

This document is for people who are evaluating Google Cloud Platform (GCP) as a platform for deploying SAP Business Suite on SAP HANA, especially for people in the following kinds of jobs:

- SAP technical architect
- Cloud architect
- SAP basis administrator
- Enterprise architect

This document also lists issues to consider prior to installation, as well as linking to SAP Notes and other documentation to facilitate deployment.

SAP Business Suite on SAP HANA is a suite of SAP systems, including SAP ECC, running on an SAP HANA database. ECC, or Enterprise Core Component, is an application suite that contains core business functions like finance, logistics, warehouse management, and sales and distribution. ECC was designed to run on various databases, including Sybase ASE, MS SQL Server, Oracle, and now SAP HANA. This document discusses architecture and deployment of SAP Business Suite on the SAP HANA database only.

Google Cloud Platform provides a cost-effective, reliable, secure, and high-performance SAP certified infrastructure for running SAP Business Suite on SAP HANA. For a complete list of supported SAP solutions on GCP, see [SAP on Google Cloud \(/sap/\)](#).

If you're an SAP customer, you can use your existing SAP HANA and other SAP application licenses to deploy SAP Business Suite on GCP under a bring-your-own-license (BYOL) model. GCP supports the BYOL model for both production and non-production use cases. Operating

system licenses are included in Compute Engine prices; alternatively, you can bring your own OS [image](/compute/docs/images#custom_images) (/compute/docs/images#custom_images) and licenses as well.

Several sizing options are available based on the implementation type. For [greenfield](https://en.wikipedia.org/wiki/Greenfield_project) (https://en.wikipedia.org/wiki/Greenfield_project) implementations, we recommend using the SAP Quick Sizer tool. For detailed information, see SAP's [Sizing](https://www.sap.com/about/benchmark/sizing.html) (https://www.sap.com/about/benchmark/sizing.html) page. SAP also provides T-shirt guides for specific solutions and tools for migrating current on-premises solutions to GCP. (For example, see [Find Certified IaaS Platforms](https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/iaas.html#categories=Google%20Cloud%20Platform) (https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/iaas.html#categories=Google%20Cloud%20Platform) and [SAP Note 2456432 - SAP Applications on Google Cloud Platform: Supported Products and Google VM types](https://launchpad.support.sap.com/#/notes/2456432) (https://launchpad.support.sap.com/#/notes/2456432).) SAP and GCP use different units to measure IOPS (input/output operations per second); consult your SI (Systems Integrator) partner to convert SAP sizing requirements into appropriately sized GCP infrastructure.

GCP offers Compute Engine instance types that are certified by SAP to match sizing requirements when you're deploying SAP HANA. For more information about sizing on GCP and the machine types that are supported, see the following pages:

- [Supported machine types for SAP HANA on GCP](/solutions/sap/docs/sap-hana-planning-guide#vm_types) (/solutions/sap/docs/sap-hana-planning-guide#vm_types).
- [Supported machine types for SAP NetWeaver on GCP](/solutions/sap/docs/netweaver-planning-guide#machine_types) (/solutions/sap/docs/netweaver-planning-guide#machine_types).

Custom machine types for SAP HANA on GCP are also certified by SAP. You can run SAP HANA instances with less than 64 vCPUs as long as you maintain a vCPU/Memory ratio of at least 6.5. To determine [SAPS](https://www.sap.com/about/benchmark/measuring.html) (https://www.sap.com/about/benchmark/measuring.html) provided by your chosen CPU architecture, see [SAP SD Standard Application Benchmark Results, Two-Tier Internet Configuration](http://global.sap.com/solutions/benchmark/sd2tier.epx) (http://global.sap.com/solutions/benchmark/sd2tier.epx).

SAP provides a certified list of GCP configurations for SAP HANA on its website as well. For more details, see the **Find Certified IaaS Platforms** page in the [SAP HANA Hardware Directory](https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/iaas.html#categories=Google%20Cloud%20Platform) (<https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/iaas.html#categories=Google%20Cloud%20Platform>)

Google Cloud Platform offers the following storage types:

- Standard (HDD) persistent disks: Low-cost block storage for large devices.
- SSD persistent disks: Fast and reliable block storage, with high IOPS and low latency.
- Local SSDs: High-performance local block storage.
- Cloud Storage buckets: Affordable object storage.

For more information, see [Storage Options \(/compute/docs/disks/\)](/compute/docs/disks/).

GCP [persistent disks \(/persistent-disk/\)](/persistent-disk/) are designed for high durability. They store data redundantly to ensure data integrity. Each persistent disk can store up to 64 TB, so you can create large logical volumes without managing arrays of disks. One key feature is that persistent disks are automatically encrypted to protect data.

Upon creation, a Compute Engine instance allocates a single root persistent disk by default that contains the operating system. You can add more storage options to the instance as needed. For SAP implementations, we recommend using persistent disks, because they are designed for high durability and compute instances can access them like physical disks on a local machine.

The following tables describe the Linux directory structures for a typical SAP deployment.

