating/modifying the Active Directory connection.	False	Yes
LM Compatibility Level	Level of supported authentication dialects for Active Directory connections. See the section “SMB authentication dialects” for more information.	ntlmv2-krb	No
Require SMB Encryption for Incoming CIFS Traffic	Requires SMB encryption for all shares. This is not used by Cloud Volumes Service; instead, set encryption on a per-volume basis (see the section “SMB share security features”).	False	No
Client Session Security	Sets signing and/or sealing for LDAP communication. This is not currently set in Cloud Volumes Service but might be needed in future releases to address . Remediation for LDAP authentication issues due to the Windows patch is covered in the section “LDAP channel binding.” .	None	No
SMB2 enable for DC connections	Uses SMB2 for DC connections. Enabled by default.	System-default	No

Security option	What it does	Default value	Can change?
LDAP Referral Chasing	When using multiple LDAP servers, referral chasing allows the client to refer to other LDAP servers in the list when an entry is not found in the first server. This is currently not supported by Cloud Volumes Service.	False	No
Use LDAPS for Secure Active Directory Connections	Enables the use of LDAP over SSL. Currently not supported by Cloud Volumes Service.	False	No
Encryption is required for DC Connection	Requires encryption for successful DC connections. Disabled by default in Cloud Volumes Service.	False	No

Dual-protocol/multiprotocol

Cloud Volumes Service offers the ability to share the same datasets to both SMB and NFS clients while maintaining proper access permissions ([dual-protocol](#)). This is done by coordinating identity mapping between protocols and using a centralized backend LDAP server to provide the UNIX identities to Cloud Volumes Service. You can use Windows Active Directory to provide both Windows and UNIX users for ease of use.

Access control

- **Share access controls.** Determine which clients and/or user and groups can access a NAS share. For NFS, export policies and rules control client access to exports. NFS exports are managed from the Cloud Volumes Service instance. SMB makes use of CIFS/SMB shares and share ACLs to provide more granular control at the user and group level. You can only configure share-level ACLs from SMB clients by using [MMC/Computer Management](#) with an account that has administrator rights on the Cloud Volumes Service instance (see the section [“Accounts with local/BUILTIN administrator/backup rights.”](#)).
- **File access controls.** Control permissions at a file or folder level and are always managed from the NAS client. NFS clients can make use of traditional mode bits (rwx) or NFSv4 ACLs. SMB clients leverage NTFS permissions.

The access control for volumes that serve data to both NFS and SMB depends on the protocol in use. For information on permissions with dual protocol, see the section [“Permission model.”](#)

User mapping

When a client accesses a volume, Cloud Volumes Service attempts to map the incoming user to a valid user in the opposite direction. This is necessary for proper access to be determined across protocols and to ensure that the user requesting access is indeed who they claim to be.

For example, if a Windows user named `joe` attempts access to a volume with UNIX permissions through SMB, then Cloud Volumes Service performs a search to find a corresponding UNIX user named `joe`. If one exists,

then files that are written to an SMB share as Windows user `joe` appears as UNIX user `joe` from NFS clients.

Alternately, if a UNIX user named `joe` attempts access to a Cloud Volumes Service volume with Windows permissions, then the UNIX user must be able to map to a valid Windows user. Otherwise, access to the volume is denied.

Currently, only Active Directory is supported for external UNIX identity management with LDAP. For more information about configuring access to this service, see [Creating an AD connection](#).

Permission model

When using dual-protocol setups, Cloud Volumes Service makes use of security styles for volumes to determine the type of ACL. These security styles are set based on which NAS protocol is specified, or in the case of dual protocol, is a choice made at the time of Cloud Volumes Service volume creation.

- If you are only using NFS, Cloud Volumes Service volumes use UNIX permissions.
- If you are only using SMB, Cloud Volumes Service volumes use NTFS permissions.

If you are creating a dual-protocol volume, you can choose the ACL style at volume creation. This decision should be made based on the desired permissions management. If your users manage permissions from Windows/SMB clients, select NTFS. If your users prefer using NFS clients and `chmod/chown`, use UNIX security styles.

Considerations for creating Active Directory connections

Cloud Volumes Service provides the ability to connect your Cloud Volumes Service instance to an external Active Directory server for identity management for both SMB and UNIX users. Creating an Active Directory connection is required to use SMB in Cloud Volumes Service.

