

Cloud Volumes Service for Google Cloud

Managed SMB Service for Enterprise Applications

SMB Reference Architectures

Niraj Verma, NetApp
November 2019

Abstract

This architecture paper provides a brief overview of SMB services provided by Netapp Cloud Volumes Service for Google Cloud.

Table of Contents

Introduction	3
Customer Problem	3
Customer Value of CVS.....	3
What are Key Values of CVS?.....	3
High performance with low latency	3
Power of NetApp ONTAP.....	3
Increased resilience with snapshots.....	3
Speed up the time to market: spin up cloud volumes in seconds	4
Data durability	4
High availability	4
Security and encryption.....	4
Support for hybrid scenarios	4
Detailed Architectures	6
SMB Service in Hybrid Cloud	
SMB Service in Google Cloud	
The Cost Benefits of CVS	12

Introduction

As organizations aim to run enterprise apps in Google Cloud, or to modernize existing apps and build new ones, the available options for managing SMB requirements in the application's data layer workflows can seem limiting. Factors such as data management, high availability, scaling capacity and performance, and reliability at the right cost can create barriers to successful implementation. Users don't want to compromise on all the application best practices they've learnt over the years by running SMB-based workloads on-premises. Public cloud should simplify the management and complexity of the data layer, while giving end users control and a seamless experience. NetApp and Google shared this mindset when we launched our joint service: NetApp Cloud Volumes Service for Google Cloud. The service is a fully managed offering that supports both NFS and SMB-based workloads. In this document, we'll review how customers can leverage NetApp Cloud Volumes Service, the fastest storage on Google Cloud.

Why Netapp Cloud Volumes Service for SMB workloads in GCP?

The following subsections outline the key benefits of Cloud Volumes Service for Google Cloud Platform.

High performance

With consistently high performance of over 400k IOPS, Cloud Volumes Service provides shared persistent storage with high throughput and low latency. It easily meets the demands of large-scale application environments. Some examples of applications that leverage shared files systems and databases in their workflows are:

- Energy/oil & gas: Petrel application workflows
- Media and entertainment: Rendering and Digital Asset Management workflows
- Retail and finance: Core business applications, CRM, ECM, SAP-based applications
- Analytics and machine learning pipelines

Increased resilience with snapshots

You can easily create a snapshot of an application's volume using NetApp® Snapshot™ technology. Snapshots act as logical backups. They're point-in-time representations of your data that allow you to restore your database or shared file systems without downtime. You can create a snapshot manually or

schedule creation by using the Cloud Volumes Service API, gcloud CLI, or the graphical user interface (GUI).

Snapshots are fast, easy, and nondisruptive. A NetApp snapshot creates a “frozen” read-only view of a volume that enables your applications to access older versions of files and directory hierarchies without additional workflows. Snapshot creation takes only a few seconds (typically less than 1 second), regardless of the size of the volume or the level of activity within the environment.

Faster time-to-market with easy copies

Most organizations need multiple copies of their data for testing and development. Software development pipelines require multiple copies for a variety of uses; creating and refreshing those copies is cumbersome. Typically, creating copies is a time-consuming and tedious process to manage as the number of copies increases. Cloud Volumes Service for Google Cloud enables you to create a near-instant copy using snapshots. It significantly improves the release cycle from development to stage and production. The process can be scripted using Cloud Volumes Service APIs, which leads to a quicker time to market.

High availability

NetApp Cloud Volumes Service for Google Cloud is a regional service. That means application instances in any of the Google Cloud zones in each region can access the same NFS volume, without having to manually create a zonal Linux cluster that you have to manage. The regional design also leads to cost savings (avoiding unnecessary zonal egress charges) and more importantly, the service is unaffected by zonal outages. For protection from a regional outage, you can also choose to replicate the data in a cloud volume in each region to another region of your choice (egress charges will apply).

Security and encryption

NetApp Cloud Volumes Service uses at-rest encryption, relying on the XTS-AES 256-bit encryption algorithm. Cloud Volumes Service encrypts your data without compromising your storage application performance. NetApp manages and rotates encryption keys for you. This single-source solution can

increase your organization's overall compliance with industry and government regulations without compromising user experience.

Zero impact changes: average cost savings of around 70%

By using Cloud Volumes Service for Google Cloud, you can control your cloud performance by dynamically adjusting performance among three service levels. For an increase in performance, you can increase the capacity allocation (for example, 10TB provides 160MB/s); you can also choose a higher service level.

