

High Availability and Disaster Recovery for SAP on Google Cloud Platform Using Veritas InfoScale Enterprise

SAP® Certified
Integration with SAP S/4HANA®

Contents

Abstract	3
Introduction to Veritas InfoScale Enterprise	4
Types of InfoScale Enterprise configurations	4
InfoScale capabilities for SAP ecosystems	5
About Veritas Virtual Business Services (VBS)	5
Flexible Storage Sharing feature of InfoScale Enterprise	6
InfoScale agent for SAP HANA	6
InfoScale SAPNW agent for SAP S/4HANA/NetWeaver	7
ENSA2 Support	7
SystemD Support	8
InfoScale agent for Google IP	8
Typical S/4HANA/NetWeaver Deployment Architecture on GCP	8
Common Dependency Between SAP S/4HANA and SAP HANA Database	9
Cost Optimization of SAP Instances on GCP	10
Supported Use Cases for SAP HANA on GCP	10
SAP HANA instances in the same Availability Zone	10
SAP HANA instances across AZs in the same GCP region	11
Single-instance SAP HANA databases across GCP region	12
SAP HANA database instances across GCP regions (cascading scenario)	14
Configuration Procedure	16
Summary	17
References	18

Abstract

An ever-increasing number of organizations are transitioning their mission-critical applications and services to run not just on-premises but in the public cloud as well. Nowhere is this more evident than with solutions such as SAP HANA and SAP NetWeaver (S/4HANA), which when combined represent nearly 25% of the global ERP market share.

As organizations seek to expand their SAP HANA footprint in the public cloud, the need to address and ultimately improve its availability becomes of paramount importance. In response to this, the HANA database framework provides a built-in high availability (HA) and System Replication (HSR) function. While seemingly adequate for most outage scenarios, when a failure event does in fact occur, the recovery process requires a manual takeover, one which involves server downtime and inevitably leads to loss of revenue opportunities and a potential impact to your business reputation.

The purpose of this document is to provide the necessary guidance for implementing a highly available SAP HANA and SAP NetWeaver / S/4HANA environment deployed in the Google Cloud Platform (GCP) using Veritas InfoScale™ Enterprise.

It is intended for the following audiences:

- Organizations who deploy small SAP systems on GCP for development, testing, training, sandboxing, demonstration, or production purposes, and want to monitor and manage SAP HANA and SAP NetWeaver / S/4HANA based products for high availability and disaster recovery
- SAP consultants who are familiar with GCP and want to manage the availability of their SAP Landscape on GCP using Veritas InfoScale Enterprise

Note: This document is not meant to replace any standard SAP documentation or GCP documentation. For information on basic SAP high availability configurations on GCP, refer to the GCP documentation at:

<https://cloud.google.com/solutions/sap/docs/overview-of-sap-on-google-cloud>

<https://cloud.google.com/solutions/sap/docs/netweaver-ha-planning-guide>

Also, refer to the standard SAP documentation and notes for the respective SAP solution:

<http://service.sap.com/instguides>

<http://service.sap.com/notes>

Introduction to Veritas InfoScale Enterprise

Veritas InfoScale Enterprise is an industry leading high availability and disaster recovery clustering technology which works across leading enterprise operating systems, server virtualization and all cloud technologies.

It is easy to deploy inside GCP instances as well as to maintain and manage within GCP. GCP offers protection from major outages when you deploy applications in multiple Availability Zones (AZs) with 99.99% SLA. However, such solutions are not application-aware, and they do not support traditional clustering for SAP application instances, which depends on 3rd party clustering solutions like Veritas InfoScale Enterprise. Veritas InfoScale Enterprise provides 99.99% SLA for all enterprise applications running inside the GCP instances.

The following components are protected, and any associated single point of failure is eliminated, because they are critical for overall application availability in a distributed SAP environment:

- Databases
 - SAP HANA DB
 - Oracle RDBMS/RAC
 - SAP MaxDB
 - SAP Sybase ASE
 - IBM DB2
 - SQL server
- Central Services instance (ENQUEUE)
- Enqueue replication server (ERS)
- Primary application server (PAS)
- Additional application servers (AAS)

SAP HANA databases and SAP Central services are installed on one instance at a time and are therefore considered single points of failure. Since multiple SAP application server (dialog) instances can run in parallel, they do not form a single point of failure. However, when an application deployed in the same AZ or across AZs fails, GCP restarts the application VM instances or redeploys the instance where the application has failed. This impacts the overall Recovery Time Objective (RTO). InfoScale Enterprise ensures the proper failover of all SAP components to a designated GCP instance on which the application can reside in order to achieve the overall Recovery Time Objective (RTO) and Recovery Point Objective (RPO).

Organizations commonly use OS-specific custom scripts to monitor applications. InfoScale Enterprise supports the largest number of pre-built enterprise application agents that simplify such deployments, thereby eliminating the need for custom scripts. With intelligent failover capabilities, organizations can reduce the cost of redeploying instances by reducing the number of standby servers inside GCP.

An InfoScale disaster recovery (DR) solution is simpler and faster because it provides automated DR and non-disruptive DR testing across GCP regions.

InfoScale Enterprise provides agents for each of the following application and infrastructure components to provide a solution: SAP HANA (SAPHDB) SAP NetWeaver / S/4HANA (SAPNW) SAP Components (SAP Components) Google IP (GCP IP)

Types of InfoScale Enterprise configurations

The various availability and recovery configurations created using InfoScale Enterprise are depicted in Figure 1.

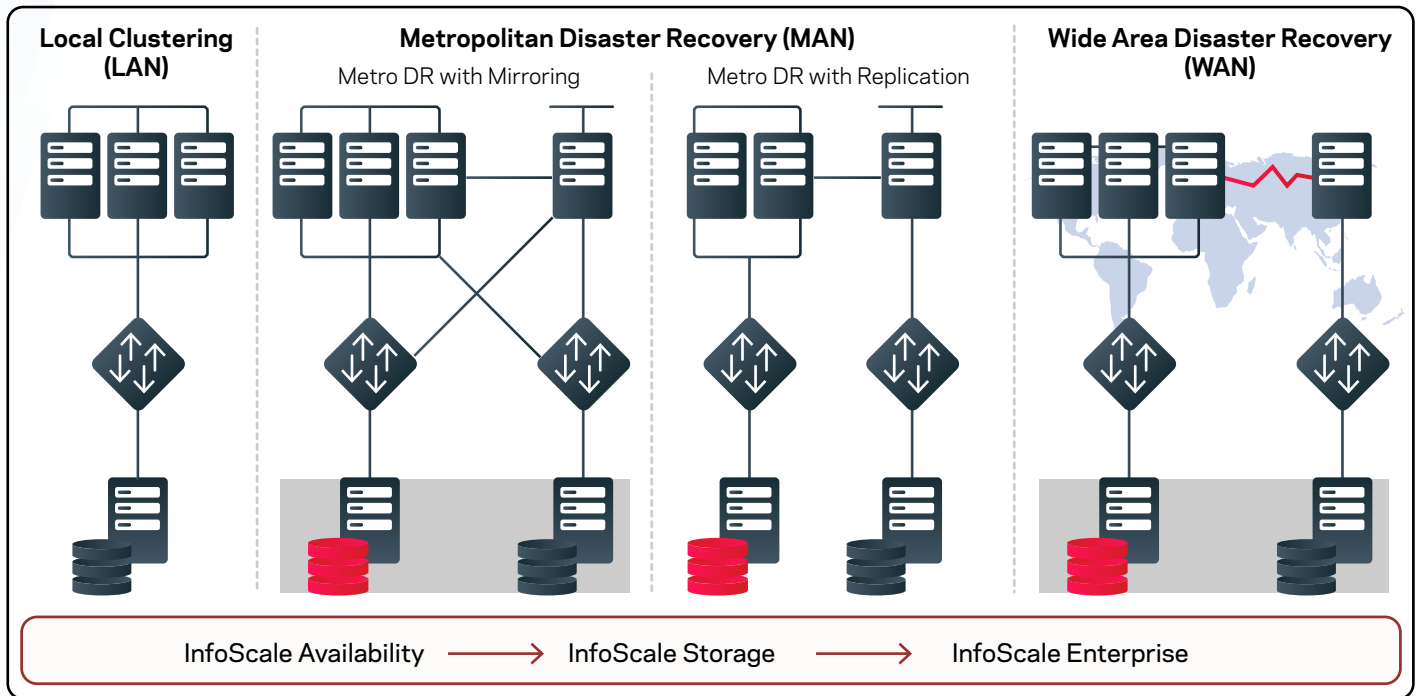


Figure 1. InfoScale Availability and Recovery Configurations

InfoScale capabilities for SAP ecosystems

The following HA and DR capabilities of InfoScale Enterprise make it ideal for managing an SAP ecosystem in Google Cloud:

- Business continuity with minimal application downtime during failures through complete automation for the SAP landscape
- Support for on-premises to GCP failover
- Optimal server utilization due to cost-effective failover configurations for development, testing, or production environments
- Support for cost optimization, Recovery Point Objective (RPO), and Recovery Time Objective (RTO) requirements for SAP workloads in GCP
- Support for disaster recovery (DR), which includes:
 - On-premises to GCP switchover
 - Supported on-premises SAP agents are still relevant in GCP, because they offer:
 - Similar customer experience as on-premises
 - Flexible Storage Sharing (FSS) for data sharing in GCP across instances
 - Replication across GCP AZs and regions

About Veritas Virtual Business Services (VBS)

IT services are no longer standalone applications running on single servers. Business services or multi-tier applications like SAP Business Suite applications make up most of an IT organization's critical services, with different components of the application running on different tiers of infrastructure with their own unique availability needs. A failure at any tier can bring down the entire business service and managing the recovery is time consuming and complex. Virtual Business Services are aware of the complete business service and takes action in the event of a failure to restore the entire service. When an individual component fails, Virtual Business Service provides automated orchestration of the connections to other computing resources, on-site or even across sites. This means faster recovery and minimal downtime—with no manual intervention. Figure 2 shows a sample VBS architecture. 7

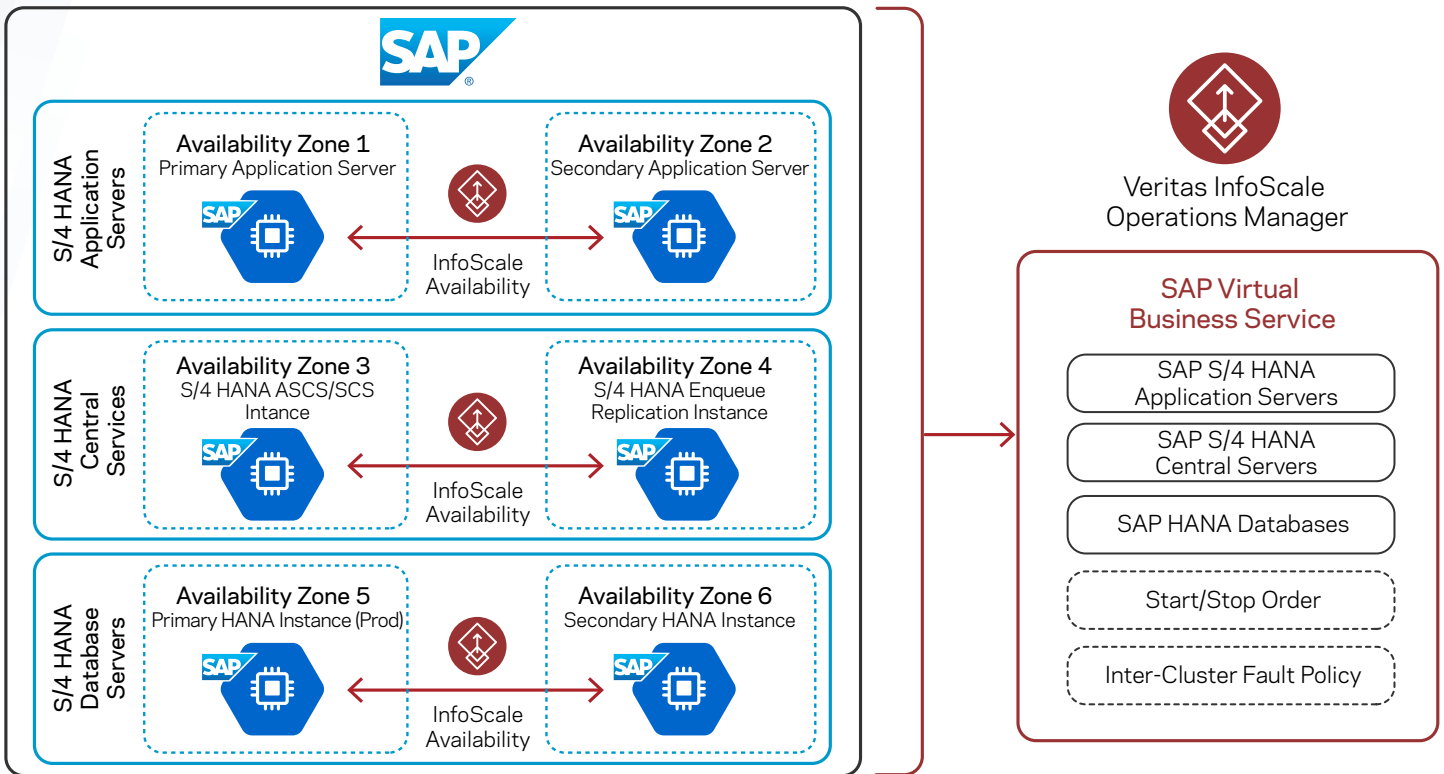


Figure 2. Sample VBS Architecture

Flexible Storage Sharing feature of InfoScale Enterprise

The FSS feature of InfoScale Enterprise lets you combine your distributed, high-performance, and highly available file systems with the latest storage and networking technologies. FSS lets you to unlock the true potential of Direct Attached Storage (DAS), without sacrificing performance or availability. It enables you to drive up to four times the performance at less than 20 percent of the cost of a traditional Storage Area Network (SAN) environment. FSS is not limited to a DAS-only deployment but can be used in conjunction with SAN in a hybrid deployment. Thus, you can also continue to use the GCP volumes.

Note: Veritas InfoScale FSS supports all GCP storage volumes and can be used for SAP NetWeaver (S/4HANA) Application server instance mounts like sapmnt, trans, SAP Central service instance (ENQUEUE) and Enqueue Replication instance (ERS) mounts for fast failover and high availability.

For details, see the Veritas FSS datasheet at: https://www.veritas.com/content/dam/Veritas/docs/data-sheets/21327034_GA_ENT_DS-Veritas-Flexible-Storage-Sharing-EN.pdf

InfoScale agent for SAP HANA

The Cluster Server agents monitors specific resources within an enterprise application. They determine the status of resources and start or stop them based on external events that may affect application availability.

The Cluster Server agent for SAP HANA (SAPHDB) provides high availability for HANA instances where the data is replicated with HANA System Replication. The agent brings a HANA instance online, monitors the instance, and takes the instance offline. It also monitors the system processes and the server state and shuts down the server in case of a failover.

The SAPHDB agent supports the following HANA features:

- Fast failover of faulted Instance within an AZ (Availability Set)
- Primary instance takeover in case of an AZ failover
- Re-registration of the old primary as the secondary
- Auto-restart of a HANA instance before takeover
- In-depth trace level monitoring and IMF support

The agent supports the following types of HANA System Replication scenarios:

- Replication between AZs within the same GCP region (local failover)
- Replication between AZs across regions (instances in different regions)
- Replication across AZs within one GCP region and then to an AZ in a different region (two instances in the same region and a third instance in a remote region)

InfoScale SAPNW agent for SAP S/4HANA/NetWeaver

Veritas InfoScale SAPNW agent is a certified cluster server agent from the SAP ICC program. The Cluster Server agent for SAP, S/4HANA (SAPNW), provides high availability for S/4HANA instances. The agent can monitor and manage the status (online/offline) of an S/4HANA instance. It also monitors the system processes and the server state and shuts down the instance in case of a failover.

The SAPNW agent supports the following S/4HANA / NetWeaver features:

- Fast Failover of faulted instances within an AZ (Availability Zone)
- Primary instance takeover in case of an AZ failover
- Auto-restart of an S/4HANA instance before takeover
- In-depth monitoring and IMF support
- Support for in-depth TRACE and Debug Log levels for troubleshooting

The agent supports the following types of S/4HANA / NetWeaver scenarios:

- High availability between AZs within the same GCP region (local failover)
- Disaster Recovery between AZs across regions (instances in different regions)
- Availability Across AZs within one GCP region and then to an AZ in a different region (two instances in the same region and a third instance in a remote region)

ENSA2 Support:

Under high availability it is mandatory that the old mechanism of Standalone Enqueue Server (ENSA1) has to fail over to the cluster node where the active ERS is running in order to acquire a replicated enqueue table which resides in the shared memory of the active ERS node.

The new Standalone Enqueue Server 2 and Enqueue Replicator 2 provide an improved high availability architecture with robust and fast replication, and failover.

In ENSA2, if the ASCS fails it can start on a separate node in the cluster and copy the lock entries from the enqueue replicator 2, it is not mandatory that it failover over to the active ERS2 node.

The Veritas InfoScale SAPNW agent supports both of the Enqueue replication methods and is certified by SAP.