The configuration for this provides several options that require some consideration for security. The external Active Directory server can be an on-premises instance or cloud native. If you are using an on-premises Active Directory server, don't expose the domain to the external network (such as with a DMZ or an external IP address). Instead, use secure private tunnels or VPNs, one-way forest trusts, or dedicated network connections to the on-premises networks with [Private Google Access](#). See the Google Cloud documentation for more information about [best practices using Active Directory in Google Cloud](#).



CVS-SW requires Active Directory servers to be located in the same region. If a DC connection is attempted in CVS-SW to another region, the attempt fails. When using CVS-SW, be sure to create Active Directory sites that include the Active Directory DCs and then specify sites in Cloud Volumes Service to avoid cross-region DC connection attempts.

Active Directory credentials

When SMB or LDAP for NFS is enabled, Cloud Volumes Service interacts with the Active Directory controllers to create a machine account object to use for authentication. This is no different from how a Windows SMB client joins a domain and requires the same access rights to Organizational Units (OUs) in Active Directory.

In many cases, security groups do not allow the use of a Windows administrator account on external servers such as Cloud Volumes Service. In some cases, the Windows Administrator user is disabled entirely as a security best practice.

Permissions needed to create SMB machine accounts

To add Cloud Volumes Service machine objects to an Active Directory, an account that either has administrative rights to the domain or has [delegated permissions to create and modify machine account objects](#) to a specified OU is required. You can do this with the Delegation of Control Wizard in Active Directory by creating a custom task that provides a user access to creation/deletion of computer objects with the following access permissions provided:

- Read/Write
- Create/Delete All Child Objects
- Read/Write All Properties
- Change/Reset Password

Doing this automatically adds a security ACL for the defined user to the OU in Active Directory and minimizes the access to the Active Directory environment. After a user has been delegated, that username and password can be provided as Active Directory Credentials in this window.



The username and password that is passed to the Active Directory domain leverages Kerberos encryption during the machine account object query and creation for added security.

Active Directory connection details

The [Active Directory Connection Details](#) provide fields for administrators to give specific Active Directory schema information for machine account placement, such as the following:

- **Active Directory Connection Type.** Used to specify whether the Active Directory connection in a region is used for volumes of either Cloud Volumes Service or CVS-Performance service type. If this is set incorrectly on an existing connection, it might not work properly when used or edited.
- **Domain.** The Active Directory domain name.
- **Site.** Limits Active Directory servers to a specific site for security and performance [considerations](#). This is necessary when multiple Active Directory servers span regions because Cloud Volumes Service does not currently support allowing Active Directory authentication requests to Active Directory servers in a different region than the Cloud Volumes Service instance. (For instance, the Active Directory domain controller is in a region that only CVS-Performance supports but you want an SMB share in a CVS-SW instance.)
- **DNS servers.** DNS servers to use in name lookups.
- **NetBIOS name (optional).** If desired, the NetBIOS name for the server. This what is used when new machine accounts are created using the Active Directory connection. For instance, if the NetBIOS name is set to CVS-EAST then the machine account names will be CVS-EAST-{1234}. See the section "[How Cloud Volumes Service shows up in Active Directory](#)" for more information.
- **Organizational Unit (OU).** The specific OU to create the computer account. This is useful if you're delegating control to a user for machine accounts to a specific OU.
- **AES Encryption.** You can also check or uncheck the Enable AES Encryption for AD Authentication checkbox. Enabling AES encryption for Active Directory authentication provides extra security for Cloud Volumes Service to Active Directory communication during user and group lookups. Before enabling this option, check with your domain administrator to confirm that the Active Directory domain controllers support AES authentication.



By default, most Windows servers do not disable weaker ciphers (such as DES or RC4-HMAC), but if you choose to disable weaker ciphers, confirm Cloud Volumes Service Active Directory connection has been configured to enable AES. Otherwise, authentication failures occur. Enabling AES encryption doesn't disable weaker ciphers but instead adds support for AES ciphers to the Cloud Volumes Service SMB machine account.

Kerberos realm details

This option does not apply to SMB servers. Rather, it is used when configuring NFS Kerberos for the Cloud Volumes Service system. When these details are populated, the NFS Kerberos realm is configured (similar to a `krb5.conf` file on Linux) and is used when NFS Kerberos is specified on the Cloud Volumes Service volume creation, as the Active Directory connection acts as the NFS Kerberos Distribution Center (KDC).



Non-Windows KDCs are currently unsupported for use with Cloud Volumes Service.


Region

A region enables you to specify the location where the Active Directory connection resides. This region must be the same region as the Cloud Volumes Service volume.