- The **Standard service level** offers economical cloud storage, at just \$0.10 per gibibyte per month. It enables throughput of up to 16 MB/s for each tebibyte allocated. This level is ideal as a low-cost solution for infrequently accessed data.
- The **Premium service level** delivers a good mix of cost and performance. At a cost of \$0.20 per gibibyte per month, it offers four times the performance of the Standard level, with 64 MB/s for each tebibyte allocated. This is a good fit for many applications where data capacity and performance needs are balanced.
- The **Extreme service level** provides the best performance. At a cost of \$0.30 per gibibyte per month, it enables up to 128 MB/s for each tebibyte allocated, and Cloud Volumes Service can scale to deliver several GB/s for reads and writes. The Extreme service level is the best fit for high-performance workloads.

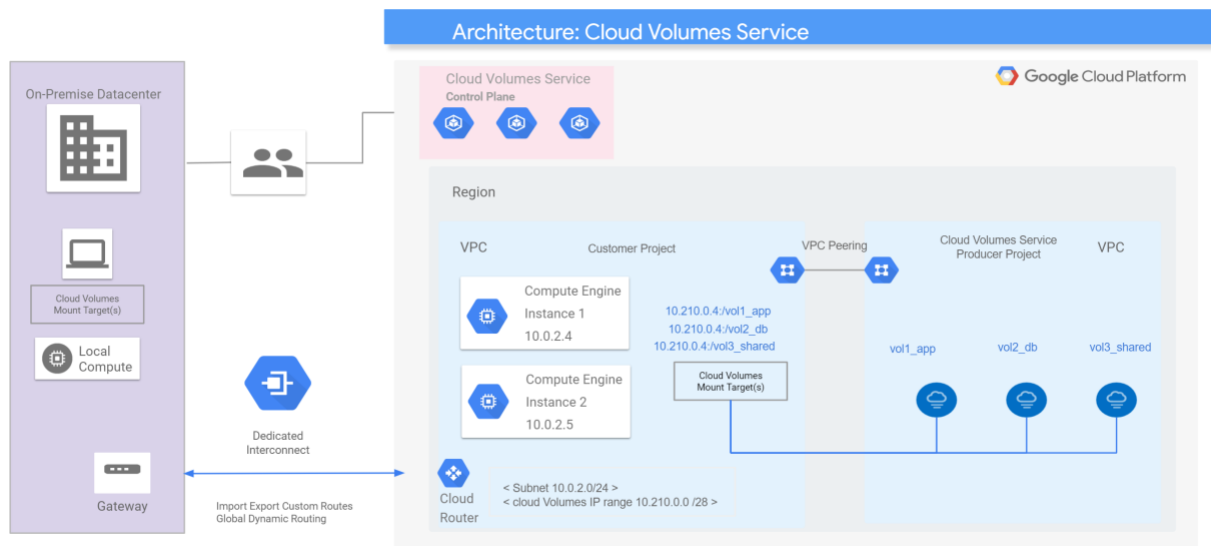
One of the unique features of NetApp Cloud Volumes Service for Google Cloud is the capability to change performance on demand without impacting the availability to the application or users. Thus, if users need Extreme performance for just two hours a day and Standard performance for the remainder, you can use API calls or a scheduler in Linux to automate that process, resulting in significant cost savings at scale.

SMB Solution Architectures

SMB service in a hybrid cloud

Often, organizations require hybrid application topology, spanning on-premises application clusters and instances in Google Compute Engine to access the same data set.

A client in an on-premises data center can access an SMB cloud volume with a Dedicated Interconnect, as depicted in the reference architecture below. Global Dynamic Routing needs to be enabled on the VPC for on-premises access, allowing the Cloud Routers to import and export custom routes that can be added.



In the above diagram, you can see that we have multiple VM application instances in Google Compute Engine, VM Instance-1 and VM Instance-2

VM instances							
CREATE INSTANCE IMPORT VM REFRESH START STOP RESET DELETE							
Filter VM instances							
Name	Zone	Recommendation	In use by	Internal IP	External IP	Connect	
<input type="checkbox"/> instance-1	us-central1-a			10.128.0.3 (nic0)	35.232.92.174	RDP	⋮
<input type="checkbox"/> instance-2	us-central1-a			10.128.0.4 (nic0)	35.223.175.191	RDP	⋮

In NetApp's producer project and VPC, we have three application volumes that were created to serve database needs, core application needs, and shared files repository.