The following is a list of Veritas InfoScale certifications for SAP HA:

Veritas InfoScale Availability / InfoScale Enterprise 7.4 for AIX (with SAPNW agent) was certified with the SAP ICC integration scenario S/4-HA-CLU 1.0 (SAP S/4HANA 1809) in 2020

- Veritas InfoScale Availability / InfoScale Enterprise 7.4 for Linux (with SAPNW agent) was certified with the SAP ICC integration scenario S/4-HA-CLU 1.0 (SAP S/4HANA 1809) in 2019
- Veritas InfoScale Availability / InfoScale Enterprise 7.4 for Linux (with SAPNW agent) was certified with the SAP ICC integration scenario NW-HA-CLU 7.50 in 2019
- Veritas InfoScale Availability / InfoScale Enterprise 7.4 for AIX (with SAPNW agent) was certified with the SAP ICC integration scenario NW-HA-CLU 7.50 in 2019
- Veritas Cluster Server 6.2 for AIX (with SAPNW agent) was certified with the SAP ICC integration scenario NW-HA-CLU 7.30 in 2015
- Veritas Cluster Server 6.2 for Solaris (with SAPNW agent) was certified with the SAP ICC integration scenario NW-HA-CLU 7.30 in 2015
- Veritas InfoScale Availability / InfoScale Enterprise 7.0 for Linux (with SAPNW agent) was certified with the SAP ICC integration scenario NW-HA-CLU 7.40 in 2017

An overview of the SAP NetWeaver High Availability Cluster (NW-HA-CLU) integration certification is also available at <https://wiki.scn.sap.com/wiki/display/SI/Certified+HA-Interface+Partners>.

SystemD Support:

SystemD is a system and service manager for the latest Enterprise Linux operating systems. It manages the application operations in system space. By default, SAP Application servers run in user space (init). As such, the SAP application servers do not stop gracefully during system reboots which cause application crashes in high availability systems.

Veritas InfoScale SAPNW/SAPHDB agents can manage restarts of SAP applications gracefully without any crashes during system reboots.

InfoScale agent for Google IP

InfoScale Enterprise provides the GoogleIP agent, which lets you monitor and manage the following networking resources in GCP:

- Private IP: A private IP is a private numerical address that networked devices use to communicate with one another
- GCP IP: A GCP IP address is a static IPv4 address designed for dynamic cloud computing, and it is associated with your GCP account
- Overlay IP: GCP allows you to redirect IP address traffic to a GCP instance in a Virtual Private Network (VPC) regardless of the subnet or AZ to which it belongs. An overlay IP lets you fail over IP addresses between cluster nodes when they are spread across multiple subnets or AZs

Typical S/4HANA / NetWeaver Deployment Architecture on GCP

Figure 3 describes the overall deployment of an SAP S/4HANA application server with an SAP HANA database on GCP.

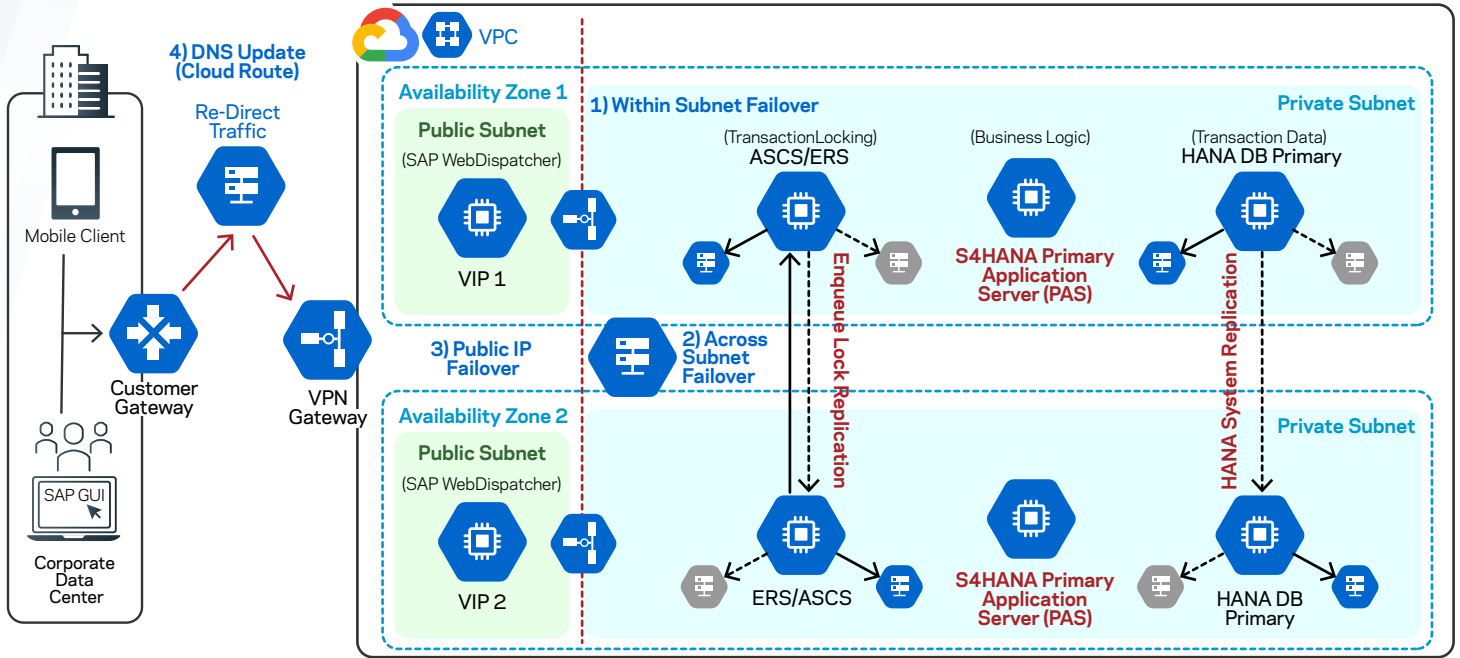


Figure 3. Deployment of an SAP S4HANA application server with an SAP HANA database on GCP

In addition to configuring SAP HANA for HA, DR, or both, you need to ensure that the client applications (for example, the S/4HANA / NetWeaver application server, JDBC, ODBC connection, and so on) can re-establish their connection with the SAP HANA system after the failover. To do so, you can configure either network-based IP redirection or network-based DNS redirection of your SAP HANA system. InfoScale agents for SAP support end-to-end HA and DR for SAP NetWeaver and SAP S/4HANA.

Common Dependency Between SSP S/4HANA and SAP HANA Database

Figure 4 depicts the typical dependency between an SAP application and an SAP HANA database:

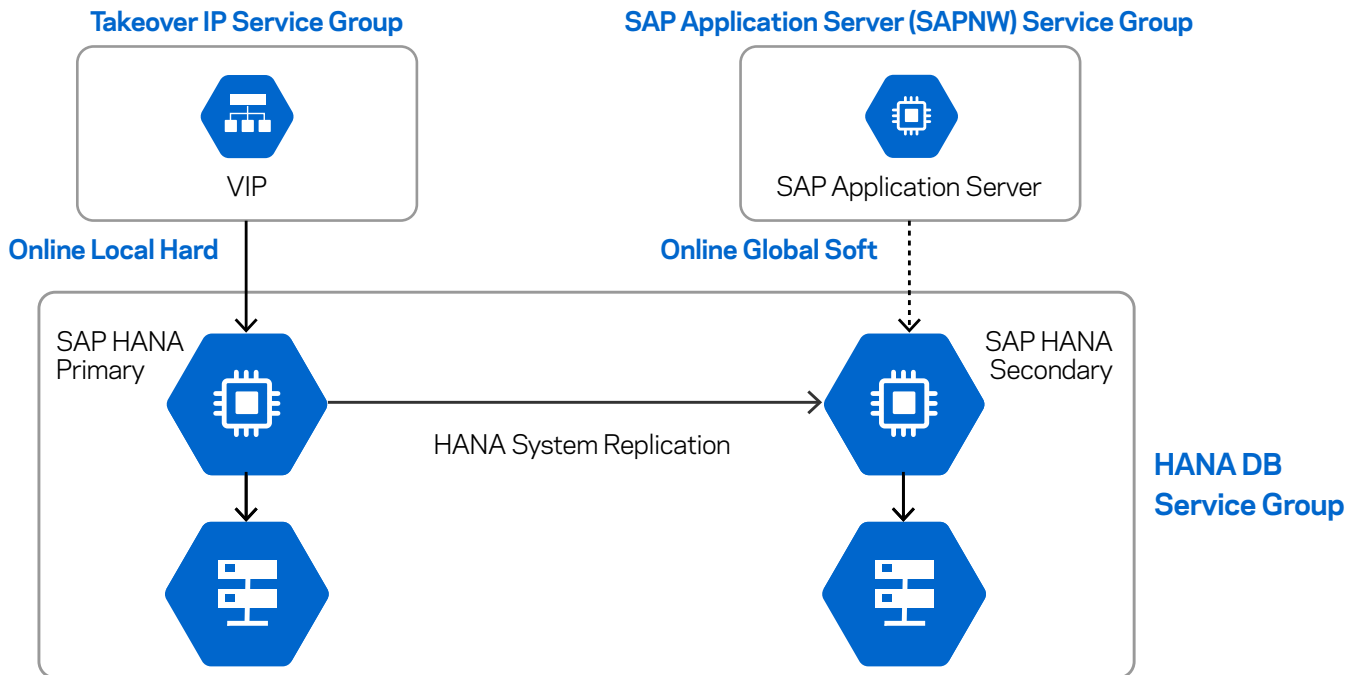


Figure 4. Dependency between an SAP application and SAP HANA database

Cost Optimization of SAP Instances on GCP

GCP lets you create and optimize SAP instances for development, testing, or production environments. If an SAP application instance outage occurs, InfoScale Enterprise fails over the instances between the designated SAP systems without disrupting the client connections. Thus, it reduces the overall Total Cost of Ownership (TCO) in case of a disaster or an outage of SAP instances on GCP.