- **Local NFS Users with LDAP.** In this section, there is also an option to Allow Local NFS Users with LDAP. This option must be left unselected if you want to extend your UNIX user group membership support beyond the 16-group limitation of NFS (extended groups). However, using extended groups requires a configured LDAP server for UNIX identities. If you don't have an LDAP server, leave this option unselected. If you have an LDAP server and want to also use local UNIX users (such as root), select this option.

Backup users

This option enables you to specify Windows users that have backup permissions to the Cloud Volumes Service volume. Backup privileges (`SeBackupPrivilege`) are necessary for some applications to properly backup and restore data in NAS volumes. This user has a high level of access to data in the volume, so you should consider [enabling auditing of that user access](#). After it is enabled, audit events display in Event Viewer > Windows Logs > Security.



Security privilege users

This option enables you to specify Windows users that have security modification permissions to the Cloud Volumes Service volume. Security privileges (SeSecurityPrivilege) are necessary for some applications (such as [SQL Server](#)) to properly set permissions during installation. This privilege is needed to manage the security log. Although this privilege is not as powerful as SeBackupPrivilege, NetApp recommends [auditing user access of users](#) with this privilege level if needed.

For more information, see [Special privileges assigned to new logon](#).

How Cloud Volumes Service shows up in Active Directory

Cloud Volumes Service shows up in Active Directory as a normal machine account object. The naming conventions are as follows.

- CIFS/SMB and NFS Kerberos create separate machine account objects.
- NFS with LDAP enabled creates a machine account in Active Directory for Kerberos LDAP binds.
- Dual protocol volumes with LDAP share the CIFS/SMB machine account for LDAP and SMB.
- CIFS/SMB machine accounts use a naming convention of NAME-1234 (random four digit ID with hyphen appended to <10 character name) for the machine account. You can define NAME by the NetBIOS name setting on the Active Directory connection (see the section "[Active Directory connection details](#)").
- NFS Kerberos uses NFS-NAME-1234 as the naming convention (up to 15 characters). If more than 15 characters are used, the name is NFS-TRUNCATED-NAME-1234.

- NFS-only CVS-Performance instances with LDAP enabled create an SMB machine account for binding to the LDAP server with the same naming convention as CIFS/SMB instances.
- When an SMB machine account is created, default hidden admin shares (see the section [“Default hidden shares”](#)) are also created (c\$, admin\$, ipc\$), but those shares have no ACLs assigned and are inaccessible.
- The machine account objects are placed in CN=Computers by default, but you can specify a different OU when necessary. See the section [“Permissions needed to create SMB machine accounts”](#) for information about what access rights are needed to add/remove machine account objects for Cloud Volumes Service.

When Cloud Volumes Service adds the SMB machine account to Active Directory, the following fields are populated:

- cn (with the specified SMB server name)
- dnsHostName (with SMBserver.domain.com)
- msDS-SupportedEncryptionTypes (Allows DES_CBC_MD5, RC4_HMAC_MD5 if AES encryption is not enabled; if AES encryption is enabled, DES_CBC_MD5, RC4_HMAC_MD5, AES128_CTS_HMAC_SHA1_96, AES256_CTS_HMAC_SHA1_96 are allowed for Kerberos ticket exchange with the machine account for SMB)
- name (with the SMB server name)
- sAMAccountName (with SMBserver\$)
- servicePrincipalName (with host/smbserver.domain.com and host/smbserver SPNs for Kerberos)

If you want to disable weaker Kerberos encryption types (enctype) on the machine account, you can change the msDS-SupportedEncryptionTypes value on the machine account to one of the values in the following table to allow AES only.

msDS-SupportedEncryptionTypes value	Enctype enabled
2	DES_CBC_MD5
4	RC4_HMAC
8	AES128_CTS_HMAC_SHA1_96 only
16	AES256_CTS_HMAC_SHA1_96 only
24	AES128_CTS_HMAC_SHA1_96 and AES256_CTS_HMAC_SHA1_96
30	DES_CBC_MD5, RC4_HMAC, AES128_CTS_HMAC_SHA1_96 and AES256_CTS_HMAC_SHA1_96

To enable AES encryption for SMB machine accounts, click Enable AES Encryption for AD Authentication when creating the Active Directory connection.

To enable AES encryption for NFS Kerberos, [see the Cloud Volumes Service documentation](#).

Other NAS Infrastructure service dependencies (KDC, LDAP, and DNS)

When using Cloud Volumes Service for NAS shares, there might be external dependencies required for proper functionality. These dependencies are in play under specific circumstances. The following table shows various configuration options and what,

if any, dependencies are required.

