

## Solutions

### [Deploying SAP on Google Cloud with NetApp Cloud Volumes Service](#)

Niraj Verma, Cloud Solution Architect NetApp

This architecture brief provides an overview of SAP databases and gives a reference architecture for deploying SAP workloads on Cloud Volumes Service for Google Cloud.

### Introduction

Some of the largest businesses run SAP applications. Increasingly, customers are running mission-critical SAP applications on Google Cloud to take advantage of Google Cloud's secure, modern, and reliable global infrastructure as well as its capabilities in AI, ML, and analytics.

Every enterprise customer is different, and each has its unique journey and challenges to the cloud when it comes to migrating SAP applications. NetApp and Google have teamed up to simplify this journey by offering customers best practices that are tied to reference architectures underpinned by [NetApp Cloud Volumes Service for Google Cloud](#). As customers look to transform their SAP landscapes, NetApp Cloud Volumes Service will help them modernize and create experiences using the highly secure and highly available Google Cloud infrastructure.

With better and more reliable hardware, simple upgrades, and easy migration, the promise of VM-based agility for their SAP data is attractive to many customers. Google also offers many integrated services for collaborating, analyzing, and innovating with the SAP data to drive new business value, all while keeping downtime to a minimum.

This document provides a brief overview of SAP use cases targeted for Google Cloud. It also gives a reference architecture for deploying SAP on NetApp Cloud Volumes Service for Google Cloud Platform and explains the added benefits and value.

## SAP challenges in the cloud

Many customers that have on-premises SAP deployments are looking to simplify operations and move to cloud faster. With a hybrid SAP landscape, there needs to be a seamless way for data management across the landscape. Customers are also looking to minimize and eliminate extraneous on-premises infrastructure used for sporadic, non-production requirements such as development, quality of service, and training. In parallel, new SAP customers are looking for virtualized SAP deployments in the cloud to avoid distributed landscapes (on-premises and cloud).

For scenarios of on-premises, existing hybrid, or new SAP deployments, customers need to think about their requirements for no-impact testing, training, and prototype development. They want to integrate with Google Cloud services for a consistent experience without sacrificing performance or data management capabilities.

## Why NetApp Cloud Volumes Service for SAP

NetApp Cloud Volumes Service for Google Cloud is great for customers who are looking to run SAP databases in the cloud. Instead of waiting for months or years to refactor an SAP-based application, customers can easily migrate their data to Cloud Volumes Service on Google Cloud Platform.

Cloud Volumes Service for Google Cloud is an enterprise-class storage offering that is running on the same technology proven for SAP deployments on-premises. It is available as a cloud-native offering. Multiple Fortune 100 enterprise SAP deployments (development, QAS, and PRD landscapes) are running on Cloud Volumes Service for Google Cloud. With high performance and near instant copies, customers can simplify operations, resulting in up to 50% cost savings and accelerating time-to-market by 70%.

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

### High performance

With consistently high performance of over 250k IOPS, Cloud Volumes Service provides shared persistent storage with high throughput and low latency. It easily meets the demands of large-scale SAP deployments.

### Increased resilience with snapshot copies

You can easily create a snapshot copy of an SAP database using NetApp™ Snapshot technology. Snapshot copies 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 snapshot copies manually or schedule creation by using the Cloud Volumes Service API, gcloud CLI, or graphical user interface (GUI).

Snapshot copies are fast, easy, and nondisruptive. A NetApp snapshot copy 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 copy 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. Since they are read-only, incremental copies, you only pay for the space consumed by new data written.

### Faster time-to-market with snapshot copy-based environments

Most organizations need multiple copies of data for testing and development. SAP landscapes are littered with system copies for a variety of uses; creating and refreshing those copies is cumbersome. Typically, creating copies of SAP landscapes is a time-consuming and tedious process. 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. It means that application instances in any of the Google Cloud zones in a given region can access the SAP database and shared files. More importantly, the service is unaffected by zonal outages. You don't have to replicate content and incur additional charges for data egress either.

## 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 the user experience.

## Zero impact changes: average cost savings of around 70%

By leveraging Cloud Volumes Service for Google Cloud, you can control your cloud performance by dynamically adjusting 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 gibibytes 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.

## SAP storage classification

To understand the specific storage requirements and different use cases for SAP, we need to take a closer look at the different types of SAP applications, as shown in Figure 1.