SAP application directory structure	Storage type
<code>/sapmnt</code> ^{Note} (<code>#nfs-note-sapmnt</code>)	Standard persistent disk (HDD)
<code>/usr/sap</code>	Standard persistent disk (HDD)

Note: In distributed deployments, `/sapmnt` can also be mounted as a network file system using an NFS solution, such as [Filestore](#) (`/filestore/docs/`).

SAP HANA directory structure	Storage type
<code>/usr/sap</code>	SSD persistent disk
<code>/hana/data</code>	SSD persistent disk
<code>/hana/log</code>	SSD persistent disk
<code>/hana/shared</code> ^{Note} (<code>#nfs-note-hana</code>)	SSD persistent disk
<code>/hanabackup</code> ^{Note} (<code>#nfs-note-hana</code>)	Standard persistent disk (HDD)

Note: In distributed deployments, `/hana/shared` and `/hanabackup` can also be mounted as a network file system using an NFS solution, such as [Filestore](#) (`/filestore/docs/`).

Create all the directories for SAP HANA and SAP Business Suite instances by following the processes described in GCP deployment guides. We recommend using a [Deployment Manager](#) (`/deployment-manager/docs/fundamentals`) template to deploy a supported and certified instance. For the SAP NetWeaver deployment guides, see [All SAP NetWeaver Guides](#) (`/solutions/sap/docs/netweaver-guides`). For SAP HANA deployment, see the [SAP HANA Deployment Guide](#) (`/solutions/sap/docs/sap-hana-deployment-guide`).

SAP Business Suite on SAP HANA consists of the following technical components:

Application layer:

- **ASCS—ABAP SAP Central Services.** Contains the following components:
 - **Message server (MS):** Acts as a communication channel between application servers. Also handles load distribution.
 - **Enqueue server (ES):** Controls the lock mechanism.

- PAS—Primary application server.
 - The first or only application server for the SAP system.
- AAS—Additional application server.
 - Usually deployed for application level load balancing. You can install multiple AAS to achieve higher availability from an application layer perspective as well. If one of the application servers goes down, all user sessions connected to that application server are terminated, but users can log in again to the other associated AAS in the environment.
- WD—Web Dispatcher (optional).
 - Intelligent software load balancer that distributes HTTP and HTTPS requests, based on application type, to PAS and AAS.

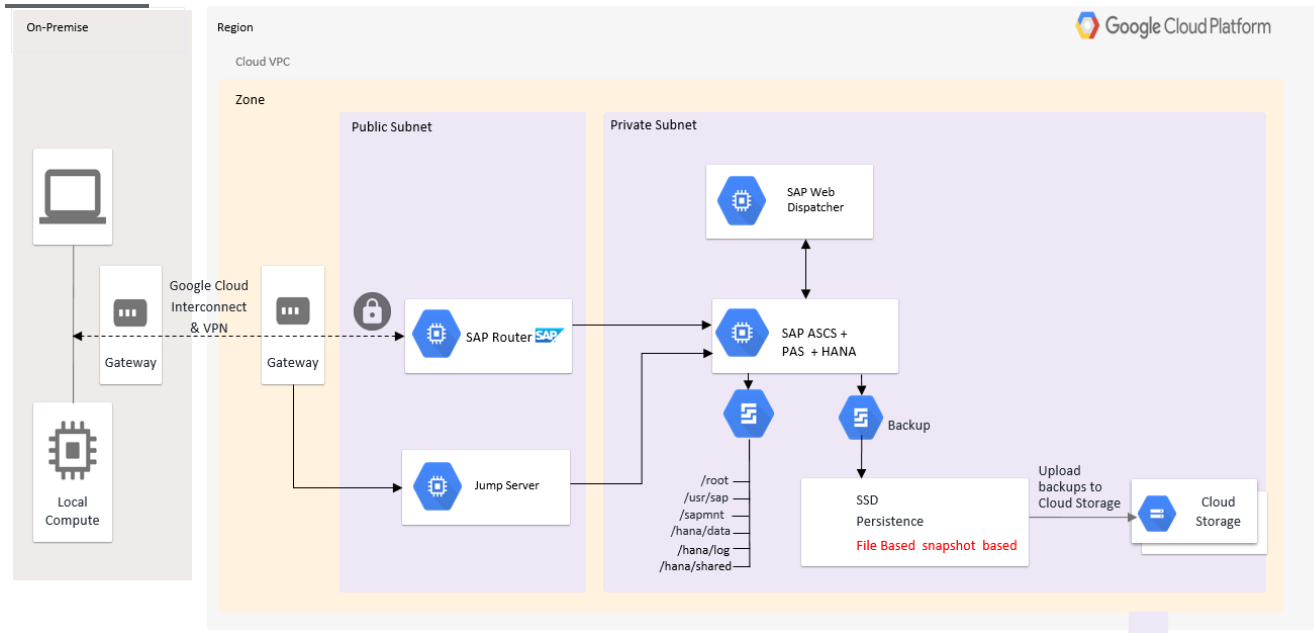
Database layer:

- SAP HANA

You can deploy SAP Business Suite on SAP HANA on GCP in either of two models: centralized deployment or distributed deployment.