Configuration	Dependencies required
NFSv3 only	None
NFSv3 Kerberos only	Windows Active Directory: * KDC * DNS * LDAP
NFSv4.1 only	Client ID mapping configuration (/etc/idmap.conf)
NFSv4.1 Kerberos only	<ul style="list-style-type: none">• Client ID mapping configuration (/etc/idmap.conf)• Windows Active Directory: KDC DNS LDAP
SMB only	Active Directory: * KDC * DNS
Multiprotocol NAS (NFS and SMB)	<ul style="list-style-type: none">• Client ID mapping configuration (NFSv4.1 only; /etc/idmap.conf)• Windows Active Directory: KDC DNS LDAP

Kerberos keytab rotation/password resets for machine account objects

With SMB machine accounts, Cloud Volumes Service schedules periodic password resets for the SMB machine account. These password resets occur using Kerberos encryption and operate on a schedule of every fourth Sunday at a random time between 11PM and 1AM. These password resets change the Kerberos key versions, rotate the keytabs stored on the Cloud Volumes Service system, and help maintain a greater level of security for SMB servers running in Cloud Volumes Service. Machine account passwords are randomized and are not known to administrators.

For NFS Kerberos machine accounts, password resets take place only when a new keytab is created/exchanged with the KDC. Currently, this is not possible to do in Cloud Volumes Service.

Network ports for use with LDAP and Kerberos

When using LDAP and Kerberos, you should determine the network ports in use by these services. You can find a complete list of ports in use by Cloud Volumes Service in the [Cloud Volumes Service documentation on security considerations](#).

LDAP

Cloud Volumes Service acts as an LDAP client and uses standard LDAP search queries for user and group lookups for UNIX identities. LDAP is necessary if you intend to use users and groups outside the standard default users provided by Cloud Volumes Service. LDAP is also necessary if you plan on using NFS Kerberos with user principals (such as [user1@domain.com](#)). Currently, only LDAP using Microsoft Active Directory is

supported.

To use Active Directory as a UNIX LDAP server, you must populate the necessary UNIX attributes on users and groups you intend to use for UNIX identities. Cloud Volumes Service uses a default LDAP schema template that queries attributes based on [RFC-2307-bis](#). As a result, the following table shows the bare minimum necessary Active Directory attributes to populate for users and groups and what each attribute is used for.

For more information about setting LDAP attributes in Active Directory, see [Managing dual-protocol access](#).

Attribute	What it does
uid*	Specifies the UNIX user name
uidNumber*	Specifies the UNIX user's numeric ID
gidNumber*	Specifies the UNIX user's primary group numeric ID
objectClass*	Specifies what type of object is being used; Cloud Volumes Service requires "user" to be included in the list of object classes (is included in most Active Directory deployments by default).
name	General information about the account (real name, phone number, and so on—also known as geocos)
unixUserPassword	No need to set this; not used in UNIX identity lookups for NAS authentication. Setting this puts the configured unixUserPassword value in plaintext.
unixHomeDirectory	Defines path to UNIX home directories when a user authenticates against LDAP from a Linux client. Set this if you want to use LDAP for UNIX home directory functionality.
loginShell	Defines path to the bash/profile shell for Linux clients when a user authenticates against LDAP.

*Denotes attribute is required for proper functionality with Cloud Volumes Service. Remaining attributes are for client-side use only.

Attribute	What it does
cn*	Specifies the UNIX group name. When using Active Directory for LDAP, this is set when the object is first created, but it can be changed later. This name cannot be the same as other objects. For instance, if your UNIX user named user1 belongs to a group named user1 on your Linux client, Windows doesn't allow two objects with the same cn attribute. To work around this, rename the Windows user to a unique name (such as user1-UNIX); LDAP in Cloud Volumes Service uses the uid attribute for UNIX user names.
gidNumber*	Specifies the UNIX group numeric ID.

Attribute	What it does
objectClass*	Specifies what type of object is being used; Cloud Volumes Service requires group to be included in the list of object classes (this attribute is included in most Active Directory deployments by default).
memberUid	Specifies which UNIX users are members of the UNIX group. With Active Directory LDAP in Cloud Volumes Service, this field is not necessary. The Cloud Volumes Service LDAP schema uses the Member field for group memberships.
Member*	Required for group memberships/secondary UNIX groups. This field is populated by adding Windows users to Windows groups. However, if the Windows groups don't have UNIX attributes populated, they are not included in the UNIX user's group membership lists. Any groups that need to be available in NFS must populate the required UNIX group attributes listed in this table.

*Denotes attribute is required for proper functionality with Cloud Volumes Service. Remaining attributes are for client-side use only.

