

This guide provides an overview for planning for and managing high availability and disaster recovery for SAP HANA systems deployed on Google Cloud by following the [SAP HANA on Google Cloud deployment guide](/solutions/sap/docs/sap-hana-deployment-guide) (/solutions/sap/docs/sap-hana-deployment-guide). This guide is not intended to replace the standard SAP documentation.

You can obtain high availability for SAP HANA on Google Cloud by using a combination of Google Cloud and SAP features that handle failures at the infrastructure or software levels. The following tables describe SAP and Google Cloud features that are used to provide high availability.

Feature	Description
Compute Engine live migration (/compute/docs/instances/live-migration)	<p>Compute Engine monitors the state of the underlying infrastructure and automatically migrates your instance away from an infrastructure maintenance event. No user intervention is required.</p> <p>Compute Engine keeps your instance running during the migration if possible. In the case of major outages, there might be a slight delay between when the instance goes down and when it is available.</p> <p>In multi-host systems, shared volumes, such as the <code>/hana/shared`</code> volume used in the deployment guide (/solutions/sap/docs/sap-hana-deployment-guide), are persistent disks attached to VM that hosts the master host, and are NFS-mounted to the worker hosts. The NFS volume is inaccessible for up to a few seconds.</p>

Feature	Description
	<p>in the event of the master host's migration. When the master host has restarted, the NFS volume functions again on all hosts, and normal operation resumes automatically.</p> <p>A recovered instance is identical to the original instance, including the instance ID, private IP address, and all instance metadata and storage. By default, standard instances are set to live migrate. We recommend not changing this setting.</p> <p>For more information, see Live migrate (/compute/docs/instances/setting-instance-scheduling-options#live_migrate)</p>
<p><u>Compute Engine automatic restart</u> (/compute/docs/instances/setting-instance-scheduling-options#autorestart)</p>	<p>If your instance is set to terminate when there is a maintenance event or if your instance crashes because of an underlying hardware issue, you can set up Compute Engine to automatically restart the instance. By default, instances are set to automatically restart. We recommend not changing this setting.</p> <p>For more information about automatic restart, see Automatic restart of VM instances on Google Cloud (#auto_restart)</p>

Feature	Description
<p><u>SAP HANA Service Auto-Restart</u> (https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.01/en-US/cf10efba8bea4e81b1dc1907ecc652d3.html)</p>	<p>SAP HANA Service Auto-Restart is a fault recovery solution provided by SAP.</p> <p>SAP HANA has many configured services running all the time for various activities. When any of these services is disabled due to a software failure or human error, the SAP HANA service auto-restart watchdog function restarts it automatically. When the service is restarted, it loads all the necessary data back into memory and resumes its operation.</p>
<p><u>SAP HANA Backups</u> (https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.01/en-US/158071e0c455487a89ea56ac53ad4b31.html)</p>	<p>SAP HANA backups create copies of data from your database that can be used to reconstruct the database to a point in time.</p> <p>For more information about using SAP HANA backups on Google Cloud, see the SAP HANA operations guide (/solutions/sap/docs/sap-hana-operations-guide#backup_and_recovery)</p>

Feature	Description
<p><u>SAP HANA Storage Replication</u> (https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.01/en-US/2a3b86c65f0d485cb39ff10181986125.html)</p>	<p>SAP HANA storage replication provides storage-level disaster recovery support through certain hardware partners. SAP HANA storage replication isn't supported on Google Cloud. You can consider using Compute Engine persistent disk snapshots instead.</p> <p>For more information about using persistent disk snapshots to back up SAP HANA systems on Google Cloud, see the SAP HANA operations guide (/solutions/sap/docs/sap-hana-operations-guide#backup_and_recovery)</p> <p>.</p>

