

High Availability and
Disaster Recovery for
SAP HANA on AWS
using Veritas InfoScale™
Enterprise

SAP® Certified
Integration with SAP S/4HANA®

Table of Contents

ABSTRACT	3
INTRODUCTION TO VERITAS INFOSCALE ENTERPRISE	4
Types of InfoScale Enterprise configurations.....	5
InfoScale capabilities for SAP ecosystems.....	6
About Veritas Virtual Business Services (VBS).....	6
Flexible Storage Sharing feature of InfoScale Enterprise.....	7
InfoScale agent for SAP HANA.....	8
InfoScale SAPNW agent for SAP S4/HANA / NetWeaver.....	8
InfoScale agent for AWS IP.....	10
InfoScale agent for AWS Route 53.....	10
TYPICAL S/4HANA / NETWEAVER DEPLOYMENT ARCHITECTURE ON AWS	11
COMMON DEPENDENCY BETWEEN SAP S/4HANA AND SAP HANA DATABASE	12
COST OPTIMIZATION OF SAP INSTANCES ON AWS	12
SUPPORTED USE CASES FOR SAP HANA ON AWS	13
SAP HANA instances in the same AZ.....	13
SAP HANA instances across AZs in the same AWS region.....	14
Single-instance SAP HANA databases across AWS regions.....	16
SAP HANA database instances across AWS regions (cascading scenario).....	18
Configuration Procedure.....	20
SUMMARY	21
REFERENCES	22

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 AWS cloud using Veritas InfoScale™ Enterprise

It is intended for the following audiences:

- Organizations who deploy SAP systems on AWS 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 Basis and SAP implementation consultants who are familiar with AWS and want to manage the availability of SAP systems on AWS using Veritas InfoScale Enterprise.
- This document does not replace any standard SAP or AWS documentation. For information on basic SAP high availability configurations on AWS, refer to the AWS documentation at: <https://aws.amazon.com/sap/>

When installing SAP solutions on AWS, always refer to the standard SAP documentation and notes for the respective SAP solution:

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

For more information about SAP on AWS, refer to the AWS documentation at:

<http://aws.amazon.com/sap>

INTRODUCTION TO VERITAS INFOSCALE ENTERPRISE

Veritas InfoScale Enterprise represents over 30 years of innovation and market leadership in both Software Defined Storage, High Availability as well as Disaster Recovery. Composed of Veritas Volume Manager, Veritas File System and Veritas Cluster Server, InfoScale Enterprise is responsible for protecting the world's most mission-critical applications and databases across several industry verticals. Whether it's the financial sector, healthcare or state and local government, InfoScale enables customers to rapidly deploy highly available, performant, multi-tier services irrespective of the operating system, underlying storage, compute infrastructure or cloud platform. Moreover, InfoScale provides extensive integration with S4/HANA and SAP HANA database environments, with an emphasis on automation, data integrity, scale and minimizing downtime.

When considering whether or not to deploy InfoScale Enterprise within the AWS public cloud, you'll want to recognize that we provide not only an intuitive configuration model, with inclusion in the Marketplace, but a more granular, application-centric view of your critical services. AWS, by design, addresses the potential for infrastructure outages only, if applications are installed across multiple Availability Zones (AZ's), for 99.99% uptime. What this model does not account for however is the need to be application aware and thus AWS necessitates a dependency on 3rd party clustering solutions. With InfoScale, you gain out-of-the-box availability support for all SAP HANA and S4/HANA components while simultaneously benefiting from the resilience of the AWS infrastructure.

InfoScale Enterprise is easily deployed and managed inside EC2 instances on AWS, thereby offering operational parity with similar configurations as on-premises. AWS does not directly integrate with the embedded resiliency features from SAP, which in turn requires the use of a solution like InfoScale to automate the recovery procedures.

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
 - Microsoft 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 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, AWS restarts the application

VM instances or redeploys the EC2 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 or alternate EC2 instance on which the application can reside in order to achieve the overall Recovery Time Objective (RTO) and Recovery Point Objective (RPO).

Organizations often resort to OS-specific or customized scripts in order to monitor their critical applications. InfoScale Enterprise supports the largest catalog of pre-built enterprise agents that simplify such deployments, thereby eliminating the need for custom monitoring tools. With intelligent failover capabilities, organizations can reduce the cost of redeploys instances by reducing the number of standby or passive servers within AWS.