LDAP bind information

To query users in LDAP, Cloud Volumes Service must bind (login) to the LDAP service. This login has read-only permissions and is used to query LDAP UNIX attributes for directory lookups. Currently, LDAP binds are possible only by using an SMB machine account.

You can only enable LDAP for `CVS-Performance` instances and use it for NFSv3, NFSv4.1, or dual-protocol volumes. An Active Directory connection must be established in the same region as the Cloud Volumes Service volume for successful deployment of the LDAP-enabled volume.

When LDAP is enabled, the following occurs in specific scenarios.

- If only NFSv3 or NFSv4.1 is used for the Cloud Volumes Service project, then a new machine account is created in the Active Directory domain controller, and the LDAP client in Cloud Volumes Service binds to Active Directory by using the machine account credentials. No SMB shares are created for the NFS volume and default hidden administrative shares (see the section [“Default hidden shares”](#)) have share ACLs removed.
- If dual-protocol volumes are used for the Cloud Volumes Service project, then only the single machine account created for SMB access is used to bind the LDAP client in Cloud Volumes Service to Active Directory. No additional machine accounts are created.
- If dedicated SMB volumes are created separately (either before or after NFS volumes with LDAP are enabled), then the machine account for LDAP binds is shared with the SMB machine account.
- If NFS Kerberos is also enabled, two machine accounts are created—one for SMB shares and/or LDAP binds and one for NFS Kerberos authentication.

LDAP queries

Although LDAP binds are encrypted, LDAP queries are passed over the wire in plaintext by using the common LDAP port 389. This well-known port cannot currently be changed in Cloud Volumes Service. As a result,

someone with access to packet sniffing in the network can see user and group names, numeric IDs, and group memberships.

However, Google Cloud VMs cannot sniff other VM's unicast traffic. Only VMs actively participating in LDAP traffic (that is, being able to bind) can see traffic from the LDAP server. For more information about packet sniffing in Cloud Volumes Service, see the section [“Packet sniffing/trace considerations.”](#)

LDAP client configuration defaults

When LDAP is enabled in a Cloud Volumes Service instance, an LDAP client configuration is created with specific configuration details by default. In some cases, options either do not apply to Cloud Volumes Service (not supported) or are not configurable.

LDAP client option	What it does	Default value	Can change?
LDAP Server List	Sets LDAP server names or IP addresses to use for queries. This is not used for Cloud Volumes Service. Instead, Active Directory Domain is used to define LDAP servers.	Not set	No
Active Directory Domain	Sets the Active Directory Domain to use for LDAP queries. Cloud Volumes Service leverages SRV records for LDAP in DNS to find LDAP servers in the domain.	Set to the Active Directory domain specified in the Active Directory connection.	No
Preferred Active Directory Servers	Sets the preferred Active Directory servers to use for LDAP. Not supported by Cloud Volumes Service. Instead, use Active Directory sites to control LDAP server selection.	Not set.	No
Bind using SMB Server Credentials	Binds to LDAP by using the SMB machine account. Currently, the only supported LDAP bind method in Cloud Volumes Service.	True	No
Schema Template	The schema template used for LDAP queries.	MS-AD-BIS	No
LDAP Server Port	The port number used for LDAP queries. Cloud Volumes Service currently uses only the standard LDAP port 389. LDAPS/port 636 is not currently supported.	389	No

LDAP client option	What it does	Default value	Can change?
Is LDAPS Enabled	Controls whether LDAP over Secure Sockets Layer (SSL) is used for queries and binds. Currently not supported by Cloud Volumes Service.	False	No
Query Timeout (sec)	Timeout for queries. If queries take longer than the specified value, queries fail.	3	No
Minimum Bind Authentication Level	The minimum supported bind level. Because Cloud Volumes Service uses machine accounts for LDAP binds and Active Directory does not support anonymous binds by default, this option does not come into play for security.	Anonymous	No
Bind DN	The user/distinguished name (DN) used for binds when simple bind is used. Cloud Volumes Service uses machine accounts for LDAP binds and does not currently support simple bind authentication.	Not set	No
Base DN	The base DN used for LDAP searches.	The Windows domain use for the Active Directory connection, in DN format (that is, DC=domain, DC=local).	No
Base search scope	The search scope for base DN searches. Values can include base, onelevel, or subtree. Cloud Volumes Service only supports subtree searches.	Subtree	No
User DN	Defines the DN where user searches start for LDAP queries. Currently not supported for Cloud Volumes Service, so all user searches start at the base DN.	Not set	No