Feature	Description
<p><u>SAP HANA Host Auto-Failover</u> (https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.01/en-US/ae60cab98173431c97e8724856641207.html)</p>	<p>SAP HANA host auto-failover is a local fault recovery solution that requires one or more standby SAP HANA hosts in a scale-out system. If one of the main hosts fail, host auto-failover automatically brings the standby host online and restores the failed host as a standby host.</p> <p>On Google Cloud, automatic restoration, which is typically faster than host auto-failover, and live migration might make SAP HANA host auto-failover unnecessary. However, if your business demands it, Google Cloud provides a Deployment Manager configuration template that you can use to automate the deployment of an SAP HANA scale-out system that supports host auto-failover.</p> <p>For more information, see:</p> <ul style="list-style-type: none"> • SAP HANA host auto-failover Google Cloud (#host_auto_failover) • SAP HANA Scaleout System with host auto-failover Deployment Guide (/solutions/sap/docs/sap-hana-ha-scaleout-deployment-guide)
<p><u>SAP HANA System Replication</u> (https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.01/en-US/b74e16a9e09541749a745f41246a065e.html)</p>	<p>SAP HANA system replication allows you to configure one or more systems to take over for your primary system in high-availability or disaster recovery scenarios. You can tune replication to meet your needs in terms of performance and failover time.</p>

In the case of VM instance restart due to maintenance or other issues, Compute Engine automatic restart and SAP HANA service auto-restart work together to automatically restart the instance and application without your intervention. No client redirection is needed.

Google Cloud supports SAP HANA auto host-failover, the local fault-recovery solution provided by SAP HANA. The host auto-failover solution uses one or more standby hosts that are kept in reserve to take over work from the master or a worker host in the event of a host failure. The standby hosts do not contain any data or process any work.

The `/hana/data` and `/hana/log` volumes are mounted on the master and worker hosts only. When a takeover occurs, the host auto-failover solution uses the SAP HANA Storage Connector API and the Compute Engine `gceStorageConnector` plugin to manage the switching of these disks from the failed host to the standby host. The configuration parameters for the `gceStorageConnector` plugin, including whether fencing is enabled or disabled, are stored in the storage section of the SAP HANA `global.ini` file.

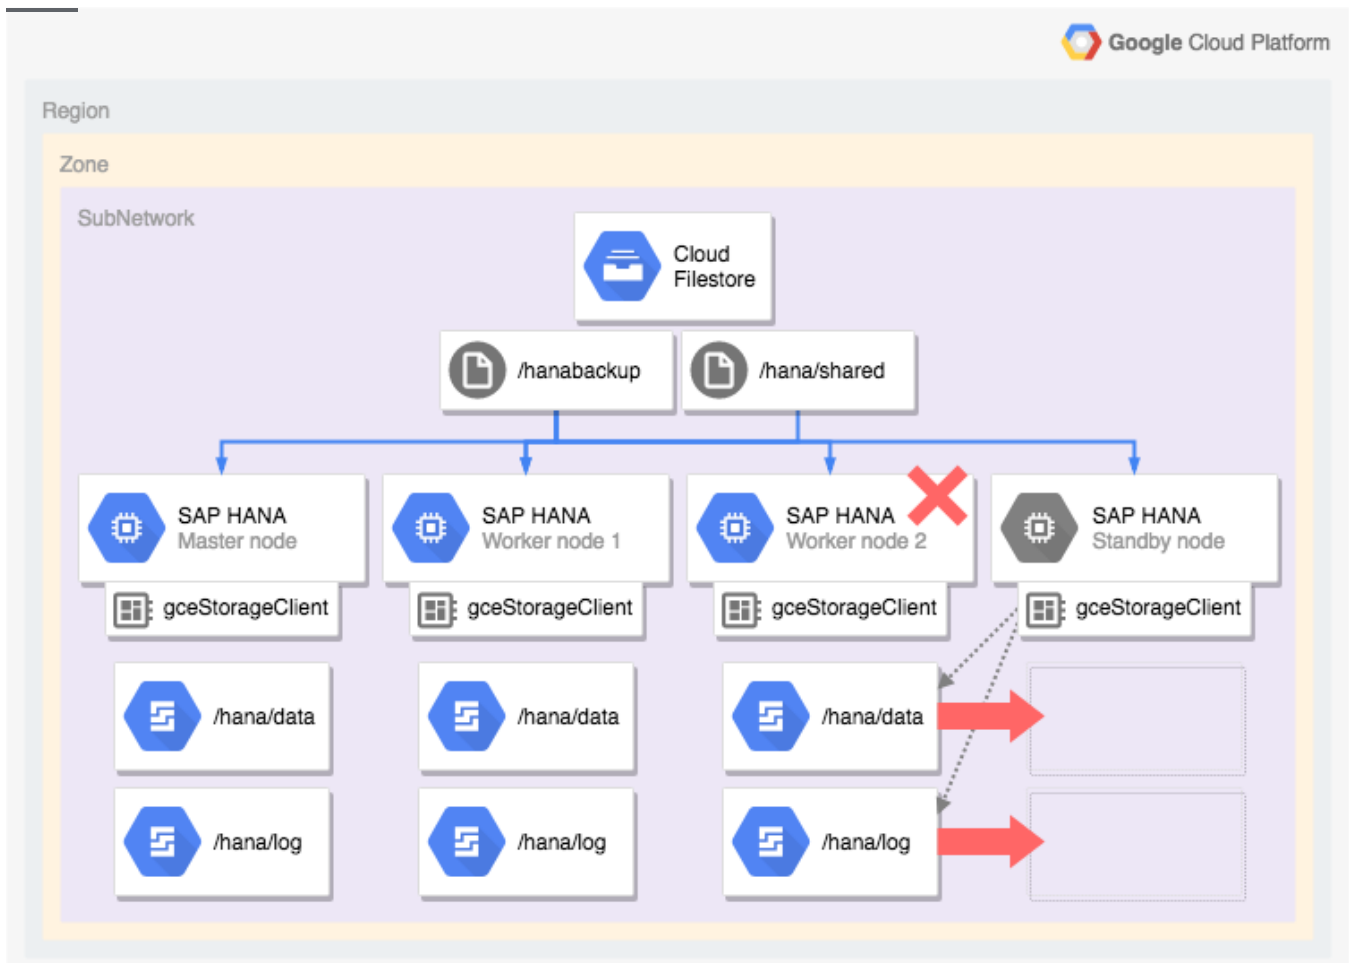
The `/hana/shared` and `/hanabackup` volumes are stored on an NFS server, which is managed by the master host and mounted on all hosts, including the standby hosts.