In a centralized deployment, you can install SAP Business Suite and the database on the same Compute Engine instance. We recommend this approach for non-production environments such as sandbox and development environments.

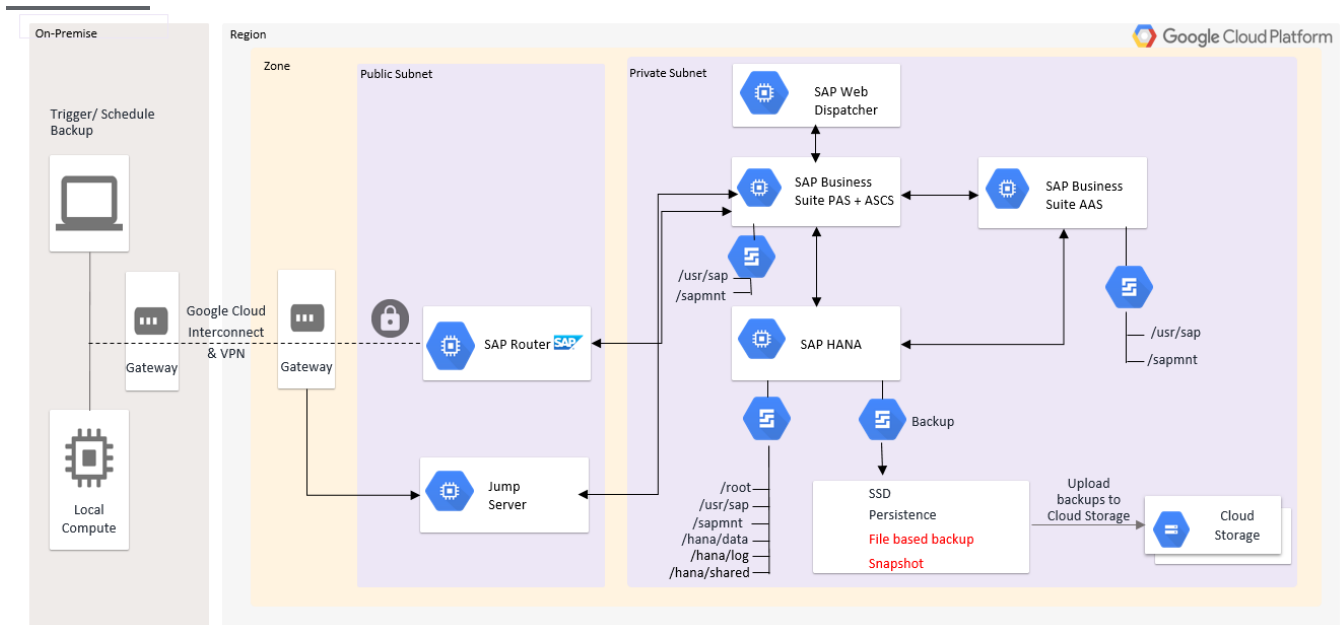
The following diagram shows a reference architecture for SAP Business Suite on SAP HANA in a centralized deployment model. Note that SAP ASCS, PAS, and HANA are all installed on the same instance.



In a distributed deployment, you can install the SAP Business Suite applications and the SAP HANA database on different Compute Engine instances. We recommend this approach for production environments or environments that require a lot of compute power to handle heavy transaction load. Each of the SAP application layer components described earlier (in "Deployment") can be independently installed on different instances.

Also, optionally, you can install one or more additional application servers (AAS), depending on your business requirements.

The following diagram shows a reference architecture for SAP Business Suite on SAP HANA in a distributed deployment model.



SAP Business Suite and the SAP HANA database are installed on different Compute Engine instances. The database must be installed using the deployment method certified by GCP. For details on installing an SAP HANA database for SAP HANA Scale-up or SAP HANA Scale-out, see the [SAP HANA Deployment Guide](/solutions/sap/docs/sap-hana-deployment-guide) (/solutions/sap/docs/sap-hana-deployment-guide). Currently, SAP Business Suite on SAP HANA is supported only in a Scale-up model.

High availability (HA) and disaster recovery (DR) are sets of techniques, engineering practices, and design principles to enable business continuity in the event of failures. These approaches work by eliminating single points of failure and providing the ability to rapidly resume operations after a system or component outage with minimal business disruption. Fault recovery is the process of recovering and resuming operations after an outage due to a failed component.

For example, here are some HA and DR tools:

- Linux clustering across zones.
- [Live Migration](/compute/docs/instances/live-migration) (/compute/docs/instances/live-migration).
- [Automatic restart](/compute/docs/instances/setting-instance-scheduling-options#autorestart) (/compute/docs/instances/setting-instance-scheduling-options#autorestart).
- [Service auto-restart](#)
(<https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.01/en-US/cf10efba8bea4e81b1dc1907ecc652d3.html>)

- Backups (#backups).

