

[Solutions](https://cloud.google.com/solutions/) (https://cloud.google.com/solutions/)

[SAP](https://cloud.google.com/solutions/sap/docs) (https://cloud.google.com/solutions/sap/docs) [Guides](#)

# Reference Architecture: SAP S/4HANA on Google Cloud Platform

## Overview

This document is for people who are evaluating Google Cloud as a platform for deploying SAP S/4HANA, 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.

Google Cloud provides a cost-effective, reliable, secure, and high-performance SAP certified infrastructure for running SAP S/4HANA on SAP HANA. For a complete list of supported SAP solutions on Google Cloud, see [SAP on Google Cloud](https://cloud.google.com/sap/) (https://cloud.google.com/sap/).

## Licensing

If you're an SAP customer, you can use your existing license to deploy SAP Business Suite on Google Cloud under a bring-your-own-license (BYOL) model. Google Cloud 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](https://cloud.google.com/compute/docs/images#custom_images) ([https://cloud.google.com/compute/docs/images#custom\\_images](https://cloud.google.com/compute/docs/images#custom_images)) and licenses as well.

## Sizing

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](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 Google Cloud. (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 Applications on Google Cloud: Supported Products and Google VM types](https://launchpad.support.sap.com/#/notes/2456432) (<https://launchpad.support.sap.com/#/notes/2456432>)). SAP and Google Cloud 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 Google Cloud infrastructure.

Before you migrate existing systems from SAP ECC to S/4HANA, SAP recommends that you run the `/SDF/HDB_SIZING` report, as described in SAP note 1872170, [Business Suite on HANA and S/4HANA sizing report](https://launchpad.support.sap.com/#/notes/1872170) (<https://launchpad.support.sap.com/#/notes/1872170>). This sizing report analyzes your source system's current memory and processing needs and provides information on the requirements for moving to S/4HANA.

## Supported machine types

Google Cloud offers Compute Engine instance types that are certified by SAP to match sizing requirements when you're deploying S/4HANA. For more information about sizing on Google Cloud and the machine types that are supported, see the following pages:

- [Supported machine types for SAP HANA on Google Cloud](https://cloud.google.com/solutions/sap/docs/sap-hana-planning-guide#vm_types) ([https://cloud.google.com/solutions/sap/docs/sap-hana-planning-guide#vm\\_types](https://cloud.google.com/solutions/sap/docs/sap-hana-planning-guide#vm_types)).
- [Supported machine types for SAP NetWeaver on Google Cloud](https://cloud.google.com/solutions/sap/docs/netweaver-planning-guide#machine_types) ([https://cloud.google.com/solutions/sap/docs/netweaver-planning-guide#machine\\_types](https://cloud.google.com/solutions/sap/docs/netweaver-planning-guide#machine_types)).

Custom machine types for SAP HANA on Google Cloud 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>) 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>).

SAP provides a certified list of Google Cloud 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>)

## Disks and file systems for S/4HANA

Google Cloud 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 (<https://cloud.google.com/compute/docs/disks/>).

Google Cloud persistent disks (<https://cloud.google.com/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 SAP HANA and ABAP on Google Cloud.

## Recommended Linux directory structure for SAP HANA

SAP HANA directory structure	Storage type
/usr/sap	SSD persistent disk
/hana/data	SSD persistent disk
/hana/log	SSD persistent disk
/hana/shared	SSD persistent disk
/hanabackup	Standard persistent disk (HDD)

## Recommended Linux directory structure for ABAP

ABAP directory structure	Storage type
/sapmnt	Standard persistent disk (HDD)
/usr/sap/	Standard persistent disk (HDD)

## Deployment

SAP S/4HANA consists of the following technical components:

- SAP HANA.
- 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.
- SAP NetWeaver Gateway.

- Fiori front end.
- WD—Web Dispatcher (optional).
  - Intelligent software load balancer that distributes HTTP and HTTPS requests, based on application type, to PAS and AAS.