After a failover completes, the failed host is restarted as a standby host.

SAP supports up to three standby hosts in scale-out systems on Google Cloud. The standby hosts do not count against the maximum of 16 active hosts that SAP supports in scale-out systems on Google Cloud.

Currently, Google Cloud supports SAP HANA host auto-failover on only the SUSE Linux Enterprise Server (SLES) for SAP public images that are available from Compute Engine in the `sles-12-sp3-sap` and `sles-12-sp2-sap` image families. To see the public images that are available from Compute Engine, see [Images \(/compute/docs/images\)](/compute/docs/images).

The following diagram shows a multi-host architecture on Google Cloud that includes support for SAP HANA host auto-failover. In the diagram, worker host 2 fails and the standby host takes over. The `gceStorageClient` plugin works with the SAP Storage Connector API (not shown) to detach the disks that contain the `/hana/data` and `/hana/logs` volumes from the failed worker and to remount them on the standby host, which then becomes worker host 2 while the failed host becomes the standby host.



If Compute Engine live migration, automatic restart, and the high monthly uptime percentage (/compute/sla) of Compute Engine VMs are not enough to meet your availability requirements, Google Cloud also provides Deployment Manager support for the following high-availability features:

- High-availability SUSE Linux Enterprise Server (SLES) cluster for SAP HANA
- SAP HANA scale-out system with SAP HANA host auto-failover

For each of these features, Google Cloud provides a configuration file template that you complete, which Deployment Manager reads to deploy a SAP HANA system for you that is fully supported by SAP and that adheres to the best practices of both SAP and Google Cloud.

For SAP HANA, Deployment Manager deploys a performance-optimized, high-availability Linux cluster that includes:

- Automatic failover
- Automatic restart
- Synchronous replication
- Memory preload
- The Pacemaker high-availability cluster resource manager
- A Google Cloud fencing mechanism
- A VM with the required persistent disks for each SAP HANA instance
- A SAP HANA instance on each VM

For more information, see the [SAP HANA High-Availability Cluster Deployment Guide](/solutions/sap/docs/sap-hana-ha-deployment-guide) (/solutions/sap/docs/sap-hana-ha-deployment-guide).