Volumes [+ CREATE](#) [🗑️ DELETE](#)

Quick reference for Cloud Volumes [Private Service Access](#) [API](#) [Shared VPC support](#) [Granular permissions](#)

🔍 Search for volumes by name, ID, region, etc.

<input type="checkbox"/>	Name	Region	Life cycle	Service level	Mount Targets	Allocated capacity	Protocol types	Show More
<input type="checkbox"/>	🟢 nfs-shared-vol1	us-central1	available	premium	10.14.133.4/confident-cocky-lewin	1,024 GiB	NFS	⋮
<input type="checkbox"/>	🟢 smb-shared-vol	us-central1	available	premium		1,024 GiB	SMB	⋮

Both application instances can share the same SMB volume , which is managed by NetApp and Google teams, and users have control over the service level and capacity, both of which can be changed on demand without impacting the workflow.

Users can easily mount the volumes to application instances by copying and pasting the instructions available under “show more” in the screenshot above.

SMB service in Google Cloud

Users can access SMB-based NetApp cloud volumes that exist in the Google cloud landscape, even across multiple regions, in a standalone or shared VPC topology.

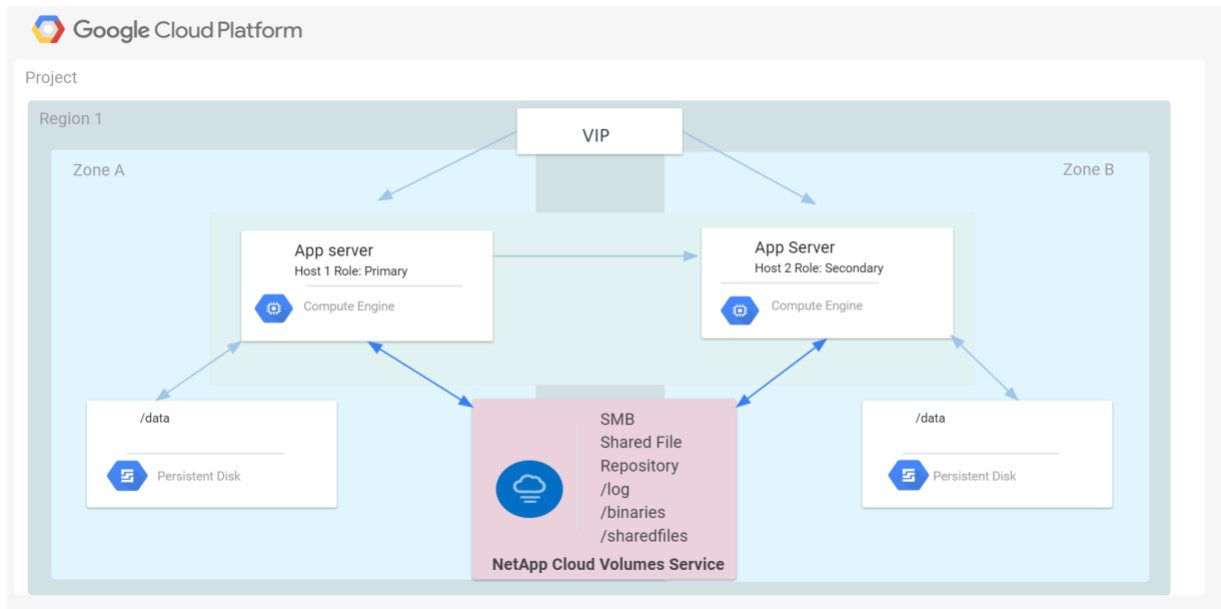
During the volume creation workflow, users must enable the shared VPC topology and select their host project. Compute resources in the service projects can then access the cloud volumes' mount targets through the host project.

- Shared VPC configuration
Provide the host project name when deploying in a shared VPC service project.

Host Project Name

Note: For production and business critical workflows, we recommend minimizing cross-region or on-premises access to avoid increased latency.