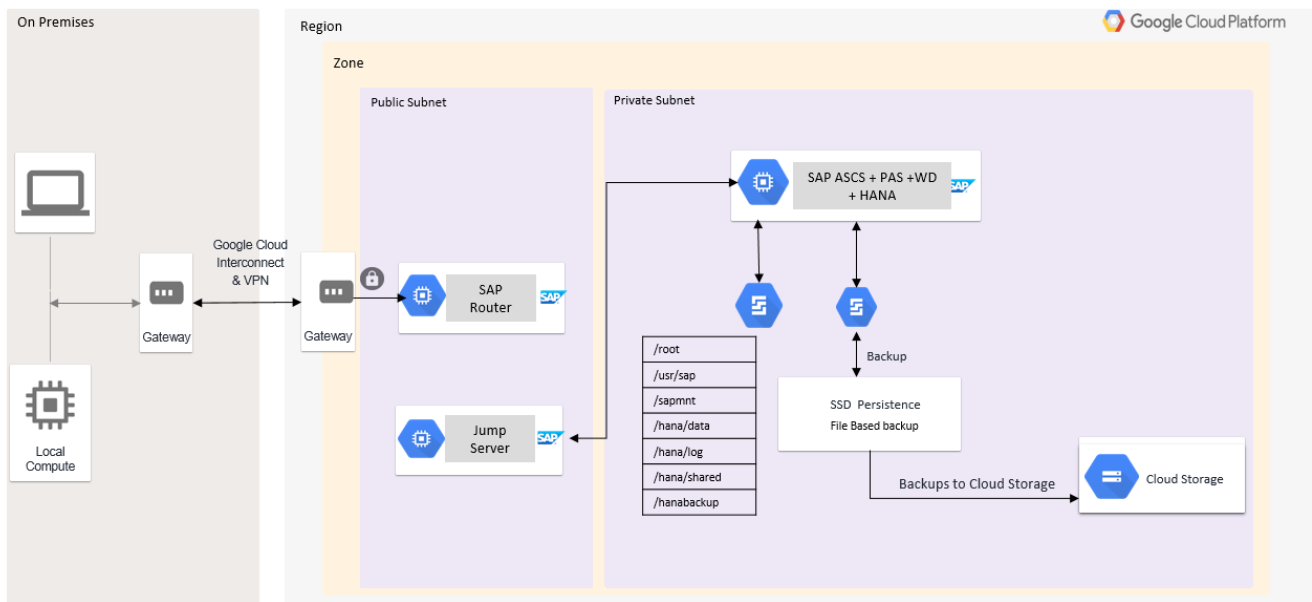
## Deployment models

You can deploy S/4HANA on Google Cloud in either of two models: centralized deployment or distributed deployment.

### Centralized deployment

In a centralized deployment, you can install S/4HANA and the SAP HANA database on the same Compute Engine instance. We recommend this approach for non-production environments such as sandbox and development environments.

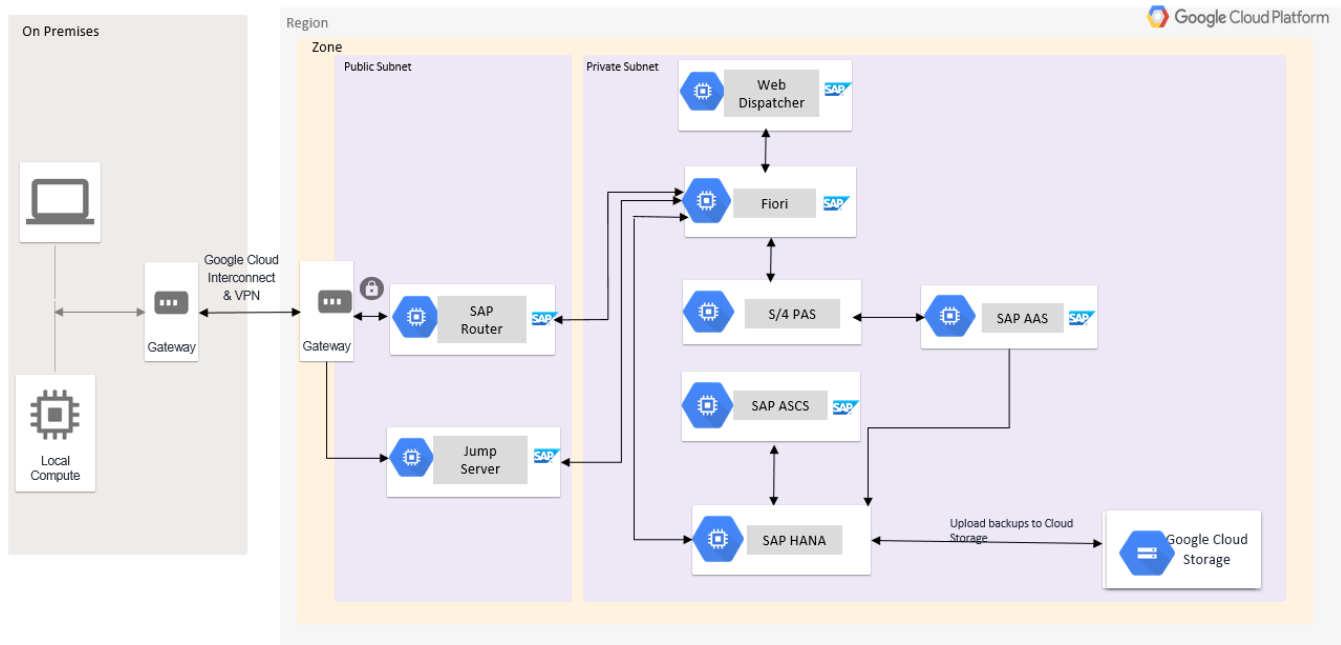
The following diagram shows a reference architecture for S/4HANA in a centralized deployment model. Note that SAP ASCS, PAS, WD, and HANA are all installed on the same instance.