Figure 1 - SAP storage performance requirements and certification

	SAP Shared File Systems	SAP AnyDB	SAP HANA
	/sapmnt	binaires	/hana/shared
	/usr/sap/trans	data	/data
	/usr/sap/SID	log	/log
Certification	None	None	Required (For Production)
KPI	None	Based on Sizing	Defined by SAP ( HWCCT)

Performance Requirement	Low	Medium	High

Figure 1 shows typical file systems that are used by an SAP application server, an SAP/AnyDB database server (such as Oracle, SQL, DB2, ASE, and MaxDB), and an SAP HANA server. The colors indicate the typical performance requirement for these type of file systems, showing a low performance requirement for standard file systems that hold binaries, log and trace files, or configuration files using a green color. In contrast, SAP HANA data and log volumes require higher storage performance; therefore, they are displayed in red when compared against traditional databases.

When you are running an SAP system in production, an important thing to remember is the SAP HANA certification or a formal support declaration for AnyDB databases, such as Oracle. In contrast, file systems used for SAP application servers do not require a certification. The SAP HANA certification includes the host and storage on which SAP HANA is running in production. Additionally, the certification includes a performance test with stringent key performance indicators (KPIs) that a system must pass to be listed in the Certified and Supported SAP HANA Hardware Directory.

## SAP use cases

Based on the classification described in the previous section, there are various use cases where NetApp Cloud Volumes Service for Google Cloud could help improve customers' SAP cloud experience:

### Category 1: SAP shared file systems

### Category 2: SAP Non-production databases (including SAP HANA)

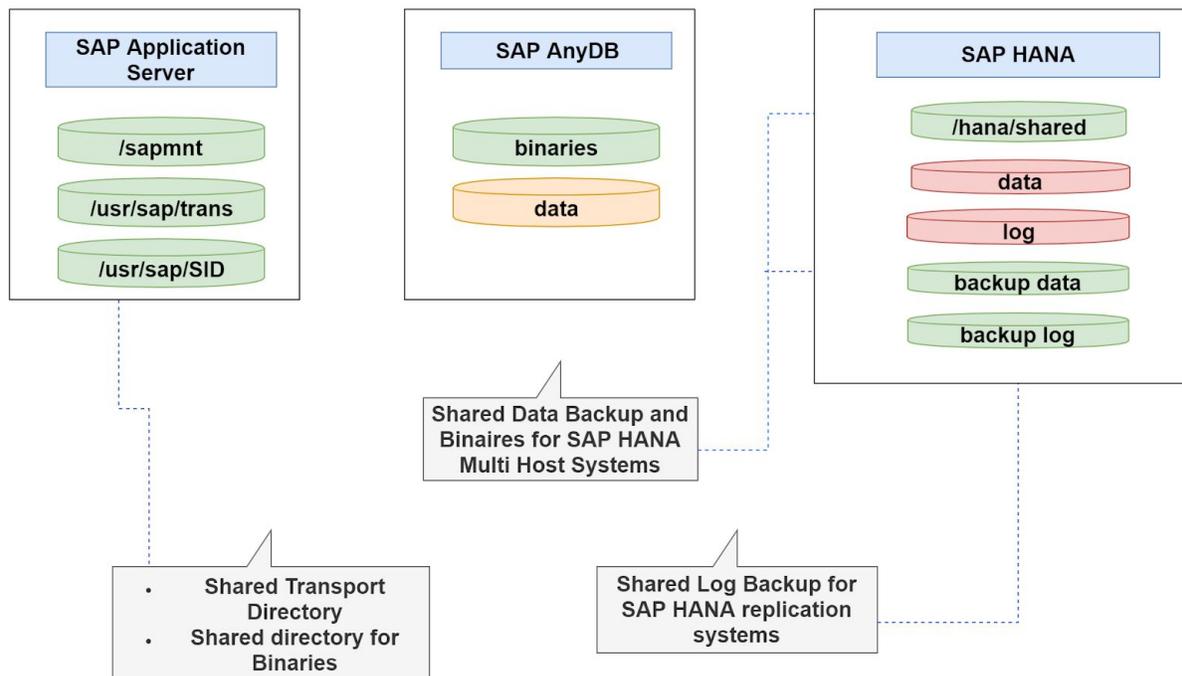
- SAP SBX (sandbox)
- SAP QAS (quality assurance) systems
- SAP training systems

Although the four use cases listed above are unique, the rest of this document will refer to them as Use Case **Category 1: SAP shared file systems**, and Use Case **Category 2: Non-production databases** (SBX, QAS, Training).