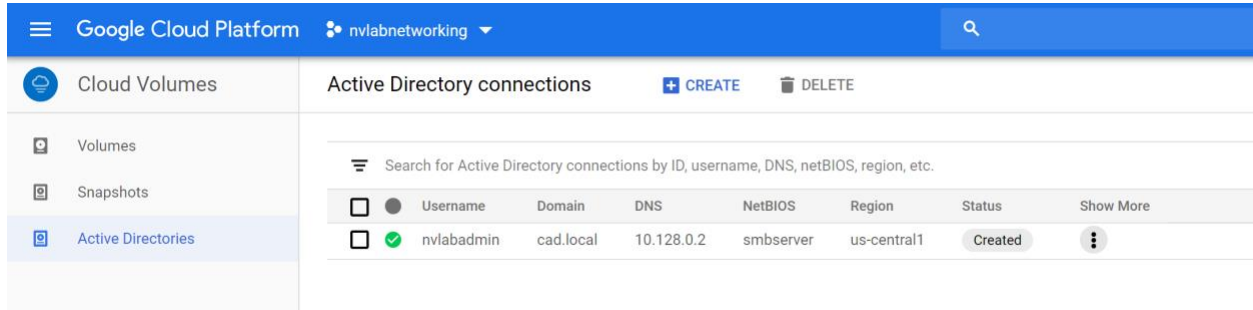
In the below example, you can see how multiple application instances can access shared binaries and log files in a secure, centralized, and cost-effective manner, running on Cloud Volumes (SMB)



When deploying SMB-based applications, users can take advantage of the following features for increased security and automation.

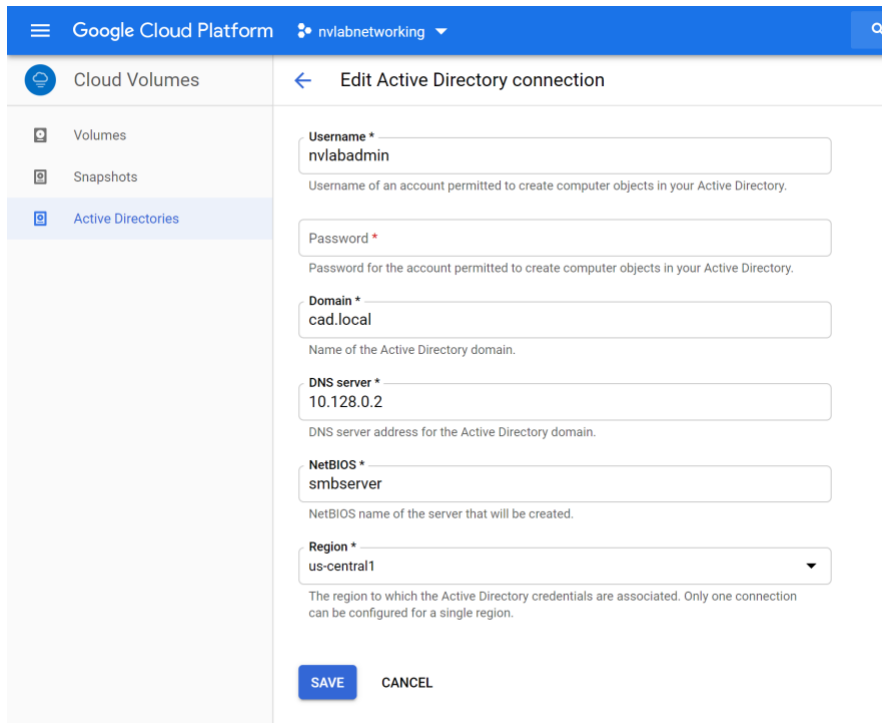
Active Directory Integration

Customers can leverage the built-in active directory (AD) integration to securely create SMB volumes. A one-time connection (per region, per project) needs to be setup with the AD server, using administrator credentials as shown below.



The screenshot shows the Google Cloud Platform console interface. The top navigation bar includes the Google Cloud Platform logo, the project name 'nvlabnetworking', and a search icon. The left sidebar shows the 'Cloud Volumes' menu with options for 'Volumes', 'Snapshots', and 'Active Directories'. The main content area is titled 'Active Directory connections' and includes '+ CREATE' and 'DELETE' buttons. Below this is a search bar and a table of connections.

	Username	Domain	DNS	NetBIOS	Region	Status	Show More
<input type="checkbox"/>	nvlabadmin	cad.local	10.128.0.2	smbserver	us-central1	Created	⋮



The screenshot shows the 'Edit Active Directory connection' form in the Google Cloud Platform console. The form fields are as follows:

- Username ***: nvlabadmin
Username of an account permitted to create computer objects in your Active Directory.
- Password ***: [Redacted]
Password for the account permitted to create computer objects in your Active Directory.
- Domain ***: cad.local
Name of the Active Directory domain.
- DNS server ***: 10.128.0.2
DNS server address for the Active Directory domain.
- NetBIOS ***: smbserver
NetBIOS name of the server that will be created.
- Region ***: us-central1
The region to which the Active Directory credentials are associated. Only one connection can be configured for a single region.

At the bottom of the form are 'SAVE' and 'CANCEL' buttons.