### Distributed deployment

In a distributed deployment, you can install the different components 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.

The following diagram shows a reference architecture for S/4HANA in a distributed deployment model. Note that SAP ASCS, PAS, WD, and HANA are all installed on different instances.



To install SAP HANA in a centralized or distributed deployment, use the deployment script. For more information, see [SAP HANA Deployment Guide](https://cloud.google.com/solutions/sap/docs/sap-hana-deployment-guide) (<https://cloud.google.com/solutions/sap/docs/sap-hana-deployment-guide>).

### A note on load balancing

In a distributed S/4HANA environment, load balancing is mandatory. You can configure application load balancing using the SAP application layer.

### High availability and disaster recovery

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 (<https://cloud.google.com/compute/docs/instances/live-migration>).
- Automatic restart  
(<https://cloud.google.com/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.

Some further details about some of those items:

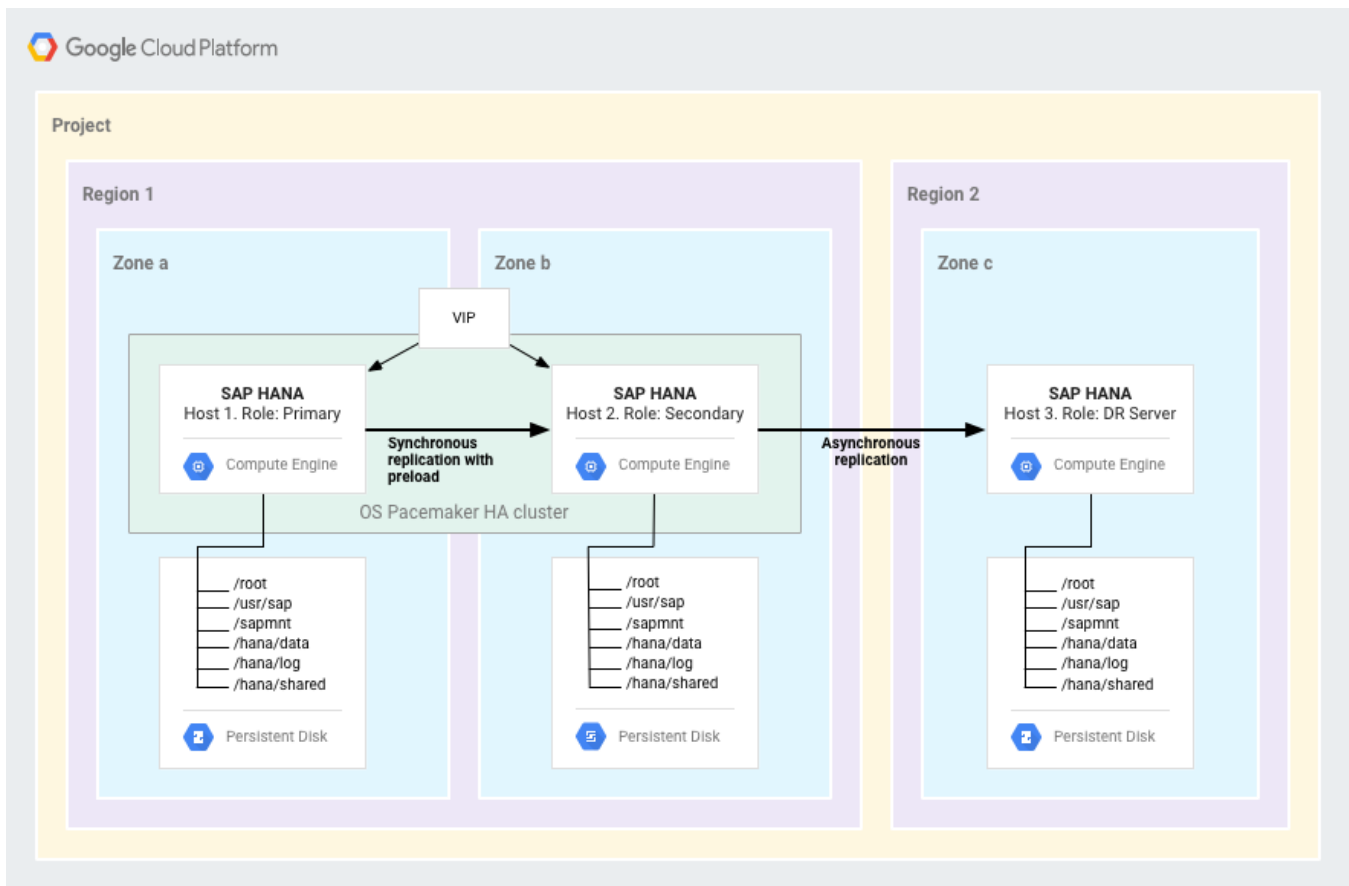
**Linux clustering across zones:** Setting up your Linux cluster across zones