For a SAP HANA scale-out system that includes the SAP HANA host auto-failover feature, Deployment Manager deploys:

- One master SAP HANA instance
- 1 to 15 worker hosts
- 1 to 3 standby hosts
- A VM for each SAP HANA host
- Persistent disks for the master and worker hosts

A SAP HANA scale-out system with host auto-failover requires an NFS solution, such as Filestore, to share the `/hana/shared` and `/hanabackup` volumes between all hosts. So that Deployment Manager can mount the NFS directories during deployment, you must set up the NFS solution yourself before you deploy the SAP HANA system.

You can set up Filestore NFS server instances quickly and easily by following the instructions at [Creating Instances](/filestore/docs/creating-instances) (/filestore/docs/creating-instances).

To deploy a scale-out system with standby hosts, see the [SAP HANA Scale-Out System with SAP HANA Host Auto-Failover Deployment Guide](/solutions/sap/docs/sap-hana-ha-scaleout-deployment-guide) (/solutions/sap/docs/sap-hana-ha-scaleout-deployment-guide).

For more information from SAP about SAP HANA high availability features, refer to the following documents:

- [SAP HANA – High Availability](https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.02/en-US/6d252db7cdd044d19ad85b46e6c294a4.html)
(<https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.02/en-US/6d252db7cdd044d19ad85b46e6c294a4.html>)
- [FAQ: High Availability for SAP HANA](https://scn.sap.com/docs/DOC-66702) (<https://scn.sap.com/docs/DOC-66702>)
- [How To Perform System Replication for SAP HANA 1.0](https://scn.sap.com/docs/DOC-47702) (<https://scn.sap.com/docs/DOC-47702>)
- [How To Perform System Replication for SAP HANA 2.0](https://assets.cdn.sap.com/sapcom/docs/2016/06/0ec37684-7a7c-0010-82c7-eda71af511fa.pdf)
(<https://assets.cdn.sap.com/sapcom/docs/2016/06/0ec37684-7a7c-0010-82c7-eda71af511fa.pdf>)
- [Network Recommendations for SAP HANA 2.0 System Replication](https://scn.sap.com/docs/DOC-56044)
(<https://scn.sap.com/docs/DOC-56044>)
- [Network Recommendations for SAP HANA 2.1 System Replication](https://archive.sap.com/documents/docs/DOC-58553)
(<https://archive.sap.com/documents/docs/DOC-58553>)

To prepare for disasters, you can use SAP HANA system replication to a secondary SAP HANA system, take backups of SAP HANA to enable recovery, or use both.

For mission critical workloads that require fast recovery times, use HANA system replication to minimize downtime. Using backups to recover a system costs less but takes longer, in that a new system must be created and then the backups restored into it to recover to the desired point in time.

In either case, you must use network-based redirection to redirect client applications that use the SAP HANA system to the IP address of the replacement system once it is available. For more information, see the [SAP HANA Administration Guide](https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.02/en-US)

(<https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.02/en-US>).

Starting with SAP HANA SPS09, you can use the Python-based API included with SAP HANA to create your own high-availability/disaster-recovery (HA/DR) provider and integrate it with the SAP HANA System Replication takeover process to automate tasks like redirecting database client connections from the primary system to the secondary system after a takeover. For more information, see [Implementing a HA/DR Provider](https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.01/en-US/1367c8fdefaa4808a7485b09815ae0f3.html)

(<https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.01/en-US/1367c8fdefaa4808a7485b09815ae0f3.html>)

Note that any restrictions defined by SAP, including distance limitation for synchronous replication, are also in effect on Google Cloud.