Global, user-accessible API

Customers can now interact with Cloud Volumes Service via API and programmatically list, create, update, and delete objects (volumes, snapshots, for example), meaning customers can provision cloud volumes as part of an automated application-deployment package, create application-consistent snapshots, and use other automation frameworks.

The API is global, which means customers can interact with a single endpoint and specify regions as location parameters in the URL. Authentication and authorization to access the Cloud Volumes Service API is integrated with GCP service accounts and IAM. See the [documentation](#) for more details.

Granular roles

To increase security for users and to enable automation via service accounts, Cloud Volumes Service supports two granular roles: `netappcloudvolumes.admin` and `netappcloudvolumes.viewer`.

These roles can be assigned to users or service accounts to perform Cloud Volumes Service actions. See the [documentation](#) for details about how to use these roles.

The screenshot shows the Google Cloud IAM console interface for adding members to a project named "ncv-demo".

- Page Title:** Add members to "ncv-demo"
- Section Header:** Add members, roles to "ncv-demo" project
- Instructions:** Enter one or more members below. Then select a role for these members to grant them access to your resources. Multiple roles allowed. [Learn more](#)
- New members:** A text input field containing the email address `nirajv@netapp.com` with a remove icon (x) and a help icon (?).
- Select a role:** A dropdown menu is open, showing a list of roles. The roles are grouped by category:
 - Service networking:** NetApp Cloud Volumes Admin, NetApp Cloud Volumes Viewer
 - Service Usage:**
 - Source:**
 - Stackdriver:** Stackdriver Debugger, Stackdriver Profiler
 - Storage:**
 - Third-party Partner:**
- MANAGE ROLES:** A button at the bottom of the role selection dropdown.

How does Cloud Volumes Compare to Google Cloud Offerings ?

	NetApp Cloud Volumes	Google Cloud Filestore	Elastifile on Google Cloud
Summary	Fully managed service Integrated directly into the GCP UI. Service is GA as of Nov. 20, 2019 Designed for scale-up workloads	Physical fully managed service in GA	NOTE: This service has been purchased by Google and will have its features integrated into Filestore Built for scale-out workloads
Protocols	NFSv3, SMB	NFSv3	NFSv3, SMB
Data Management <ul style="list-style-type: none"> • Replication • Copy • Snapshots • Data sync • Integrated backup 	Yes <ul style="list-style-type: none"> • Replication: No (coming soon) • Copy: Yes • Snapshots: Yes • Data sync: Yes • Integrated backup: No 	No	Some <ul style="list-style-type: none"> • Replication: No • Copy: No • Snapshots: Yes • Data sync: Yes • Integrated backup: No
High/extreme performance, low latency	Yes Standard (16MB/s) Premium (64MB/s) Extreme (128MB/s) change on demand, scales with capacity	Different clusters for different performance levels – fixed	High performance – no extreme level. Two levels: Standard, SSD.
Availability	AZ failure tolerant 99.99% (depends on region)	No	99.9%
Scale	100 TB – can change on demand Future multi PBs	Max 63.9 TB	Scales to Multi-PB
Billing/Support	Google Cloud	Google Cloud	Product will soon end its life after GC acquisition announcement.

The Cost Benefits of NetApp Cloud Volumes Service

In addition to obviating the need for infrastructure you need to manage, NetApp cloud volumes is also the lowest cost, highest quality storage solution for SMB volumes in Google Cloud. Instead of deploying and managing virtual machines, storage devices and operation systems, you can simply consume volumes. In addition, three performance tiers allow you to optimize for your workload and spending requirements. Choose from Standard, Premium, and Extreme tiers.

The table below compares using NetApp Cloud Volumes to renting virtual machines and running an open source Linux distribution. The “roll your own” shows the necessary doubling of the infrastructure in order to match the availability offered by NetApp.

Requirement	NetApp Cloud Volumes	“Roll your own HA File Server”
Capacity	50TiB	50TiB
Monthly Storage Costs	\$5,120 (50TiB x Standard Service Level @ \$0.10 per gibibyte/month snapshots included)	\$17,408 (50Tib x Persistent Disk – Regional Provisioned SSD, snapshots not included)
Annual Storage Costs	\$61,440	\$ 208,896
Compute/VM Costs Monthly	\$0	\$ 401.78 (2x n1-standard-4 with 2x 375gb local SSD)
Compute/VM Costs Annual	\$0	\$4,821.36
TCO (yearly)	\$61,440	\$213,717.36
Savings	~70% savings	