To ensure high availability for SAP Business Suite on SAP HANA, consider the following components:

- SAP HANA database (HDB)
- ABAP Central Services (ASCS)
- Primary application server (PAS)

SAP HANA database: We recommend the SAP HANA system replication (HSR) solution for achieving database high availability. In this scenario, system replication is configured between the primary and secondary node and data is replicated from the primary to the secondary persistent disk. For more details, see [Configuring SAP HANA System Replication](https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.02/en-US/676844172c2442f0bf6c8b080db05ae7.html) (<https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.02/en-US/676844172c2442f0bf6c8b080db05ae7.html>)

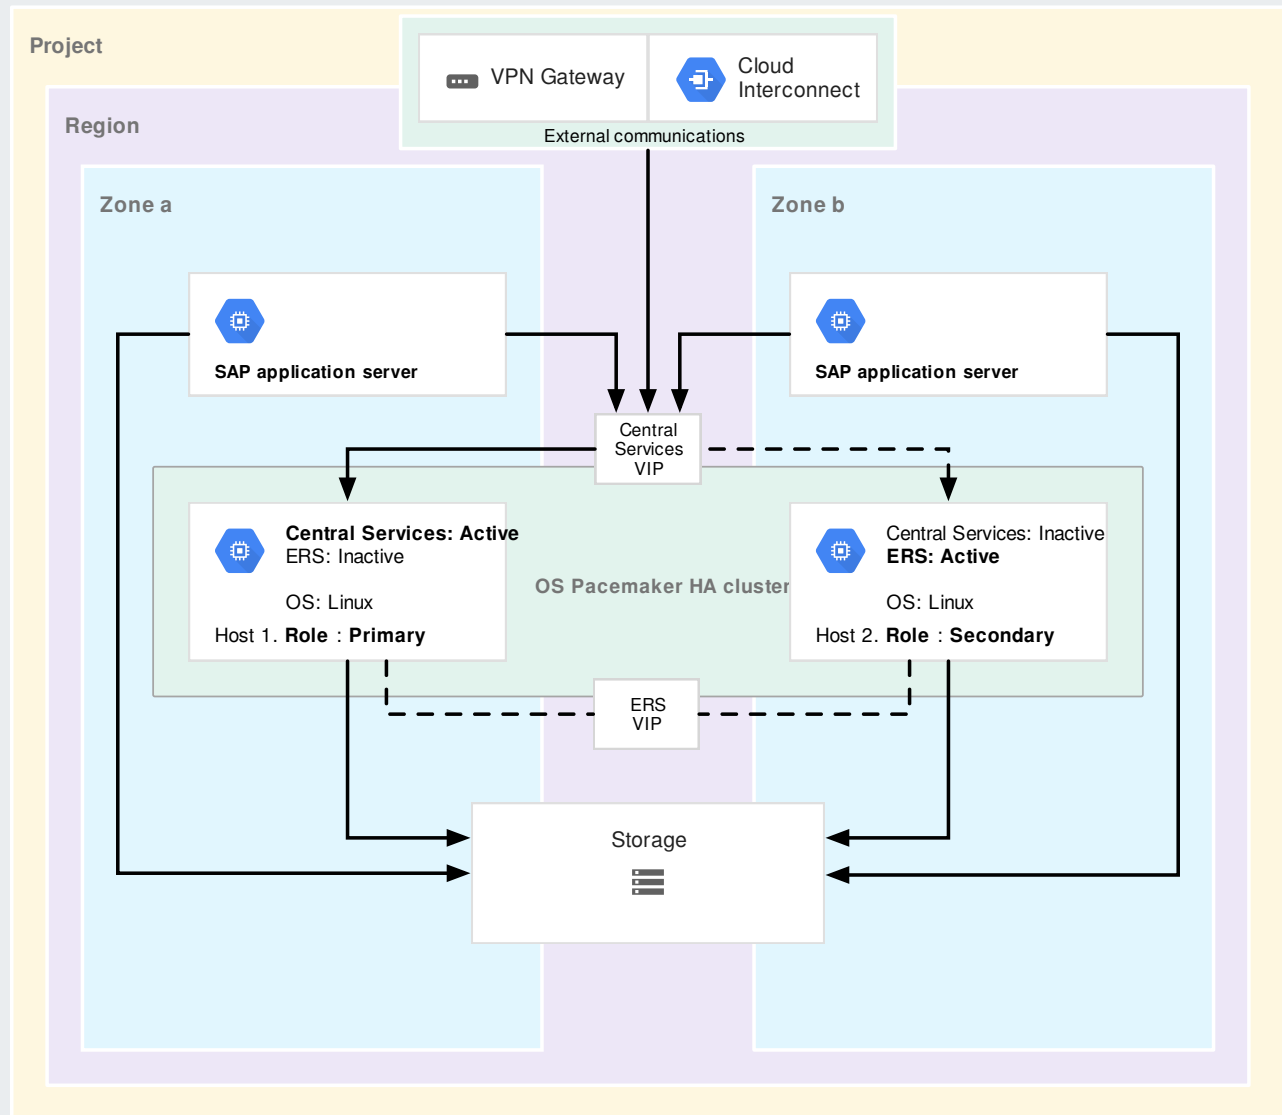
In the SAP HANA system replication scenario, failover is not automated by default. You can implement failover using operating system level clustering, which is designed to manage component failures. Clustering involves the use of multiple servers, storage devices, and interconnections to form a single highly available system. For more details about configuring HA for SAP HANA, see [SAP HANA High Availability and Disaster Recovery Planning Guide](/solutions/sap/docs/sap-hana-hadr-planning-guide) (/solutions/sap/docs/sap-hana-hadr-planning-guide).