### Category 1: SAP shared file systems

As shown in Figure 2, shared files are needed for almost any SAP system landscape, starting with the typical candidates such as /usr/sap/trans, SAP HANA data volumes, and SAP HANA log volumes. The exact performance requirements can vary based on a customer's setup; however, you can find shared file systems in any customer environment, mostly using NFS as the protocol.

Figure 2 - SAP Shared Files requirements



Almost every SAP landscape requires a shared file system provided through the NFS protocol on UNIX, the Linux operating systems, or the SMB protocol in the case of an SAP system running on Windows operating system.

This section discusses the shared file requirements for an example setup using an SAP HANA multiple-host system with SAP HANA system replication, an SAP application server using clustered SAP ABAP SAP Central Services (ASCS), and an SAP Enqueue Replication Server (ERS). This is a very common setup to achieve high availability for the SAP HANA database by using HANA system replication, as well as for the application server by implementing clustered ASCS and SAP ERS instances.

Figure 3 - Shared Files for SAP

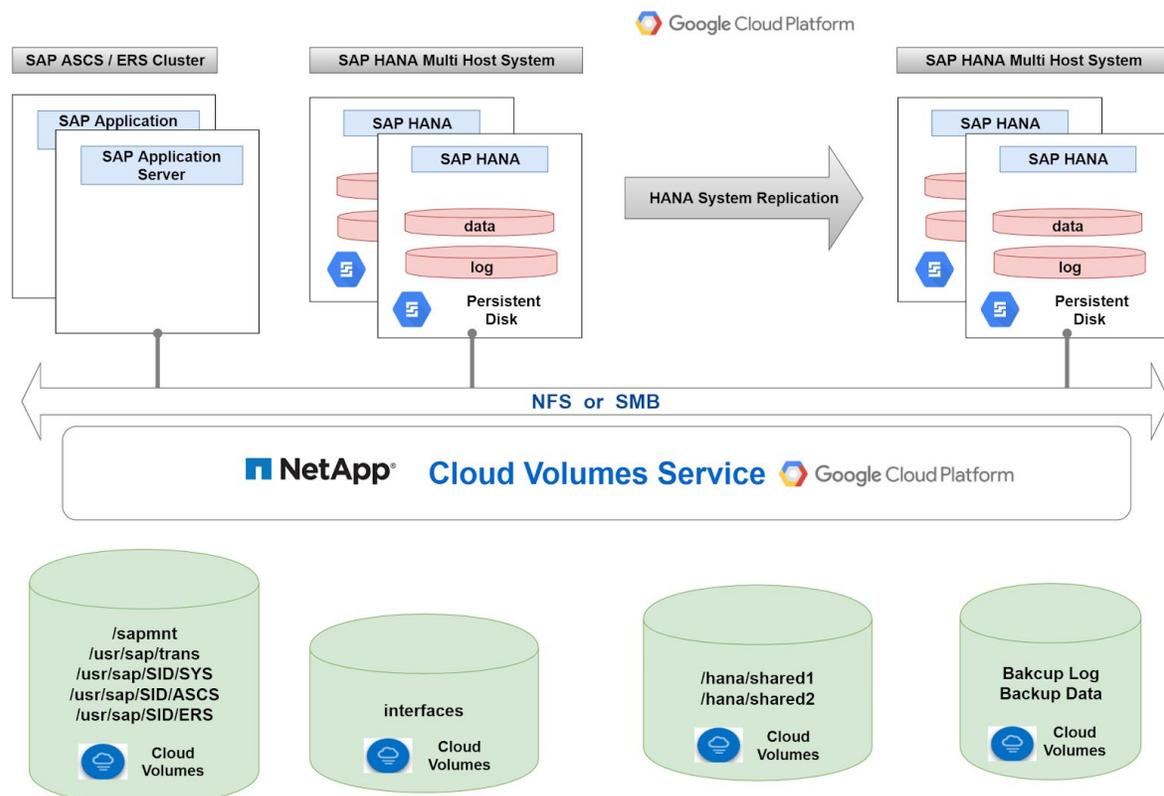


Figure 3 shows the following shared file systems that are required for the system landscape:

- /sapmnt:**  
 If you have more than one application instance, use the /sapmnt file system to store a common set of binaries and configuration files. The I/O pattern is reading the binaries and configuration files and writing view logs. In Figure 6, green indicates a lower performance requirement.