LDAP client option	What it does	Default value	Can change?
User search scope	The search scope for user DN searches. Values can include base, onelevel, or subtree. Cloud Volumes Service does not support setting the user search scope.	Subtree	No
Group DN	Defines the DN where group searches start for LDAP queries. Currently not supported for Cloud Volumes Service, so all group searches start at the base DN.	Not set	No
Group search scope	The search scope for group DN searches. Values can include base, onelevel, or subtree. Cloud Volumes Service does not support setting the group search scope.	Subtree	No
Netgroup DN	Defines the DN where netgroup searches start for LDAP queries. Currently not supported for Cloud Volumes Service, so all netgroup searches start at the base DN.	Not set	No
Netgroup search scope	The search scope for netgroup DN searches. Values can include base, onelevel, or subtree. Cloud Volumes Service does not support setting the netgroup search scope.	Subtree	No
Use start_tls over LDAP	Leverages Start TLS for certificate based LDAP connections over port 389. Currently not supported by Cloud Volumes Service.	False	No
Enable netgroup-by-host lookup	Enables netgroup lookups by hostname rather than expanding netgroups to list all members. Currently not supported by Cloud Volumes Service.	False	No

LDAP client option	What it does	Default value	Can change?
Netgroup-by-host DN	Defines the DN where netgroup-by-host searches start for LDAP queries. Netgroup-by-host is currently not supported for Cloud Volumes Service.	Not set	No
Netgroup-by-host search scope	The search scope for netgroup-by-host DN searches. Values can include base, onelevel or subtree. Netgroup-by-host is currently not supported for Cloud Volumes Service.	Subtree	No
Client session security	Defines what level of session security is used by LDAP (sign, seal, or none). LDAP signing is supported by CVS-Performance, if requested by Active Directory. CVS-SW does not support LDAP signing. For both service types, sealing is currently not supported.	None	No
LDAP referral chasing	When using multiple LDAP servers, referral chasing allows the client to refer to other LDAP servers in the list when an entry is not found in the first server. This is currently not supported by Cloud Volumes Service.	False	No
Group membership filter	Provides a custom LDAP search filter to be used when looking up group membership from an LDAP server. Not currently supported with Cloud Volumes Service.	Not set	No

Using LDAP for asymmetric name mapping

Cloud Volumes Service, by default, maps Windows users and UNIX users with identical usernames bidirectionally without special configuration. As long as Cloud Volumes Service can find a valid UNIX user (with LDAP), then 1:1 name mapping occurs. For instance, if Windows user `johnsmith` is used, then, if Cloud Volumes Service can find a UNIX user named `johnsmith` in LDAP, name mapping succeeds for that user, all files/folders created by `johnsmith` show the correct user ownership, and all ACLs affecting `johnsmith` are

honored regardless of the NAS protocol in use. This is known as symmetric name mapping.

Asymmetric name mapping is when the Windows user and UNIX user identity don't match. For instance, if Windows user `johnsmith` has a UNIX identity of `jsmith`, Cloud Volumes Service needs a way to be told about the variation. Because Cloud Volumes Service currently doesn't support creation of static name mapping rules, LDAP must be used to look up the identity of the users for both Windows and UNIX identities to ensure proper ownership of files and folders and expected permissions.

By default, Cloud Volumes Service includes LDAP in the ns-switch of the instance for the name map database, so that to provide name mapping functionality by using LDAP for asymmetric names, you only need to modify some of the user/group attributes to reflect what Cloud Volumes Service looks for.

The following table shows what attributes must be populated in LDAP for asymmetric name mapping functionality. In most cases, Active Directory is already configured to do this.

Cloud Volumes Service attribute	What it does	Value used by Cloud Volumes Service for name mapping
Windows to UNIX objectClass	Specifies the type of object being used. (That is, user, group, posixAccount, and so on)	Must include user (can contain multiple other values, if desired.)
Windows to UNIX attribute	that defines the Windows username at creation. Cloud Volumes Service uses this for Windows to UNIX lookups.	No change needed here; sAMAccountName is the same as the Windows login name.
UID	Defines the UNIX username.	Desired UNIX username.

Cloud Volumes Service currently does not use domain prefixes in LDAP lookups, so multiple domain LDAP environments do not function properly with LDAP namemap lookups.

The following example shows a user with the Windows name `asymmetric`, the UNIX name `unix-user`, and the behavior it follows when writing files from both SMB and NFS.

The following figure shows how LDAP attributes look from the Windows server.