(<https://cloud.google.com/compute/docs/regions-zones/>) helps guard against component failures in a given region. You can deploy a Linux cluster across zones using either an active/passive configuration or an active/active

(<https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.03/en-US/fe5fc53706a34048bf4a3a93a5d7c866.html>)

configuration. In both cases, you start by setting up two Compute Engine instances in separate zones, each with its own SAP HANA database.

- **Active/passive configuration:** Configure one instance as the cluster's primary node (active) and the other as the secondary node (passive). Use SAP HANA System Replication (SR) to configure the secondary node to take over as primary if the primary fails, as shown in the following diagram. For more information on how to configure and set up HANA SR, see HANA System Replication  
(<https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.02/en-US/099caa1959ce4b3fa1144562fa09e163.html>)



- **Active/active configuration** (read-enabled): Configure both instances to be active, but the secondary node as read-only. This setup is based on continuous log replay. The virtual IP (VIP) is configured to point to the current read/write node.
- For more information, see the [SAP HANA High Availability and Disaster Recovery Planning Guide](https://cloud.google.com/solutions/sap/docs/sap-hana-hadr-planning-guide) (<https://cloud.google.com/solutions/sap/docs/sap-hana-hadr-planning-guide>).

**Live migration:** Compute Engine offers live migration to keep your Compute Engine instances running even when a host system event occurs, such as a software or hardware update. In that situation, Compute Engine live-migrates your running instance to another host in the same zone, rather than requiring your running instance to be rebooted. The mechanism replicates the original instance's VM state, so when the new instance comes up, it already has the original instance's memory pre-loaded.

In the rare case when live migration doesn't happen, the failed virtual machine is automatically restarted on the new hardware within the same zone.

For more details, see [Live Migration](https://cloud.google.com/compute/docs/instances/live-migration) (<https://cloud.google.com/compute/docs/instances/live-migration>).



## Backup and recovery

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.