- **/usr/sap/trans:**  
This is a common file system that is used to share (or transport) customer developments or other transports between systems in a single SAP landscape.
- **/usr/sap/<SID>/SYS, /usr/sap/<SID>/ASCS, and /usr/sap/<SID>/ERS:**  
These file systems are used for the SAP application server instances. The performance requirements are rather low. However, for a high-available setup with the ERS, it is mandatory that the underlying file system be a high-available system also, so that the ERS locking table is preserved in case of an instance failover.
- **/hana/shared:**  
For a multiple-host SAP HANA system, /hana/shared must be an NFS shared file system.
- **Backup data:**  
For file-based backups in a multiple-host environment, all SAP HANA servers should have access to the backups, which requires an NFS share. File shares used for file-based backups require a significantly higher throughput than the previously discussed file systems to allow the backup of the SAP HANA database to finish as quickly as possible. This requirement can be partly mitigated using snapshot copies.
- **Backup log:**  
The automatic SAP HANA log backup is written to this shared location. For an SAP HANA system replication setup, the location must be a shared location between both SAP HANA systems to allow for failover. Even in an SAP HANA single-host setup, this location should still be highly available. Also, for the log backup, a medium performance requirement is recommended.

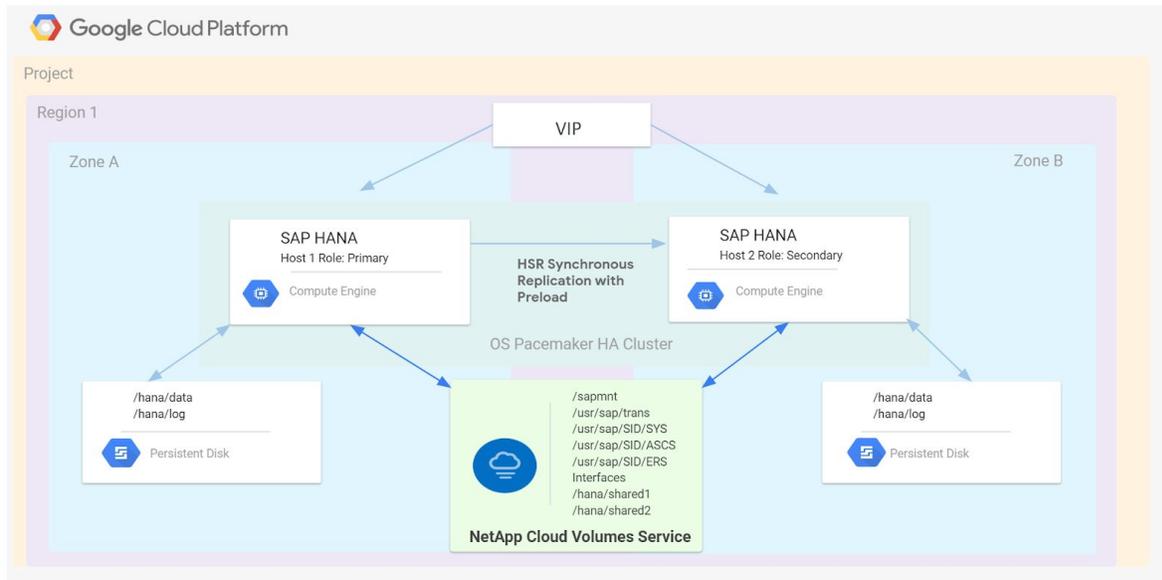
#### SAP shared file system detailed architecture

A do-it-yourself Linux cluster using a block device such as Persistent Disks is one of the solutions within Google Cloud that provides a highly available NFS service. At first glance, this solution is appealing due to the localization of files to the application instance. However, this type of solution can quickly present the following challenges when it comes to scaling:

- The level of manual administration that the deployments require.  
For example, allocating a new file system requires allocating new storage, mounting it to the compute hosts that will serve the data, and potentially initializing the new share with existing data. If the file system needs to grow, the growth must be handled manually. If the performance of the underlying disks needs to be upgraded, the allocation of the new storage and migration of existing files need to be handled with minimum downtime.
- The complexity of managing the storage over time as the deployment grows.  
Storage administrators working with production file shares need to maintain uninterrupted access to the files, provide backup or snapshot facilities, allow test copies of the data to be created, and much more. Not all storage, SAP, or application administrators have the skills to maintain and administer a Linux cluster. Providing robust support for this kind of functionality requires high-level technical expertise.