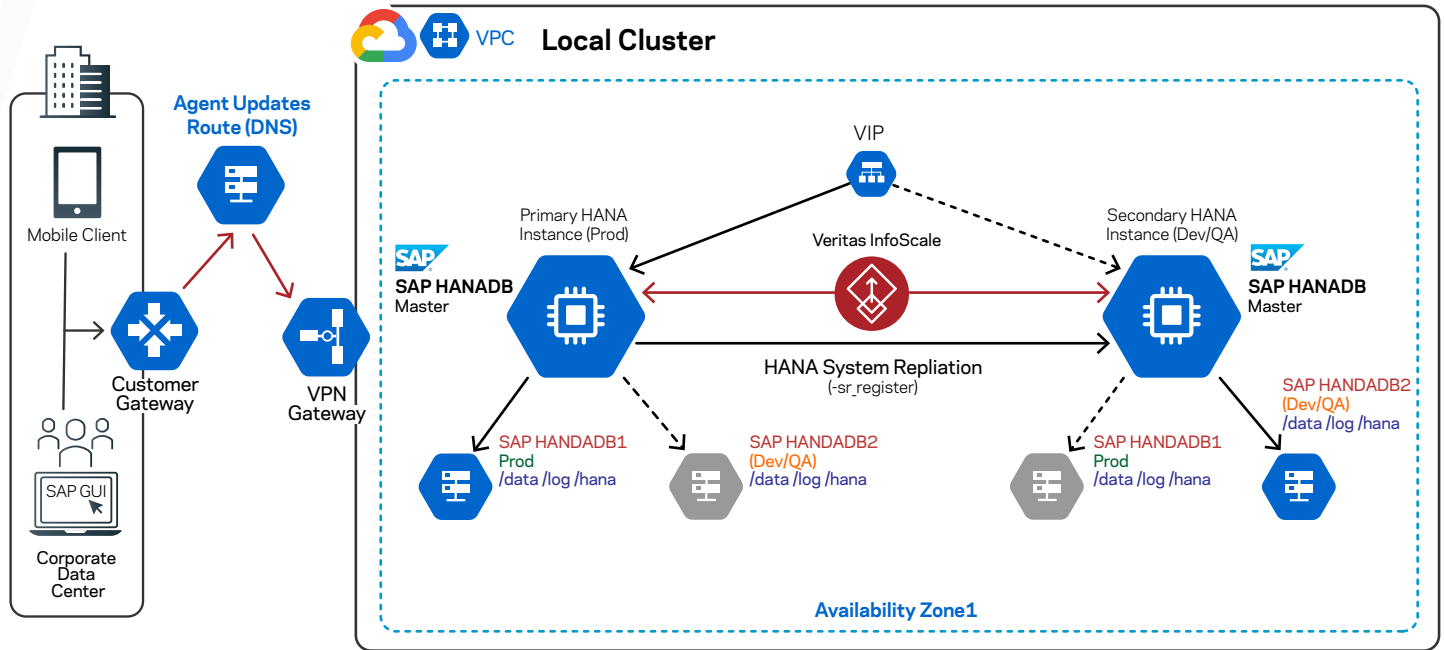


Figure 5. If an SAP HDB instance in production fails, InfoScale will detect the SAP HDB instance failures and non-disruptively fail over the instances between the designated SAP systems

Supported Uses Cases for SAP HANA on GCP

The following graphics represent the availability configuration scenarios supported by InfoScale for SAP HANA. If an SAP HDB instance in a production environment fails, InfoScale detects the failure and moves that instance to a test/development node.

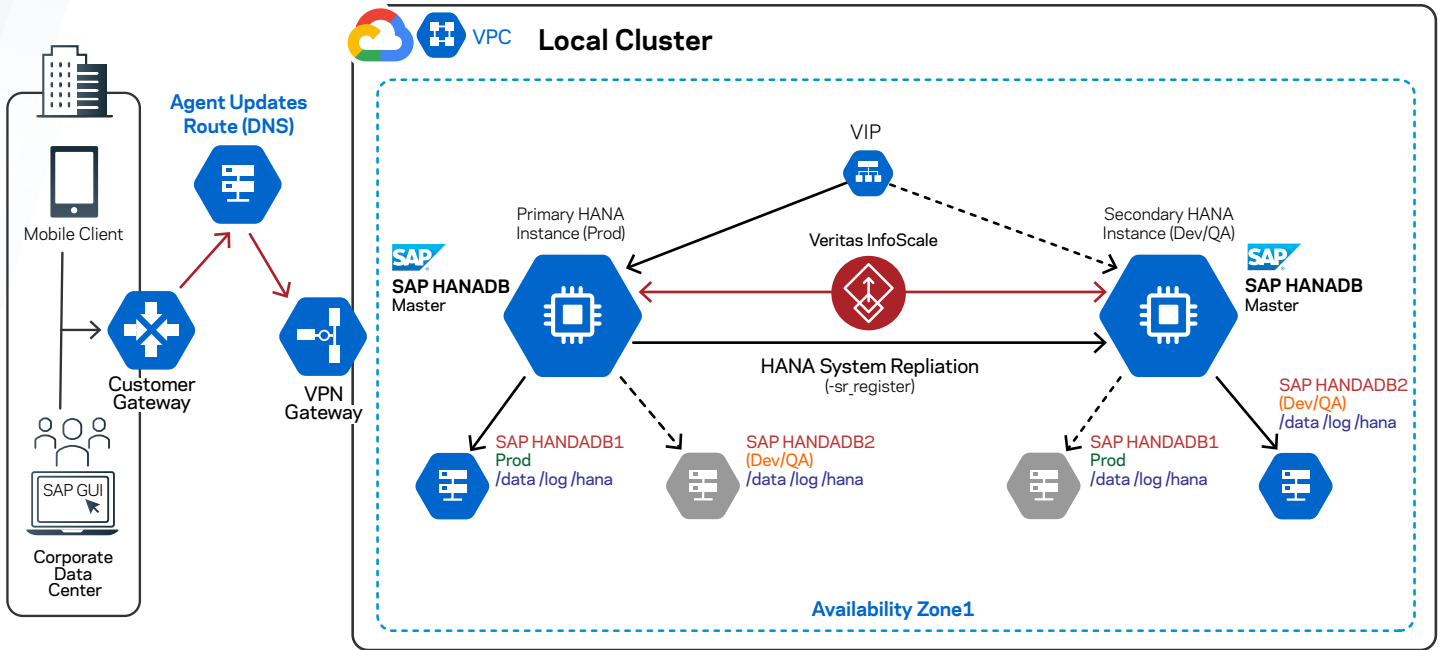
InfoScale Enterprise lets you monitor and control SAP HANA database instances for the following use cases:

- SAP HANA instances in the same AZ
- SAP HANA instances across AZs in the same GCP region
- Single-instance HANA databases across GCP regions
- SAP HANA database instances across GCP regions (cascading scenario)

SAP HANA instances in the same Availability Zone

In this scenario, a master instance (primary) and a standby instance (secondary) are configured along with HANA System Replication in the same AZ.