### Backups

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

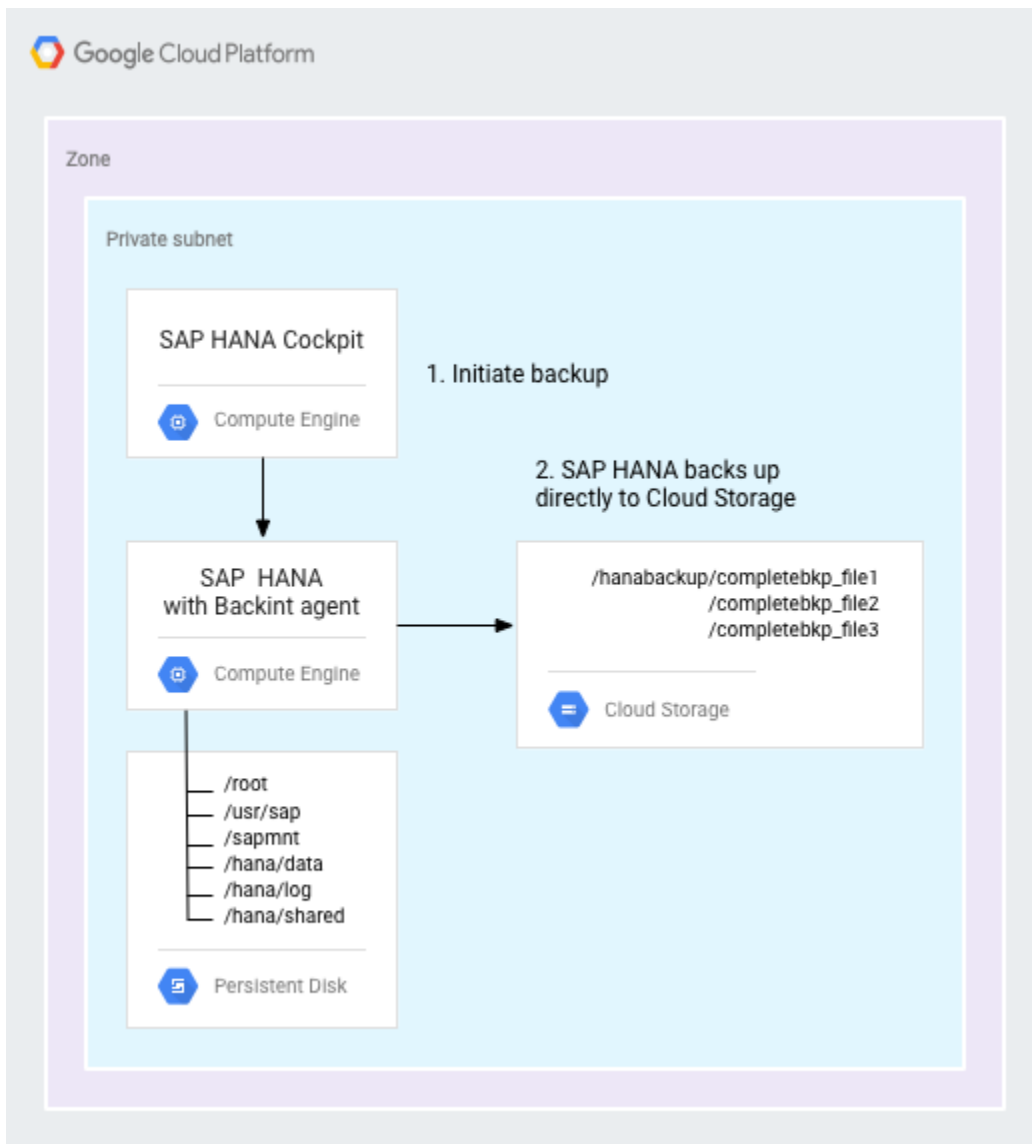
- Backing up directly to Cloud Storage using the SAP-certified Cloud Storage Backint agent for SAP HANA (Backint agent).
- Backing up to a persistent disk and then uploading the backups to Cloud Storage.
- Taking snapshots of the entire disk that contains the /hanabackup directory by using the Compute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

### The Cloud Storage Backint agent for SAP HANA

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](https://cloud.google.com/solutions/sap/docs/sap-hana-operations-guide#backint_operations) ([https://cloud.google.com/solutions/sap/docs/sap-hana-operations-guide#backint\\_operations](https://cloud.google.com/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.