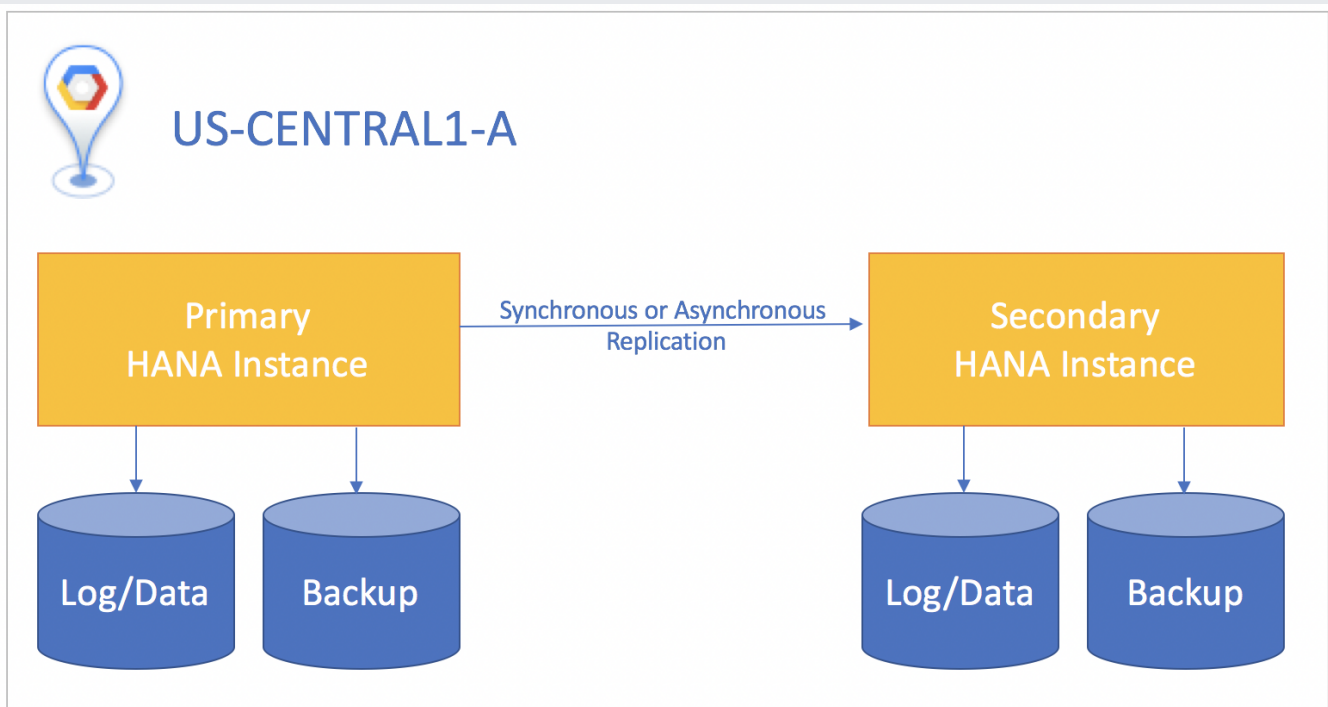
To maximize infrastructure resource utilization and to cost-optimize your DR solution, you can use the secondary system for non-production use cases, such as for a development or QA system. In this case, the secondary system isn't preloaded with data, so the failover time is longer than having the secondary system preloaded with data and kept in sync with the primary system.

HANA 2 SPS00 includes support for Active/Active (read enabled) configuration mode, which enables SAP HANA system replication to support read access on the secondary system. For more information, see [Active/Active \(Read Enabled\)](https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.02/en-US/fe5fc53706a34048bf4a3a93a5d7c866.html).