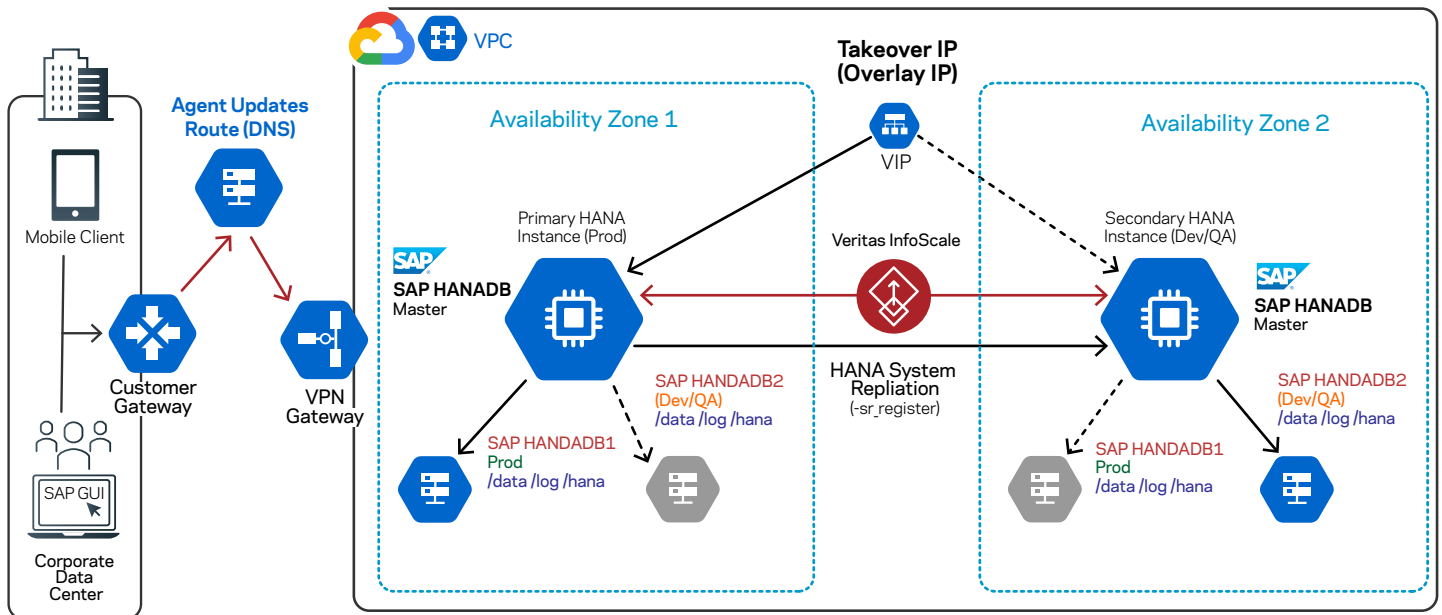
If the master instance fails, the SAPHDB agent promotes the secondary to the primary.



Note: The failover operations are performed according to the HANA HA takeover guidelines and the SAPHDB agent adheres to the HANA Takeover rules.

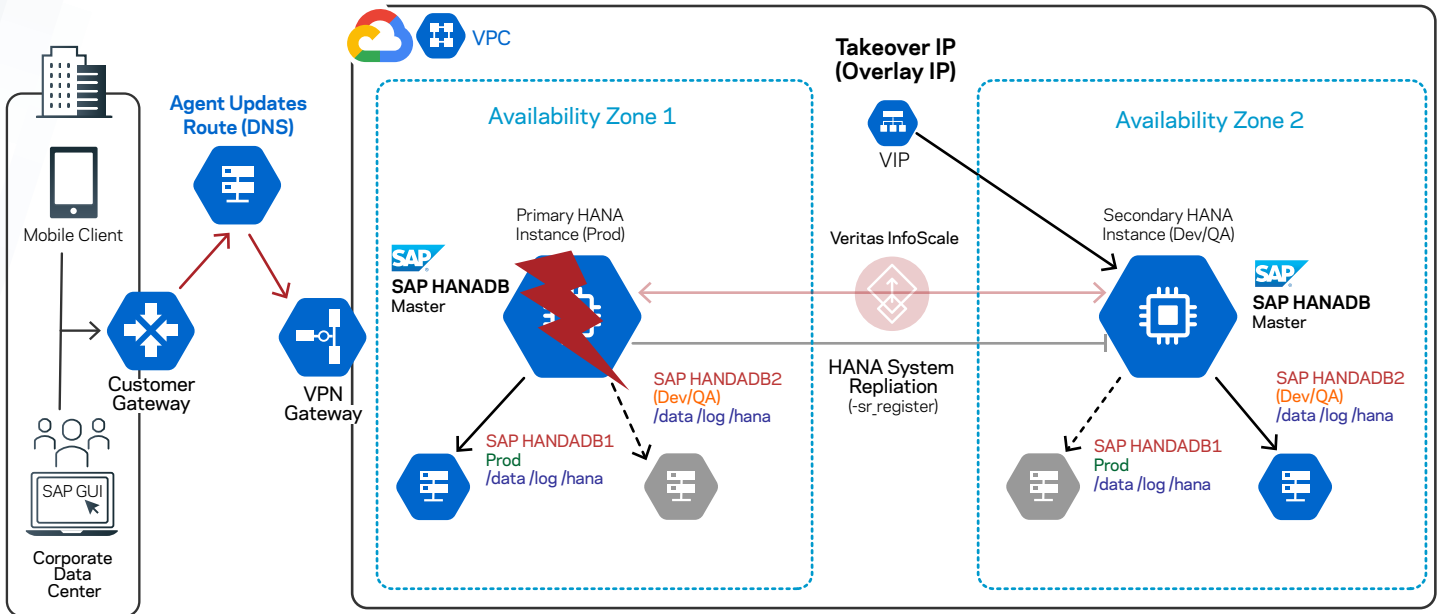
SAP HANA instances across AZs in the same GCP region

In this configuration, the primary and the secondary instances of SAP HANA exist either in the same AZ or in different AZs in the same region. When HANA System Replication is enabled between the two instances, all the data and the logs are replicated to the secondary instance.

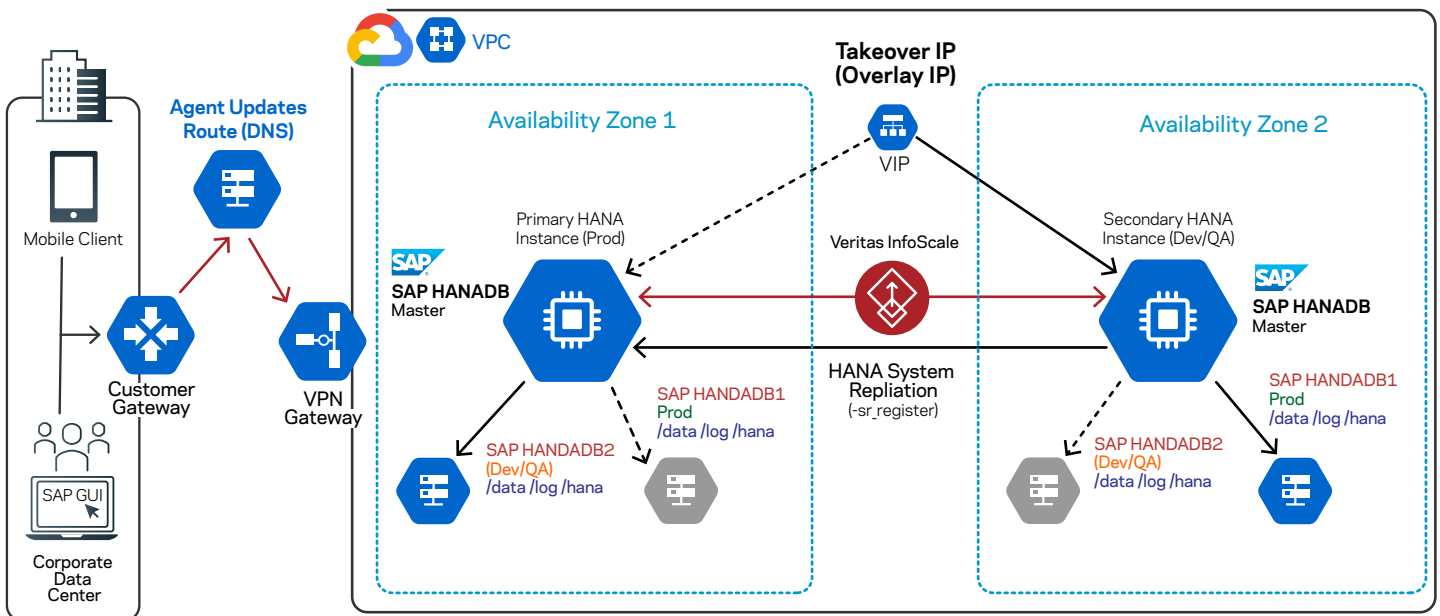


In this example, the primary (master) and the secondary (worker) instances are configured in different AZs with local clustering.