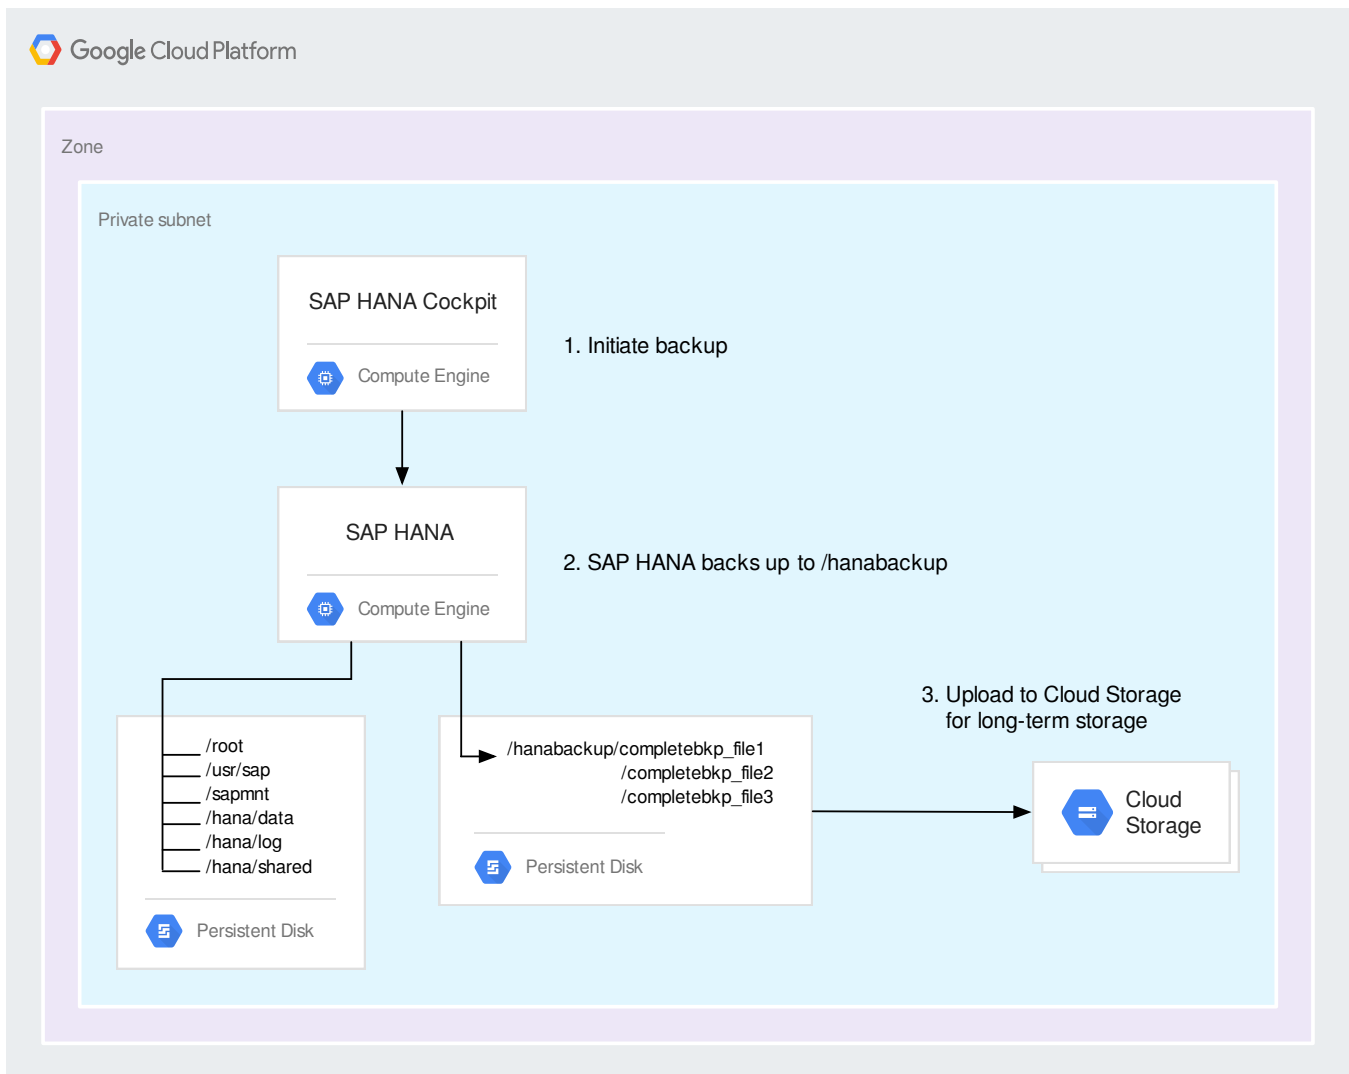
## Backing up to persistent disks

You can use the native SAP HANA backup and recovery function with Compute Engine persistent disks and 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.