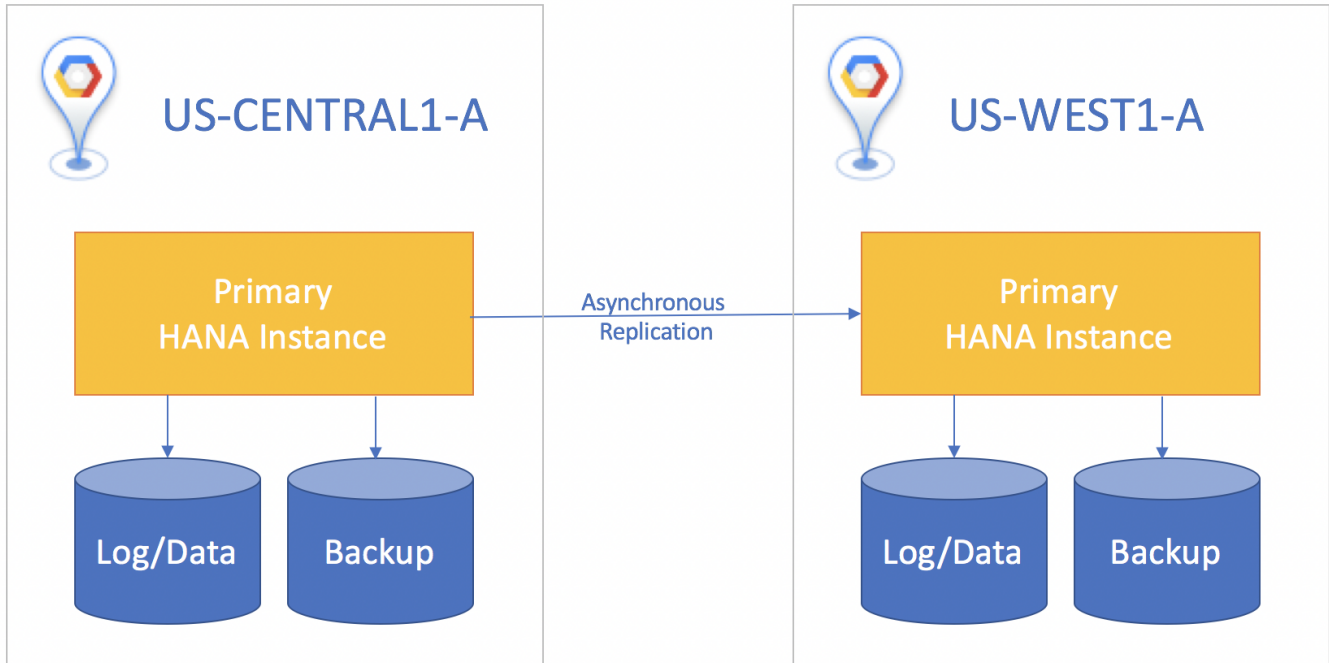
(<https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.02/en-US/fe5fc53706a34048bf4a3a93a5d7c866.html>)

Both synchronous and asynchronous replication are supported when using SAP HANA system replication with Google Cloud.

If possible, we recommend using synchronous replication, where SQL transactions are not committed on the primary database instance until they are committed on the standby instance. This keeps the standby instance 100% in sync and ensures a zero recovery point objective. Synchronous replication can be used for instances that reside in any zones within the same region.



If the standby system is in a different region than the primary system, use asynchronous replication, where there is no wait for the standby instance to acknowledge the data before the commit on the primary instance. In this scenario, you might lose small amounts of data if a disaster happens. A tradeoff is that asynchronous replication gives you a greater than zero recovery point objective.



For all replication scenarios, you must manually perform a takeover on the standby system to start disaster recovery. You also need to manually redirect any applications that use the SAP HANA database to target the instance it has failed over to in the standby system.

Choose the HANA System Replication option that best fits your business needs, such as recovery time objective (RTO), and recovery point objective (RPO). For more information, see [Replication Modes for SAP HANA System Replication](#)

(<https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.02/en-US/c039a1a5b8824ecfa754b55e0caffc01.html>)

In this scenario, your SAP HANA system is replicated to a dedicated standby system. The SAP HANA database is replicated to a Compute Engine VM that has a unique hostname and its own persistent disks attached. All of the SAP HANA data is loaded into memory on the standby system. If you have to failover, the failover time only takes around 90 seconds because all of the data is preloaded.

For more information about SAP HANA System Replication with preload, see the **System Replication** section in [SAP HANA – High Availability](#).

(<https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.02/en-US/6d252db7cdd044d19ad85b46e6c294a4.html>)

In this scenario, your SAP HANA system is replicated to a dedicated standby system. The SAP HANA database is replicated to a Compute Engine VM that has a unique hostname and its own persistent disks attached. The SAP HANA data is not loaded into memory on the standby system. If you have to fail over, the failover time can take from minutes to hours, depending on the size of your dataset.