Published Certificates	Member Of	Password Replication	Dial-in	Object
Security	Environment	Sessions	Remote control	
General	Address	Account	Profile	Telephones
Remote Desktop Services Profile		COM+	Attribute Editor	

Attributes:

Attribute	Value
name	asymmetric
objectCategory	CN=Person,CN=Schema,CN=Configuration,
objectClass	top; person; organizationalPerson; user
objectGUID	de489556-dd7b-43a3-98fa-2722f79d67ed
objectSid	S-1-5-21-3552729481-4032800560-2279794
primaryGroupID	513 = (GROUP_RID_USERS)
pwdLastSet	1/19/2017 1:56:34 PM Eastern Standard Time
replPropertyMetaData	AttID Ver Loc.USN Org.DSA
sAMAccountName	asymmetric
sAMAccountType	805306368 = (NORMAL_USER_ACCOUNT
uid	unix-user
uidNumber	1207

From an NFS client, you can query the UNIX name but not the Windows name:

```
# id unix-user
uid=1207(unix-user) gid=1220(sharedgroup) groups=1220(sharedgroup)
# id asymmetric
id: asymmetric: no such user
```

When a file is written from NFS as `unix-user`, the following is the result from the NFS client:

```
sh-4.2$ pwd
/mnt/home/ntfssh-4.2$ touch unix-user-file
sh-4.2$ ls -la | grep unix-user
-rwx----- 1 unix-user sharedgroup 0 Feb 28 12:37 unix-user-nfs
sh-4.2$ id
uid=1207(unix-user) gid=1220(sharedgroup) groups=1220(sharedgroup)
```

From a Windows client, you can see that the owner of the file is set to the proper Windows user:

```
PS C:\ > Get-Acl \\demo\home\ntfs\unix-user-nfs | select Owner
Owner
-----
NTAP\asymmetric
```

Conversely, files created by the Windows user `asymmetric` from an SMB client show the proper UNIX owner, as shown in the following text.

SMB:

```
PS Z:\ntfs> echo TEXT > asymmetric-user-smb.txt
```

NFS:

```
sh-4.2$ ls -la | grep asymmetric-user-smb.txt
-rwx----- 1 unix-user      sharedgroup  14 Feb 28 12:43 asymmetric-
user-smb.txt
sh-4.2$ cat asymmetric-user-smb.txt
TEXT
```

LDAP channel binding

Because of a vulnerability with Windows Active Directory domain controllers, [Microsoft Security Advisory ADV190023](#) changes how DCs allow LDAP binds.

The impact for Cloud Volumes Service is the same as for any LDAP client. Cloud Volumes Service does not currently support channel binding. Because Cloud Volumes Service supports LDAP signing by default through negotiation, LDAP channel binding should not be an issue. If you do have issues binding to LDAP with channel binding enabled, follow the remediation steps in ADV190023 to allow LDAP binds from Cloud Volumes Service to succeed.

DNS

Active Directory and Kerberos both have dependencies on DNS for host name to IP/IP to host name resolution. DNS requires port 53 to be open. Cloud Volumes Service does not make any modifications to DNS records, nor does it currently support the use of [dynamic DNS](#) on network interfaces.

You can configure Active Directory DNS to restrict which servers can update DNS records. For more information, see [Secure Windows DNS](#).

Note that resources within a Google project default to using Google Cloud DNS, which isn't connected with Active Directory DNS. Clients using Cloud DNS cannot resolve UNC paths returned by Cloud Volumes Service. Windows clients joined to the Active Directory domain are configured to use Active Directory DNS and can resolve such UNC paths.

To join a client to Active Directory, you must configure its DNS configuration to use Active Directory DNS.

Optionally, you can configure Cloud DNS to forward requests to Active Directory DNS. See [Why can't my client resolve the SMB NetBIOS name?](#) for more information.



Cloud Volumes Service does not currently support DNSSEC and DNS queries are performed in plaintext.

File access auditing

Currently not supported for Cloud Volumes Service.

Antivirus protection

You must perform antivirus scanning in Cloud Volumes Service at the client to a NAS share. There is currently no native antivirus integration with Cloud Volumes Service.

Service operation

The Cloud Volumes Service team manages the backend services in Google Cloud and uses multiple strategies to secure the platform and prevent unwanted access.

Each customer gets their own unique subnet that has access fenced off from other customers by default, and every tenant in Cloud Volumes Service gets their own namespace and VLAN for total data isolation. After a user is authenticated, the Service Delivery Engine (SDE) can only read configuration data specific to that tenant.

Physical security

With proper preapproval, only onsite engineers and NetApp-badged Field Support Engineers (FSEs) have access to the cage and racks for physical work. Storage and network management is not permitted. Only these onsite resources are able to perform hardware maintenance tasks.