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

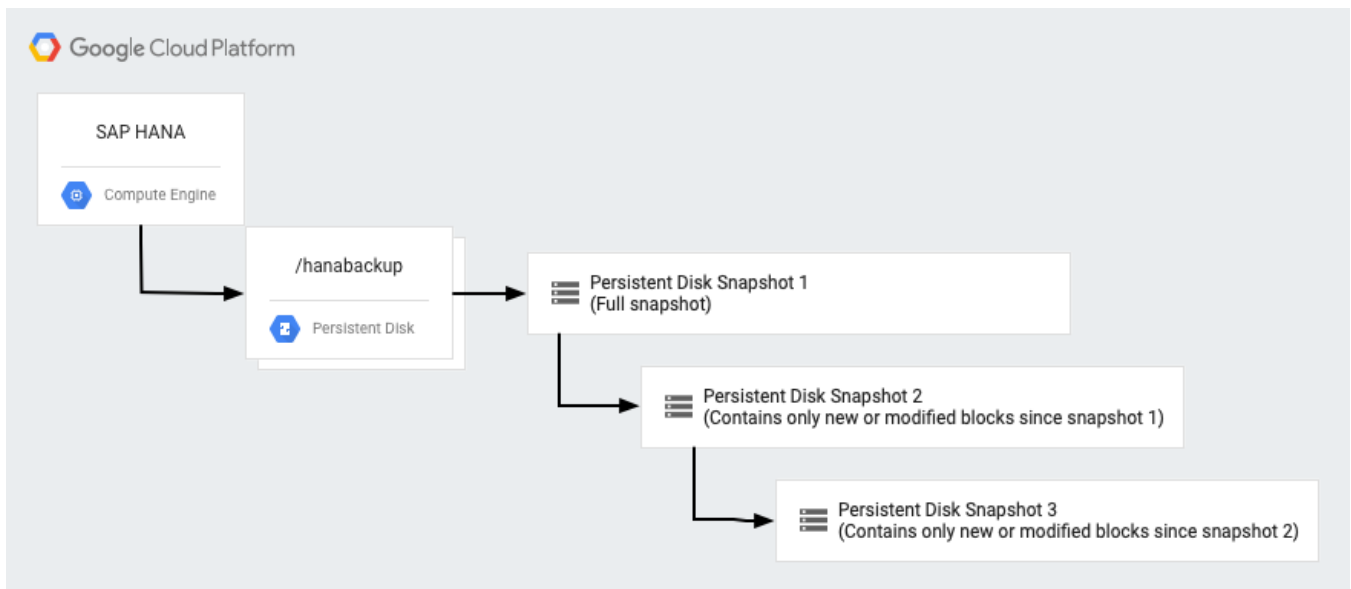
The following diagram shows the flow of the backup feature for SAP HANA.



## Backing up persistent disks by using snapshots

Another option you can add to your backup strategy is to take snapshots of entire disks by using the [persistent-disk snapshot](https://cloud.google.com/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.



## Recovery

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.
  - Log backups.
  - Redo log entries that are still available.
- Restore to a point in time in the past.
- Restore to a specified full backup.

## Important pre-deployment SAP Notes

Read the following SAP Notes before beginning to deploy SAP S/4HANA on Google Cloud. 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>).

---

*Except as otherwise noted, the content of this page is licensed under the [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/) (<https://creativecommons.org/licenses/by/4.0/>), and code samples are licensed under the [Apache 2.0 License](https://www.apache.org/licenses/LICENSE-2.0) (<https://www.apache.org/licenses/LICENSE-2.0>). For details, see our [Site Policies](https://developers.google.com/terms/site-policies) (<https://developers.google.com/terms/site-policies>). Java is a registered trademark of Oracle and/or its affiliates.*

*Last updated November 27, 2019.*