ABAP Central Services: ASCS consists of a message server (MS) and an enqueue server (ES). The message server acts as a communication channel between the application servers and handles the load distribution, while the enqueue server controls the lock mechanism. We recommend that you use a clustering solution to achieve high availability on ASCS. To implement HA, install ABAP Central Services (ASCS) and enqueue replication service (ERS) on primary and secondary nodes. When the primary node goes down, message and enqueue (MS/ES) services automatically fail over to the secondary node. When the primary node becomes available again, you can automatically or manually fail over to the original primary node. For more details, see the SUSE setup guide for a [SAP ASCS high availability cluster](#)

(https://www.suse.com/docrep/documents/zrfolmyklc/sap_netweaver_availability_cluster_740_setup_guide.pdf)

The following diagram shows an architecture for implementing HA for ASCS.

Google Cloud Platform



Primary application server: You can achieve high availability for the primary application server by installing additional application servers (AAS). You can install multiple AAS to achieve higher availability. If one of the application servers goes down, all user sessions connected to that application server are terminated, but users can log in to the other application servers. Google Cloud Platform provides a live migration feature that could be used for PAS HA. For more information, see [Live Migration](https://cloud.google.com/compute/docs/instances/live-migration/) (/compute/docs/instances/live-migration).

To recover an SAP Business Suite on SAP HANA system from a disaster, use the following methods:

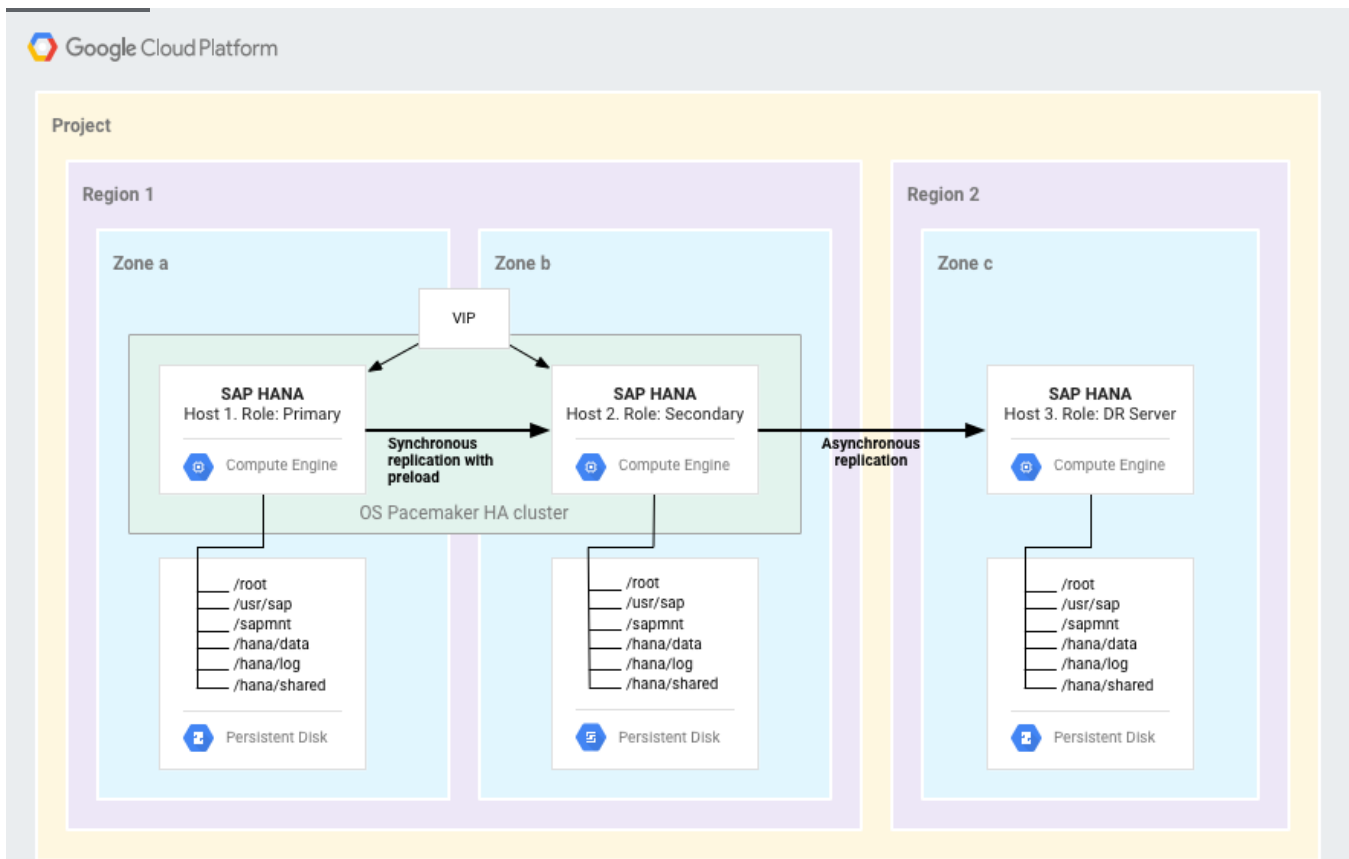
- SAP HANA system replication
- SAP HANA backup

For DR scenarios, we recommend that you have the standby system in a different region than the primary system and use asynchronous replication.