If the primary instance fails or becomes unavailable, the SAPHDB agent identifies the fault and automatically triggers the takeover operation on the secondary instance. The following graphic depicts this action:

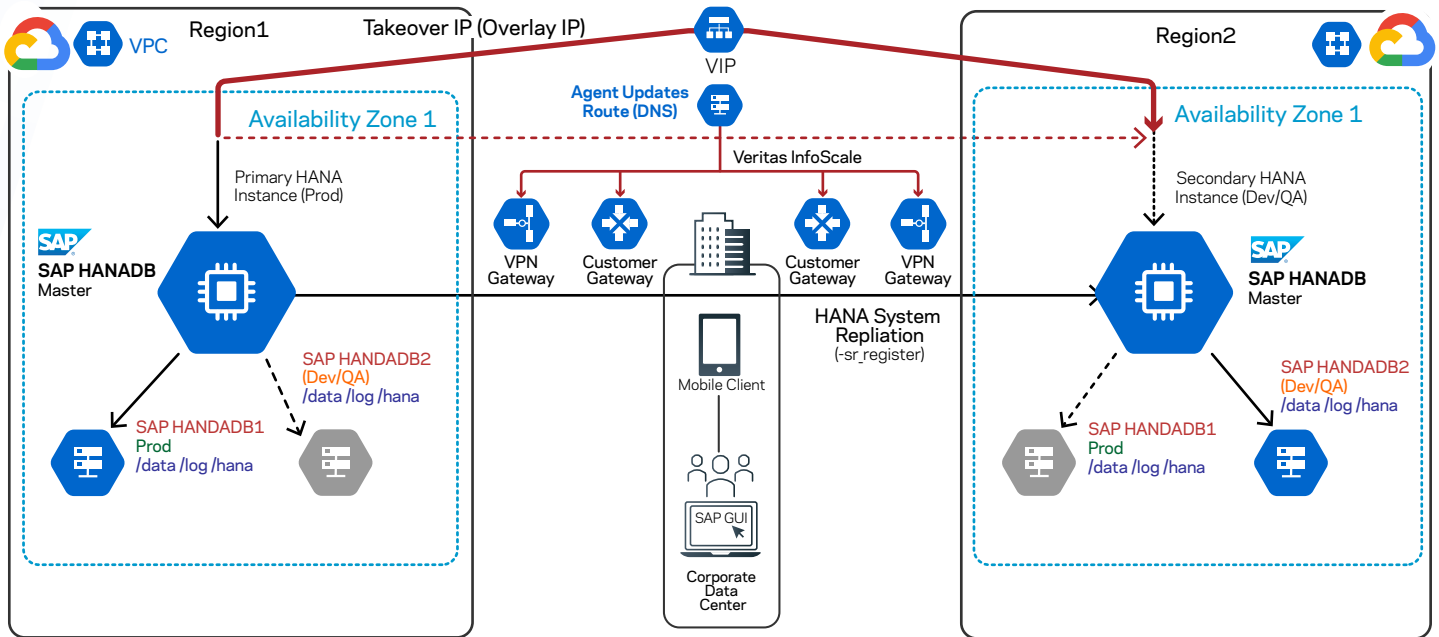


You may choose to clear the fault and perform any other necessary maintenance activities on the old primary instance. Thereafter, the SAPHDB agent can automatically designate it as the secondary instance by using the auto re-register feature. HANA System Replication can then continue the data replication in the reverse direction. The following graphic depicts this action:

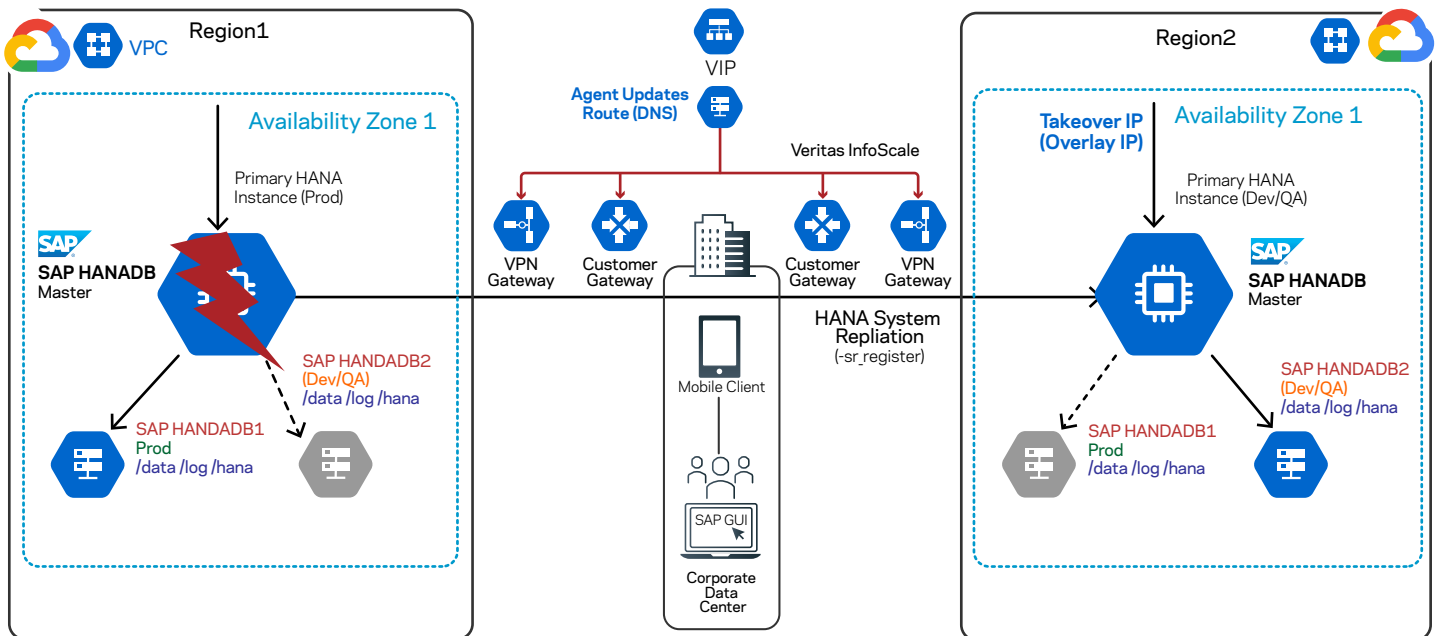


Single-instance SAP HANA databases across GCP regions

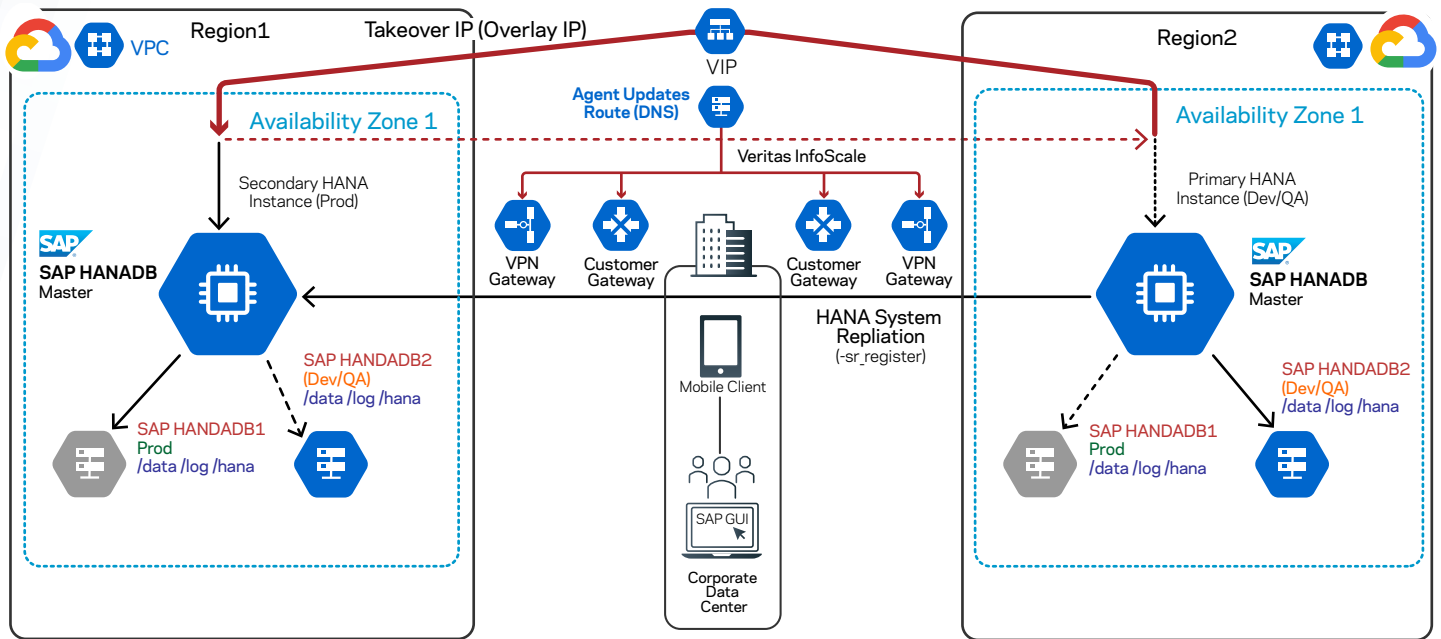
You can configure single HANA database instances across GCP regions, and they can be controlled and monitored by InfoScale agents using a GCO configuration. The following graphic depicts such a configuration:



If the primary instance fails or becomes unavailable, the SAPHDB agent identifies the fault and automatically triggers the takeover operation on the secondary instance in the other region. The following graphic depicts this action:



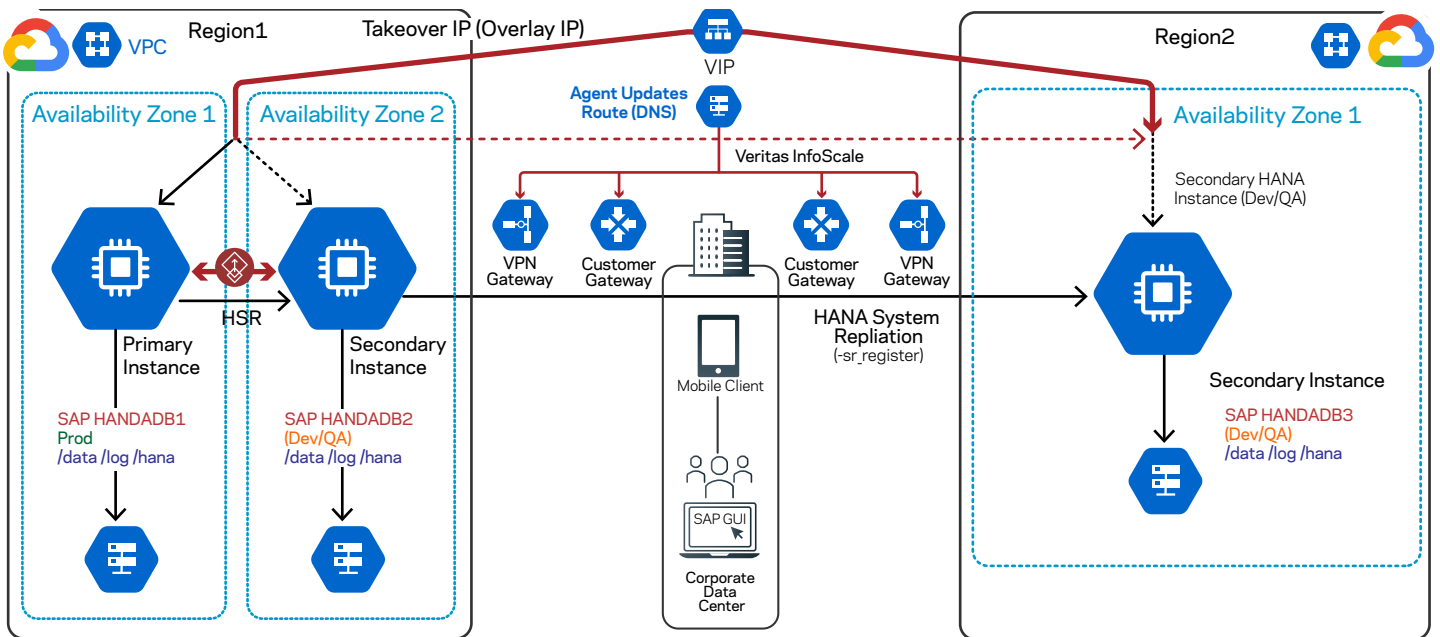
You'll need to clear the fault and perform any other necessary maintenance activities on the old primary instance. Thereafter, the SAPHDB agent automatically designates it as the secondary instance by using the auto re-register feature. The following graphic depicts this action:



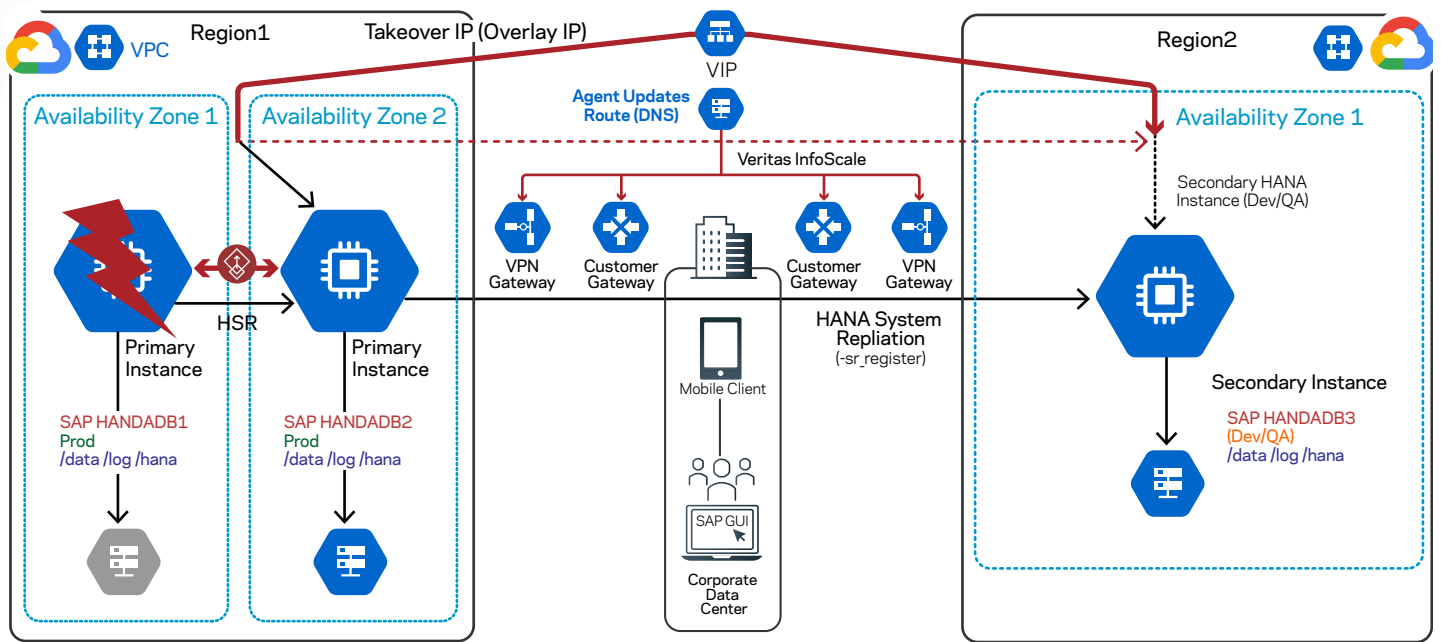
SAP HANA database instances across GCP regions (cascading scenario)

In this scenario, the HANA database primary and secondary are configured in AvailabilityZone1 and AvailabilityZone2 respectively in Region1. The third HANA database instance is in Region2.

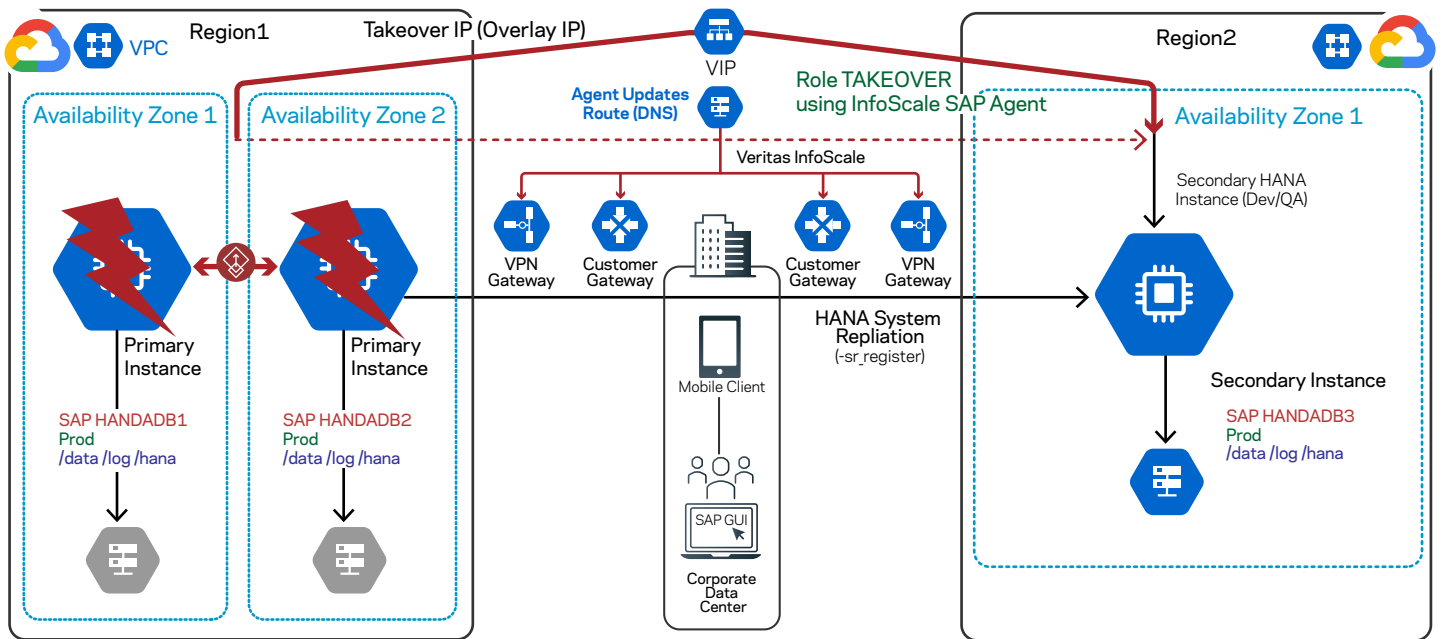
The following graphic depicts this cascading configuration:



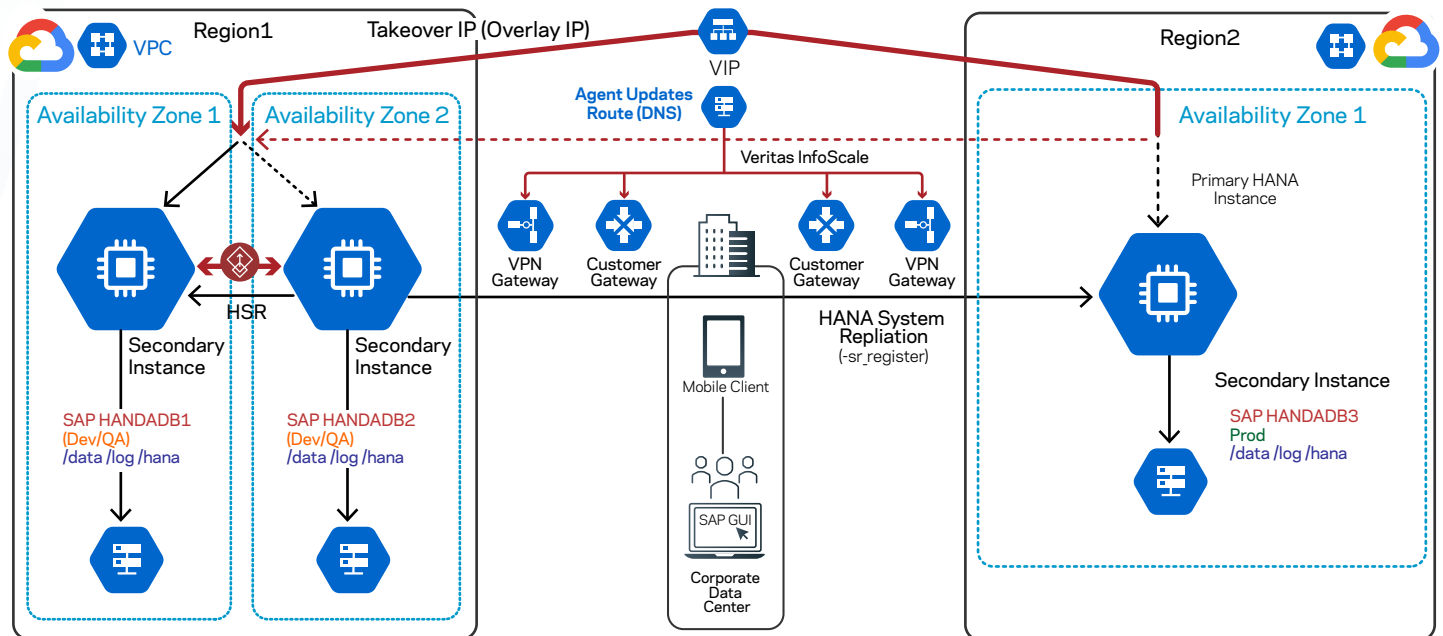
If the primary instance fails, the SAPHDB agent automatically triggers the takeover by failing over the IP resource from the primary to the secondary instance in the same region. Virtual IP failover operations are managed by the GoogleIP and the IP agents. The following graphic depicts how the secondary becomes the primary and is active in Availability Zone 2:



If all the instances within in an AZ or a region fail, the SAPHDB agent automatically triggers the takeover action by failing over the IP resource from the primary to the secondary instance in the remote region. The following graphic depicts how the secondary in the remote region becomes the primary:



Thereafter, you'll need to clear the faults on the instances in the Region1 and re-register them with the current primary in Region2. The following graphic depicts the reversed replication direction:



For details, refer to the InfoScale agent installation and configuration documentation at: <https://sort.veritas.com/agents/>

Configuration Procedure

1. Configure a VPN to connect the GCP cloud.

Follow the procedure as per GCP documentation

2. Create GCP instances based on the planning document and the sizing of Instances recommended by GCP for SAP systems.
3. Install and configure Veritas InfoScale Enterprise on all the systems that must be part of the cluster.

For details, refer to the Veritas InfoScale Installation Guide.

4. Allocate SAP-recommended Storage.
5. Prepare the cloud environment as follows:
 - a) Create the GCP instances.
 - b) Allocate SAP-recommended Storage.
 - c) Mount the following SAP mount points by using SAP recommended file system for HANA DB:
 - /hana/shared/
 - /hana/data/
 - /hana/log/

6. If installing and configuring SAP S/4HANA / NetWeaver for use with HANA, use InfoScale Enterprise for storage configuration for /sapmnt, /usr/sap directories.

Note: Veritas InfoScale FSS supports all GCP storage services. VxVM, VxFS, and FSS are supported and can be used with S/4HANA / SAP NetWeaver application servers as well as most database management systems, however these utilities are not currently supported for use with SAP HANA databases.

For details, refer to the following InfoScale documentation:

Storage Foundation and High Availability Configuration and Upgrade Guide
Storage Foundation Cluster File System High Availability Administrator's Guide
Storage Foundation Cluster File System High Availability Configuration and Upgrade Guide
Veritas InfoScale Disaster Recovery Implementation Guide

7. Install and configure SAP HANA in the GCP instances on primary and secondary sites as per your disaster recovery plan.
8. Configure SAP HANA system replication between the sites on GCP. Ensure that all the required ports are enabled for replication.
9. Configure InfoScale cluster service groups and resources for the SAP HANA instances.

For details, refer to the following InfoScale documentation:

Storage Foundation and High Availability Configuration and Upgrade Guide

VCS saphdb install guide: https://sort.veritas.com/agents/download_docs/19450/vcs_saphdb_install

Summary

With SAP HANA being a popular option for enterprise SAP deployments, Veritas InfoScale is a certified solution to ensure HA as well as simplified DR automation and testing for HANA databases running in GCP. InfoScale has direct integration with SAP HANA and GCP native tools, giving it visibility into all of the components and processes that need to be managed as part of an HA configuration for maximum application uptime. Some key benefits of using InfoScale to manage HA and DR for SAP HANA in GCP are:

- Near-instant fault detection that provides minimal RTO and RPO for HANA databases
- Automation of the entire failover process and non-disruptive DR testing across GCP regions
- Flexible configuration options that support multiple usage scenarios within GCP zones and regions
- Support for the broader SAP ecosystem that may include non-HANA SAP deployments

With the ability to provide best-in-class architectural flexibility, availability and resiliency for SAP HANA, InfoScale enables businesses to improve upon SAP application SLA's while reducing infrastructure footprints by integrating cloud into their IT strategy. Whether running on-premises, in a hybrid cloud configuration or entirely within a cloud environment, InfoScale is an enterprise software-defined availability and resiliency solution for SAP HANA in GCP that provides the tools needed to run SAP applications with maximum uptime.

References

SAP notes for SAP on GCP:

- [1588667](#): Overview of related SAP Notes and Web-Links
- [1656099](#): Supported SAP, DB/OS and GCP products
- [1656250](#): Support prerequisites
- [171356](#): SAP Software on Linux: Essential information

GCP documentation:

- [SAP on Google Cloud](#)
- [Overview of SAP on Google Cloud](#)
- [Reference Architecture: SAP S/4HANA on Google Cloud Platform](#)
- [SAP HANA High Availability and Disaster Recovery Planning Guide on Google Cloud](#)

Veritas documentation:

- [Veritas InfoScale](#)
- [Veritas InfoScale Trial License](#)
- [Veritas InfoScale Linux Documentation](#)
- [Veritas InfoScale Agents Getting Started Guide](#)

About Veritas

Veritas Technologies is a global leader in data protection and availability. Over 80,000 customers—including 87 percent of the Fortune Global 500—rely on us to abstract IT complexity and simplify data management. The Veritas Enterprise Data Services Platform automates the protection and orchestrates the recovery of data everywhere it lives, ensures 24/7 availability of business-critical applications, and provides enterprises with the insights they need to comply with evolving data regulations. With a reputation for reliability at scale and a deployment model to fit any need, Veritas Enterprise Data Services Platform supports more than 800 different data sources, over 100 different operating systems, more than 1,400 storage targets, and more than 60 different cloud platforms. Learn more at www.veritas.com. Follow us on Twitter at [@veritastechllc](https://twitter.com/veritastechllc).

VERITAS™

2625 Augustine Drive
Santa Clara, CA 95054
+1 (866) 837 4827
veritas.com

For global contact
information visit:
veritas.com/company/contact