Figure 4 shows NetApp recommended architecture for SAP shared file systems using Cloud Volumes Service for Google Cloud in a multi-host HANA deployment:

Figure 4 – Reference Architecture for SAP Multi-Host HANA Shared File system on Cloud Volumes Service for Google Cloud



By running the high-performance data and log stores locally on a persistent disk, and remainder shared file systems on Cloud Volumes Service for Google Cloud, customers can gain several advantages that will strengthen their mission-critical or enterprise-scale SAP landscapes:

- Eliminate costs associated with running Linux Instance clusters
- Significantly reduce complexity and risk associated with administration
- Accelerate deployments and time-to-market
- Design once, with easy long-term growth and scale
- Expose “shared files” to other services and workflows by using the Cloud Volumes Service API
- Easily protect, copy, and sync shared file systems

Customers have the choice to host shared file systems by performance and protection needs with the matching service level (Standard, Premium, or Extreme). The service level can also be changed on the fly with automation (REST APIs) to respond to a business process (such as a seasonal peak) or sustained growth, as discussed earlier.

### Data protection for SAP shared files

In addition to performance, ease of management, and flexibility, SAP customers also require enterprise-grade data protection not only for their databases, but also for their shared file systems. By using NetApp Cloud Volumes Service for Google Cloud, customers can implement comprehensive data protection for their SAP shared file systems by using NetApp snapshot copies.

Cloud Volumes Service for Google Cloud snapshot copies are very fast (completed in seconds), space efficient, and have no performance impact, regardless of the size of the file system. Instead of copying data, the service marks the blocks on the active file system to be part of the new snapshot copy. It

ensures that whenever a block is changed, the new content is written to an empty block, preserving the snapped data block and avoiding any additional I/O. In other words, snapshot copies are pointers to data blocks that also allow restore operations to be fast because only pointers are changed and no data will be copied.

Customers can easily create storage snapshot copies by using the Google Cloud console, or they can integrate snapshot functionality into their own scripts by using REST API. Using the console, you can schedule snapshot copies to include automated retention management. For example, customers can schedule hourly, daily, weekly, and monthly snapshot copies, and configure the number of snapshot copies they want to keep. These snapshot copies can be used to quickly restore files. They can also be used to create clone copies of the file systems in the same region for development, training, or QAS, or to create a new volume in another region for added protection.

### Snapshot Policy

Allow automatic snapshots.  
When enabled, snapshots are automatically created according to below specified schedules.

HOURLY   
  DAILY   
  WEEKLY   
  MONTHLY

Snapshots To Keep

Minute

Google Cloud Platform		CV-Product-Mgmt		Snapshots <span>BETA</span>				<a href="#">+ CREATE</a>	<a href="#">DEDELETE</a>
<input type="checkbox"/>	demo-snapshot	us-east4	available	a5ab74d5-a782-3143-b72e-16c2dc3de081	0 GiB				
<input type="checkbox"/>	hana-snapshot	us-east4	available	a5ab74d5-a782-3143-b72e-16c2dc3de081	3.29 GiB				
<input type="checkbox"/>	hana-snapshot	us-east4	available	bc791d78-6ad9-f825-ebbc-8a2aba6527ee	3.29 GiB				
<input type="checkbox"/>	hana-snapshot	us-east4	available	c1bb8b6f-bad6-f2c5-11c9-33de2b04dd9d	3.29 GiB				
<input type="checkbox"/>	practice-snapshot	us-east4	available	b91a7127-398d-5442-00ca-77ab4b079ed0	2.29 GiB				
<input type="checkbox"/>	practice-snapshot	us-east4	available	d0f9a5d6-2e2d-f8fc-a693-b041b38bbfed	2.29 GiB				
<input type="checkbox"/>	monthly-on-day-1-0-min-past-4am.2019-03-01_0400	us-east4	available	e6826162-d4a1-fb74-192b-13d55916d87f	0 GiB				
<input type="checkbox"/>	confluencedemo_snap2	us-east4	available	a46b5355-1561-8f71-1e98-0c3940f4a323	0 GiB				
<input type="checkbox"/>	monthly-on-day-1-0-min-past-4am.2019-04-01_0400	us-east4	available	e6826162-d4a1-fb74-192b-13d55916d87f	0 GiB				