Choose an SAP HANA system replication option that best fits your business requirements for recovery point objectives (RPO) and that fits your cost/benefit preferences.

The following diagram shows the flow of replication. Specifically:

1. Virtual IP (VIP) is configured to point only to the active (read/write) node.
2. In this scenario, Node 1 plays the primary role. When a failover is triggered, Node 2 takes over and acts as the primary node, with the virtual IP moving to Node 2.



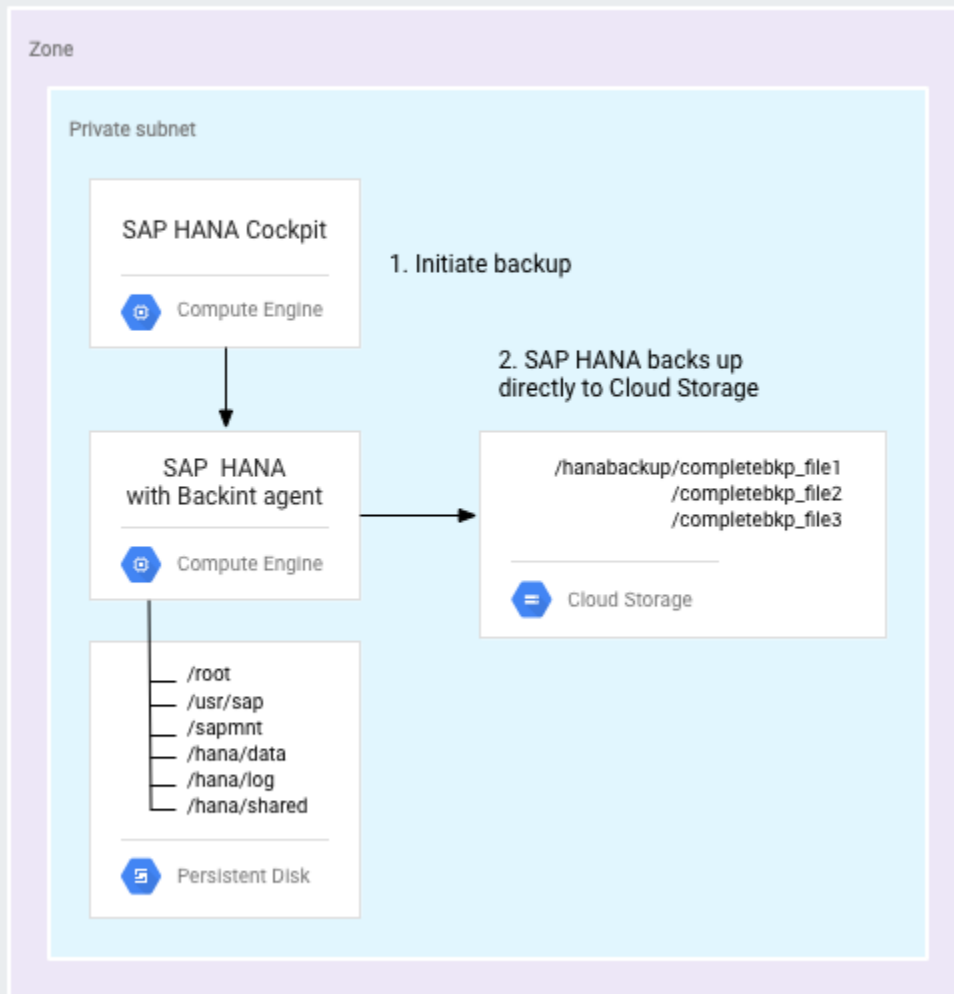

Make backups of your application server and database on a regular schedule so that you can recover in case of a system crash, data corruption, or other problems.

You have a number of options for backing up your SAP HANA data on Google Cloud, including:

- Backing up directly to Cloud Storage using the SAP-certified Cloud Storage Backint agent for SAP HANA (Backint agent).
- Backing up to a persistent disk and then uploading the backups to Cloud Storage.
- Taking snapshots of the entire disk that contains the /hanabackup directory by using the Compute Engine snapshot function.

You can install the Cloud Storage Backint agent for SAP HANA (Backint agent, which simplifies backup storage. The Backint agent integrates with the native SAP HANA backup and recovery functions so you can backup directly to, and recover from, Cloud Storage without needing persistent disk storage for your backups. For more information, see the [SAP HANA Operations Guide \(/solutions/sap/docs/sap-hana-operations-guide#backint_operations\)](/solutions/sap/docs/sap-hana-operations-guide#backint_operations).