For onsite engineers, a ticket is raised for the statement of work (SOW) that includes the rack ID and device location (RU) and all other details are included in the ticket. For NetApp FSEs, a site visitation ticket must be raised with the COLO and the ticket includes the visitor's details, date, and time for auditing purposes. The SOW for the FSE is communicated internally to NetApp.

Operations team

The operations team for Cloud Volumes Service consists of Production Engineering and a Site Reliability Engineer (SRE) for Cloud Volume Services and NetApp Field Support Engineers and Partners for hardware. All operations team members are accredited for work in Google Cloud and detailed records of work are maintained for every ticket raised. In addition, there is a stringent change control and approval process in place to ensure each decision is appropriately scrutinized.

The SRE team manages the control plane and how the data is routed from UI requests to backend hardware and software in Cloud Volumes Service. The SRE team also manages system resources, such as volume and inode maximums. SREs are not allowed to interact with or have access to customer data. SREs also provide coordination with Return Material Authorizations (RMAs), such as new disk or memory replacement requests for the backend hardware.

Customer responsibilities

Customers of Cloud Volumes Service manage their organization's Active Directory and user role management as well as the volume and data operations. Customers can have administrative roles and can delegate

permissions to other end users within the same Google Cloud project using the two predefined roles that NetApp and Google Cloud provide (Administrator and Viewer).

The administrator can peer any VPC within the customer project to Cloud Volumes Service that the customer determines to be appropriate. It is the responsibility of the customer to manage access to their Google Cloud marketplace subscription and to manage the VPCs that have access to the data plane.

Malicious SRE protection

One concern that could arise is how does Cloud Volumes Service protect against scenarios in which there is a malicious SRE or when SRE credentials have been compromised?

Access to the production environment is with a limited number of SRE individuals only. Administrative privileges are further restricted to a handful of experienced administrators. All actions performed by anyone in the Cloud Volumes Service production environment are logged and any anomalies to the baseline or suspicious activities are detected by our security information and event management (SIEM) threat intelligence platform. As a result, malicious actions can be tracked and mitigated before too much damage is done to the Cloud Volumes Service backend.

Volume life cycle

Cloud Volumes Service manages only the objects within the service—not the data within the volumes. Only clients accessing the volumes can manage the data, the ACLs, file owners, and so on. The data in these volumes is encrypted at rest and access is limited to tenants of the Cloud Volumes Service instance.

The volume lifecycle for Cloud Volumes Service is create-update-delete. Volumes retain Snapshot copies of volumes until the volumes are deleted, and only validated Cloud Volumes Service administrators can delete volumes in Cloud Volumes Service. When a volume deletion is requested by an administrator, an additional step of entering the volume name is required to verify the deletion. After a volume is deleted, the volume is gone and cannot be recovered.

In cases where a Cloud Volumes Service contract is terminated, NetApp marks volumes for deletion after a specific time period. Before that time period expires, you can recover volumes at the customer's request.

Certifications

Cloud Volumes Services for Google Cloud is currently certified to ISO/IEC 27001:2013 and ISO/IEC 27018:2019 standards. The service also recently received its SOC2 Type I attestation report. For information about the NetApp commitment to data security and privacy, see [Compliance: Data security and data privacy](#).

GDPR

Our commitments to privacy and compliance with GDPR are available in a number of our [customer contracts](#), such as our [Customer Data Processing Addendum](#), which includes the [Standard Contractual Clauses](#) provided by the European Commission. We also make these commitments in our Privacy Policy, backed by the core values set out in our corporate Code of Conduct.

Additional information and contact information

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

- [Google Cloud documentation for Cloud Volumes Service](#)

<https://cloud.google.com/architecture/partners/netapp-cloud-volumes/>

- Google private service access

https://cloud.google.com/vpc/docs/private-services-access?hl=en_US

- NetApp product documentation

<https://www.netapp.com/support-and-training/documentation/>

- Cryptographic Validation Module Program—NetApp CryptoMod

<https://csrc.nist.gov/projects/cryptographic-module-validation-program/certificate/4144>

- The NetApp Solution for Ransomware

<https://www.netapp.com/pdf.html?item=/media/16716-sb-3938pdf.pdf&v=202093745>

- TR-4616: NFS Kerberos in ONTAP

<https://www.netapp.com/pdf.html?item=/media/19384-tr-4616.pdf>

Contact us

Let us know how we can improve this technical report.

Contact us at doccomments@netapp.com. Include TECHNICAL REPORT 4918 in the subject line.

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.