## Category 2: Non-production databases (including SAP HANA)

This category includes the following systems:

- SAP SBX systems
- SAP QAS (quality assurance) systems
- SAP training systems

In this second category of use cases, customers can run database volumes (data and log) on NetApp Cloud Volumes Service for Google Cloud along with shared files (separate volumes).

The advantages of running SAP databases on Cloud Volumes Service for Google Cloud include:

- No need to deploy and manage individual persistent disks for each HANA instance in a single or multi-host deployment, resulting in significantly lower management overhead
- Cloud Volumes Service is the only enterprise-grade, fully managed, cloud-native offering for SAP
- Cloud Volumes Service is running on the same NetApp operating system that has over two decades of innovation with SAP
- Cloud Volumes Service is proven for SAP, and the most stable NFS and SMB offering within Google Cloud
- Cloud Volumes Service offers the highest performance (~128 Mb/s for the Extreme tier) for SAP databases within Google Cloud
- Unified data protection for database and log volumes

### SAP HANA non-production databases

The most basic SAP HANA configuration requires three different storage volumes—data, log, and /hana/shared.

#### *SAP HANA single-host system*

The following storage volumes are required for a single-host HANA system:

- **/hana/shared**  
This storage volume is used for shared files such as binary, logs, and configuration. There are no specific storage requirements for the shared file system for a single-host SAP HANA system. It can be either local attached storage or NFS mounted storage. In the case of a multiple-host SAP HANA deployment, the shared file system must be a shared NFS mounted file system.
- **Data volume**  
The SAP HANA data volume must have memory size allocated for the SAP HANA database, such as the memory of the SAP HANA VM in most of the setups. For the data and the log volume, SAP defines performance criteria and specific certification KPIs. For a single-host system, the data volume can be local storage or an NFS mounted storage. For a multiple-host SAP HANA system, an external enterprise storage is required. In public cloud environments, a shared NFS volume is the only viable option.

- **Log volume**

SAP HANA log volume is required to persist the most recent redo logs. SAP recommends that the size be 50% of the HANA memory, with a maximum size of 0.5 TB. For a single-host system, the log volume can be local storage or NFS mounted storage. For a multiple-host SAP HANA system, a shared NFS volume is required to enable SAP HANA multiple-host cluster functionality.

### *Multiple-Host SAP HANA*

When the memory size of a single server is not sufficient, SAP HANA allows you to combine the memory of multiple servers to run SAP HANA in a multiple-host configuration. In this setup, SAP HANA allows you to configure a standby host that takes over the role of the failed server if a host failure occurs on one of the worker hosts (see Figure 5). This SAP HANA cluster mechanism requires that the /hana/shared file system is available on all SAP HANA nodes, which in turn requires NFS. Also, the data and log volumes must be remounted on the standby host, which is a perfect use case for an NFS endpoint, such as NetApp Cloud Volumes Service for Google Cloud.

In addition to these changes, SAP HANA multiple-host setup requires that one of the following common backup volumes be shared between all SAP HANA hosts for a file-based backup:

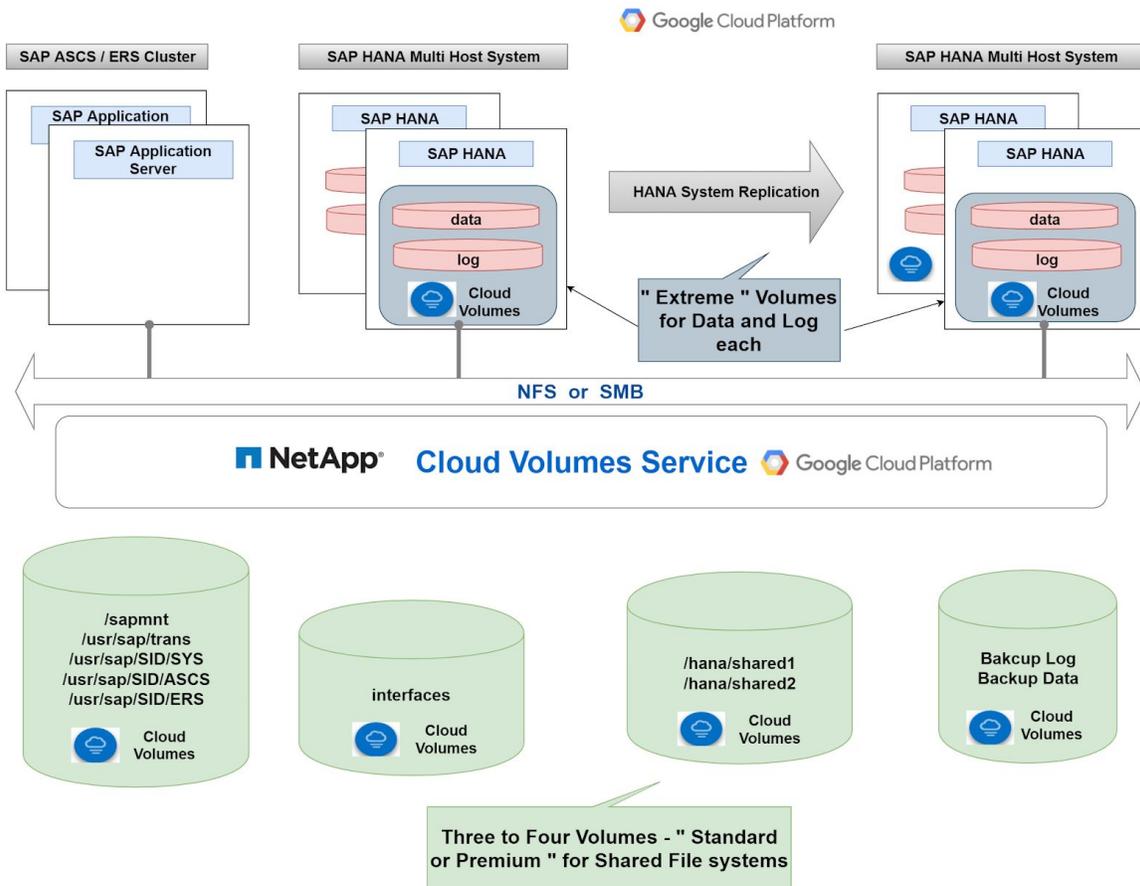
- **Backup volume -- data backup**

Depending on how many backup sets must be kept, SAP recommends that, for each set, the size of the SAP HANA data area should be allocated at the volume level.

- **Backup volume -- log backup**

SAP HANA automatically archives the redo logs from the log volume to a shared file system that all hosts must have access to. Many customers use the same volume for storing the file-based data backups as well as the automated log backups. The sizing depends on the change rate within the SAP HANA database.

*Figure 5 - Multi Host SAP Development / Non-Production Database on Cloud Volumes Service*



## Data protection for SAP HANA non-production databases

As discussed in the *Data protection for SAP shared files* section, Cloud Volumes Service offers integrated, fast, and easy snapshot copies that can be created on-demand or according to a predefined schedule. SAP architects and administrators can orchestrate application-consistent snapshots by first opening a native snapshot in HANA, and then creating a snapshot copy using the gcloud CLI or the Cloud Volumes Service REST API, and then closing the HANA snapshot to complete the process.

In the next update of this document, we will be adding sample code and illustrations to highlight how easy and powerful this process can be.

For more information about Cloud Volumes Service, technical architecture, and API specifications, see the [Cloud Volumes Service for Google Cloud user documentation](#).

## Summary

Cloud Volumes Service on Google Cloud gives customers a robust set of capabilities for migrating or deploying new SAP workloads. Cloud Volumes Service allows for an agile, future proof design by eliminating any performance or scalability challenges with zero impact, and automated and/or controlled service level changes.

Cloud Volumes Service is highly available, protects against data corruption and data loss, and features granular roles that are built into Google Cloud Identity and Access management, making it a secure choice for the most demanding SAP workloads.

To learn more about NetApp Cloud Portfolio, visit [cloud.netapp.com/cloud-volumes-service-for-gcp](https://cloud.netapp.com/cloud-volumes-service-for-gcp).