For information about the SAP certification of the Cloud Storage Backint agent for SAP HANA, see [SAP Note 2031547 \(https://launchpad.support.sap.com/#/notes/2031547\)](https://launchpad.support.sap.com/#/notes/2031547). The following diagram shows the flow of backups when you use the Backint agent.

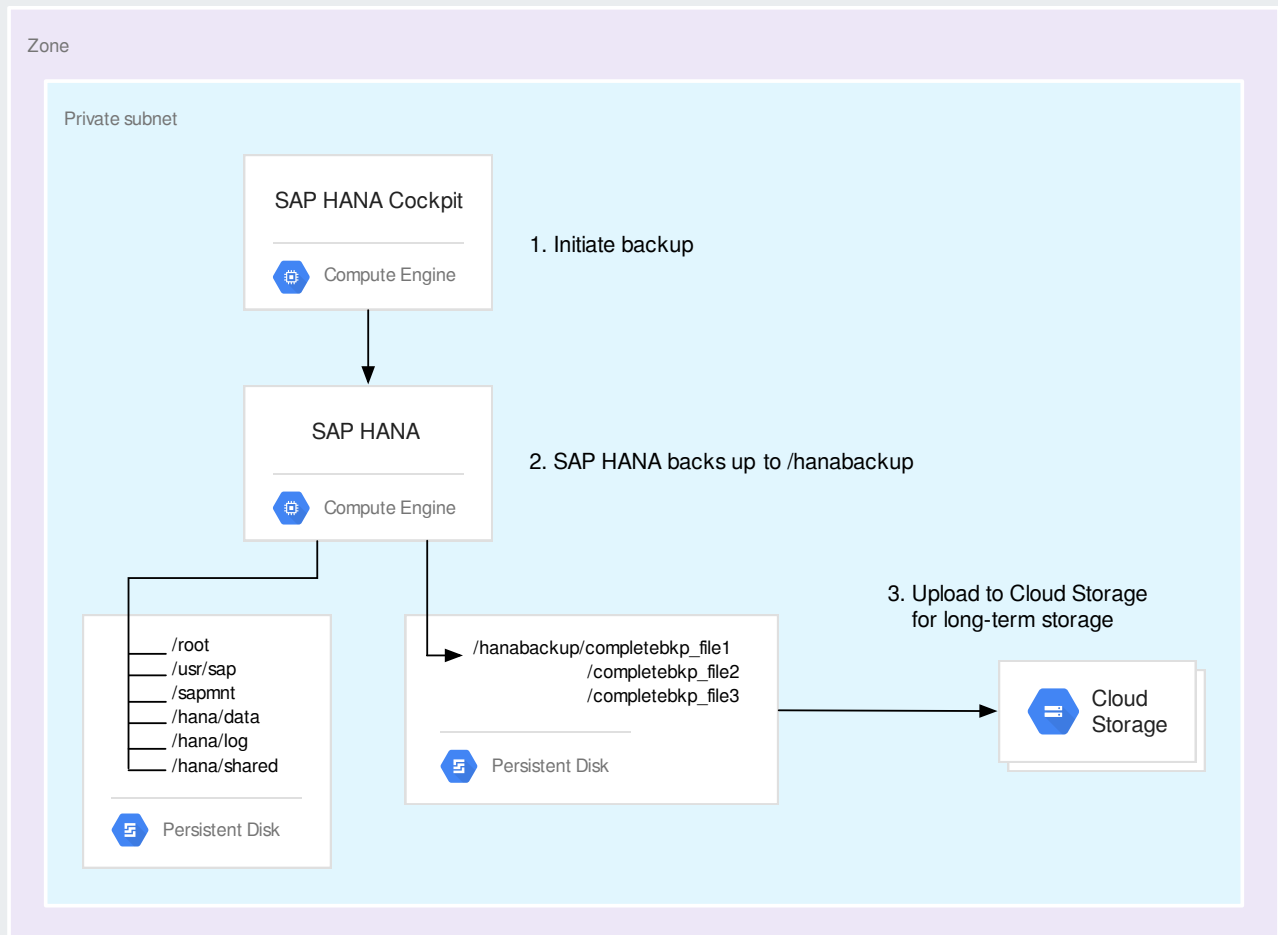


You can use the native SAP HANA backup and recovery function to store backups on Compute Engine persistent disks. You can use a Cloud Storage bucket for longer-term storage of the backups.

During normal operation, SAP HANA automatically saves data from memory to disk at regular save points. Additionally, all data changes are captured in redo log entries. A redo log entry is written to disk after each committed database transaction. You can back up the redo logs to longer term storage at regular intervals.