When you don't preload the data, the memory requirements for the Compute Engine VM that hosts the SAP HANA database are much smaller. The VM only needs either 64 GB of memory, or the amount of memory which is consumed by the rowstore on the target host, whichever is larger. You can get information about the rowstore memory footprint by running the following query:

The reduced memory requirement gives you cost-saving options when choosing a Compute Engine machine type.

- You can use a machine type that has low memory specifications for hosting the SAP HANA database in the standby system to lower your running cost. A low-memory VM isn't supported for SAP HANA in a production system, but you could use this lower-cost VM to perform a takeover in a disaster-recovery scenario, and then can modify the VM afterwards to change the machine type to one with a supported amount of memory. To do this, you must stop the VM to perform the upgrade, and so will have additional downtime before the SAP HANA system is available.
- You can use a high-memory machine type for hosting the SAP HANA database in the standby system, and can share it with development or test systems to improve your return on investment. You can set the global allocation limit for the SAP HANA database to 64 GB by following the instructions at [Change the Global Memory Allocation Limit](https://help.sap.com/viewer/1c837b3899834ddcbae140cc3e7c7bdd/1.0.11/en-US/c1a774b3bb571014af018d2352f0d757.html) (<https://help.sap.com/viewer/1c837b3899834ddcbae140cc3e7c7bdd/1.0.11/en-US/c1a774b3bb571014af018d2352f0d757.html>), leaving the rest of the memory for other systems to use. When the standby system is needed, shut down dev and test operations, perform a takeover, and then remove the global allocation limit.

You can use either synchronous and asynchronous replication without preload. However, synchronous replication requires that the source and target instances be in the same Google Cloud

region.

You can use an HA/DR provider to address issues such as shutting down the development and/or test systems in the secondary host.

To invoke disaster recovery, trigger the SAP HANA System Replication Takeover procedure in your standby system. [SAP Note 2063657](https://launchpad.support.sap.com/#/notes/2063657) (https://launchpad.support.sap.com/#/notes/2063657) provides guidelines to help you decide whether takeover is the best option.

To trigger the takeover, follow the standard SAP HANA takeover process. For more details information about this procedure, see [How To Perform System Replication for SAP HANA 1.0](https://scn.sap.com/docs/DOC-47702) (https://scn.sap.com/docs/DOC-47702) or [How To Perform System Replication for SAP HANA 2.0](https://www.sap.com/documents/2017/07/606a676e-c97c-0010-82c7-eda71af511fa.html) (https://www.sap.com/documents/2017/07/606a676e-c97c-0010-82c7-eda71af511fa.html).

In cases of data issues or software failure, there might not be automatic notifications so that you can perform the takeover. Consider creating a custom solution to send alerts using Stackdriver or HANA monitoring tools.

In cases where a longer recovery time objective is acceptable and your recovery point objective is greater than 15 minutes, you can recover from disaster by restoring from backup. To ensure successful recovery when using backups, make frequent copies of your backup files, especially log backups, to a Cloud Storage bucket, or some other long-term storage location that exists outside of the region where your SAP HANA system runs. We recommend documenting the infrastructure of your primary system and creating scripts that allow you to quickly create a replacement system to restore your backups to.

For more information, see the [SAP HANA operations guide](/solutions/sap/docs/sap-hana-operations-guide#backup_and_recovery) (/solutions/sap/docs/sap-hana-operations-guide#backup_and_recovery).

- For more information about high-availability and disaster recovery for SAP HANA on Google Cloud, see the [SAP HANA Operations Guide](/solutions/sap/docs/sap-hana-operations-guide) (/solutions/sap/docs/sap-hana-operations-guide).
- To deploy a high-availability Linux cluster for SAP HANA, see [SAP HANA High Availability Cluster on SLES Deployment Guide](/solutions/sap/docs/sap-hana-ha-deployment-guide) (/solutions/sap/docs/sap-hana-ha-deployment-guide).