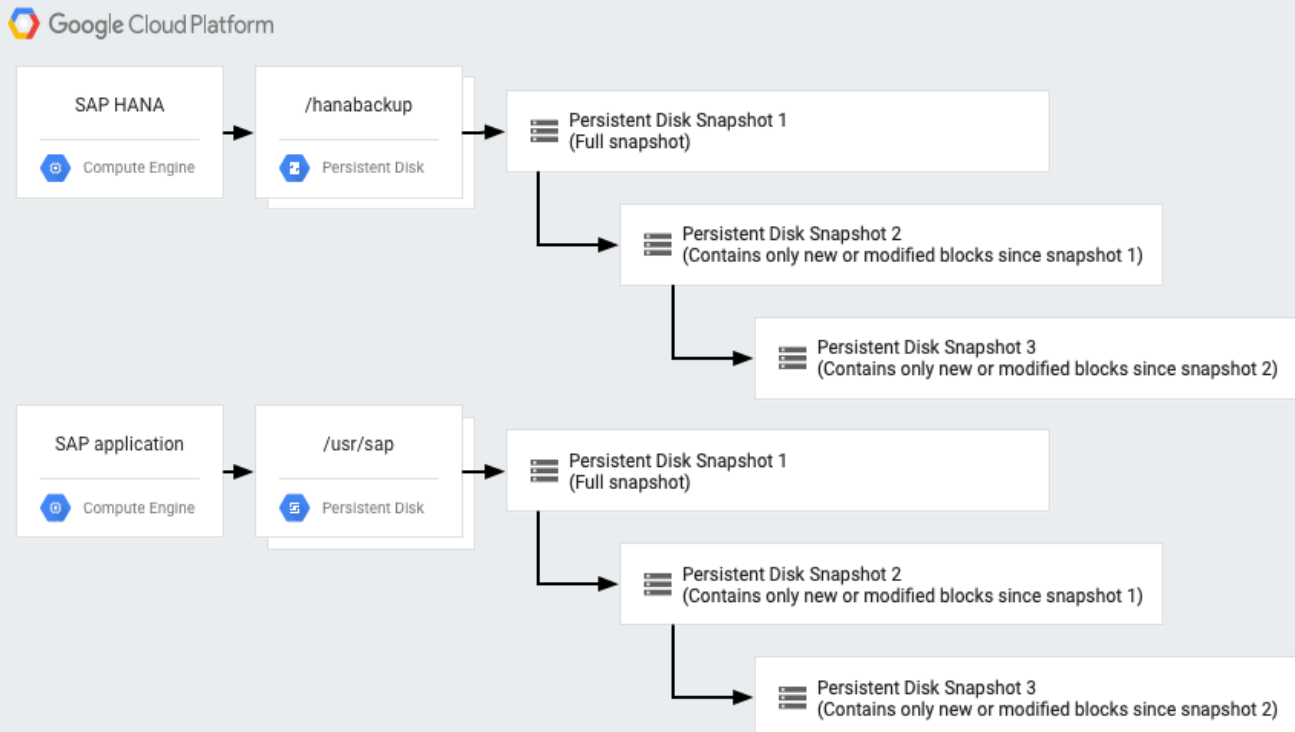
From SAP HANA 2.0 onwards, you must use SAP HANA cockpit to back up SAP HANA.

The following diagram shows the flow of the backup feature for SAP HANA.



Another option you can add to your backup strategy is to take snapshots of entire disks by using the [persistent-disk snapshot \(/compute/docs/disks/create-snapshots\)](/compute/docs/disks/create-snapshots) feature of Compute Engine. For example, you can take scheduled snapshots of your backup directory disk for use in disaster recovery scenarios. To ensure application consistency, take snapshots when no changes are being made to the target volume. Snapshots occur on the block level.

After the first snapshot, each subsequent snapshot is incremental and stores only incremental block changes, as shown in the following diagram.



The recovery tools in SAP HANA can recover to the latest point in time or a specific point in time, and you can use these tools to restore to a new system or create a copy of the database. Unlike backups, which you can run while the database is operational, you can use recovery tools only while the database is shut down. Recovery options are listed below; choose an appropriate one for your situation.

- Restore to the most recent state, using any of the following resources:
 - Full backup or snapshot.
 - Redo log entries that are still available.

- Restore to a point in time in the past.
- Restore to a specified full backup.

Before beginning to deploy your SAP systems on Google Cloud, read the SAP Notes in the following list that are relevant to your planned configuration. Before proceeding with any SAP product implementation, always check the SAP Marketplace for updated product installation guides and notes.

- [2456432 - SAP Applications on Google Cloud: Supported Products and Google VM types](https://launchpad.support.sap.com/#/notes/2456432)
(<https://launchpad.support.sap.com/#/notes/2456432>).
- [2446441 - Linux on Google Cloud \(IaaS\): Adaption of your SAP License](https://launchpad.support.sap.com/#/notes/2446441)
(<https://launchpad.support.sap.com/#/notes/2446441>).
- [2456953 - Windows on Google Cloud \(IaaS\): Adaption of your SAP License](https://launchpad.support.sap.com/#/notes/2456953)
(<https://launchpad.support.sap.com/#/notes/2456953>).
- [1380654 - SAP support in public cloud environments](https://launchpad.support.sap.com/#/notes/1380654)
(<https://launchpad.support.sap.com/#/notes/1380654>).
- [2456406 - SAP on Google Cloud Platform: Support Prerequisites](https://launchpad.support.sap.com/#/notes/2456406)
(<https://launchpad.support.sap.com/#/notes/2456406>).