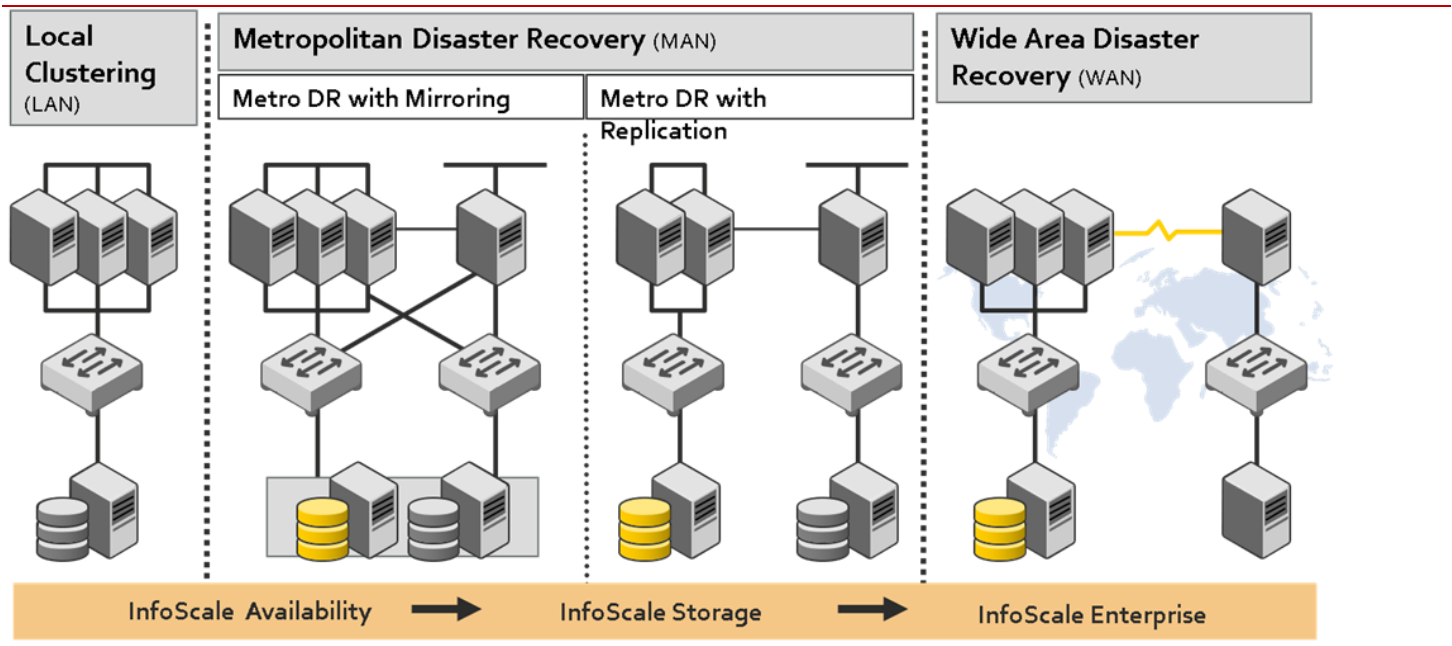
Configuring high-availability (HA) and disaster recovery (DR) with InfoScale is inherently simpler as well as faster for critical SAP workloads. Moreover, the need for non-disruptive recovery validation across AWS regions can be met with InfoScale's FireDrill DR testing capability.

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)
- AWS IP (AWSIP)
- AWS Route 53 (AWSRoute53)

Types of InfoScale Enterprise configurations

The following graphic depicts various availability and recovery configurations created using InfoScale Enterprise:



InfoScale capabilities for SAP ecosystems

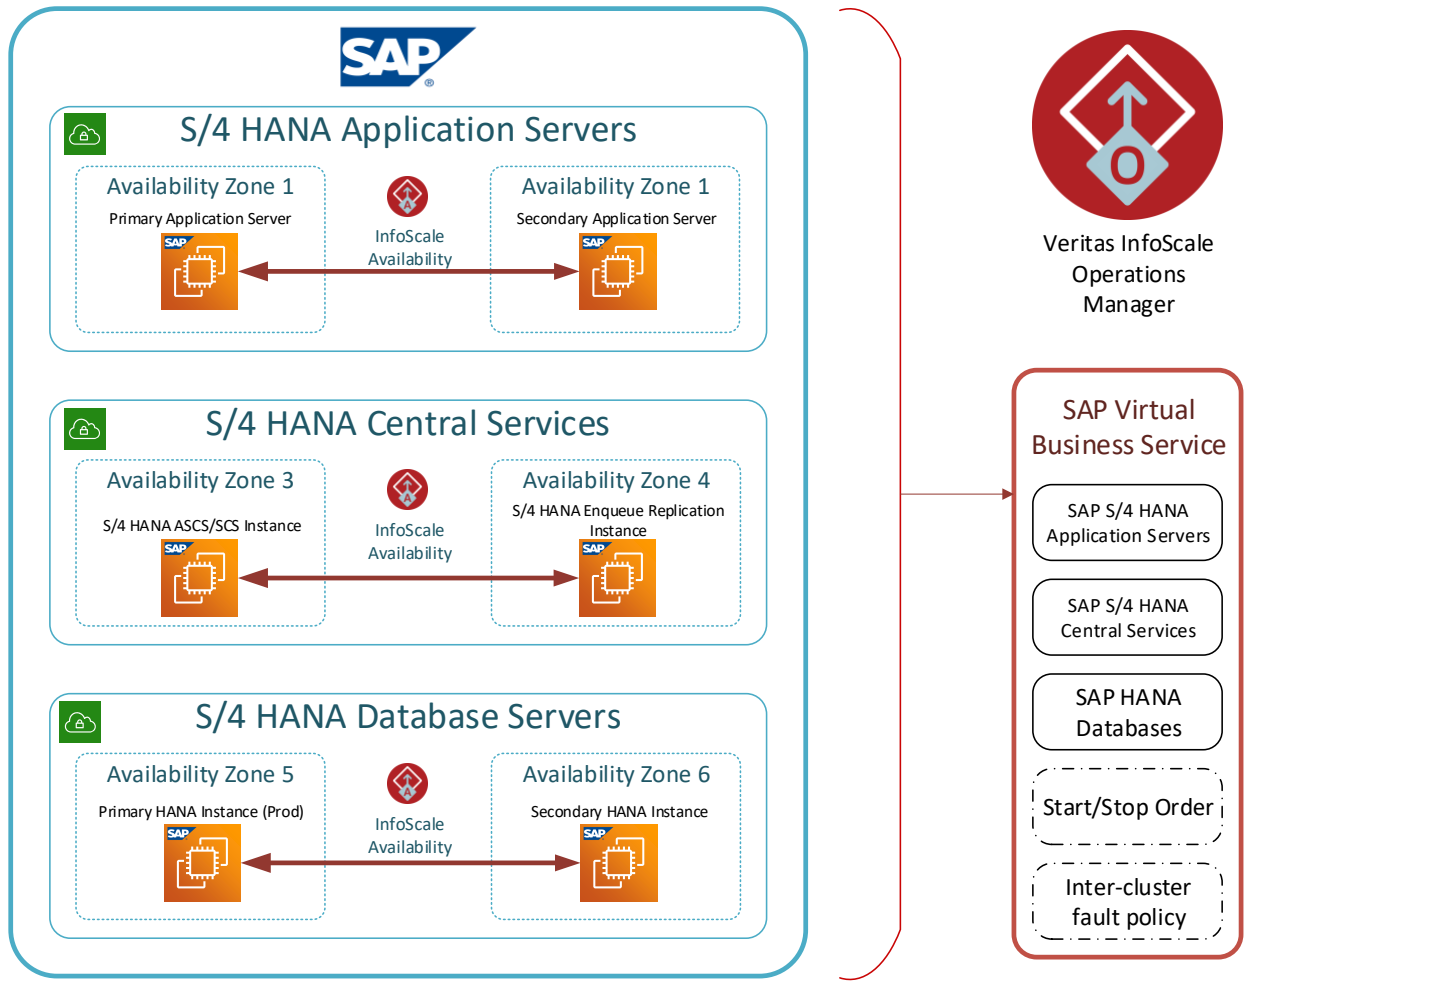
The following HA and DR capabilities of InfoScale Enterprise make it ideal for managing an SAP ecosystem in AWS:

- Business continuity with minimal application downtime during failures through complete automation of the SAP landscape
- Support for on-premises to AWS 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 AWS
- Support for DR, which includes:
 - On-premises to AWS
 - Supported on-premises SAP agents continue to be relevant in AWS by providing:
 - Operational parity between on-premises and the cloud
 - Flexible Storage Sharing (FSS) for data sharing in AWS across instances & AZ's
 - Replication across AWS 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 the SAP Business Suite applications have different components of the application running on different tiers of infrastructure, each with their own unique availability requirements. A failure in 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 the appropriate action in the event of a failure to restore the entire service. When an individual component fails, the Virtual Business Service provides automated orchestration of the connections to other computing resources, on-site or even across sites. This results in a faster service recovery and minimal downtime—with no manual intervention.

Sample VBS Architecture



Flexible Storage Sharing feature of InfoScale Enterprise

The FSS feature of InfoScale Enterprise allows you to combine distributed, high-performance and highly available file systems with the latest storage and networking technologies. FSS unlocks the 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.

Note: Veritas InfoScale FSS supports all AWS 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, refer to 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 monitor 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 fault within an AZ
- 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 AWS region (local failover)
- Replication between AZs across regions (instances in different regions)
- Replication across AZs within one AWS 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 S4/HANA / 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 will shut down the instance in the case of a failover.

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

- Fast Failover of faulted Instances within an AZ
- 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 AWS region (local failover)
- Disaster Recovery between AZs across regions (instances in different regions)
- Availability Across AZs within one AWS 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 to acquire the replicated enqueue table which resides in the shared memory of the active ERS node.

The new Standalone Enqueue Server 2 and Enqueue Replicator 2 provides 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 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/Sl/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 and this can cause application crashes in high availability systems.

Veritas InfoScale SAPNW/SAPHDB agents can manage restarts of the SAP applications gracefully during system reboots

InfoScale agent for AWS IP

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

- Private IP: A private IP is a private numerical address that networked devices use to communicate with one another
- Elastic IP: An elastic IP address is a static IPv4 address designed for dynamic cloud computing, and it is associated with your AWS account
- Overlay IP: AWS allows you to redirect IP address traffic to an EC2 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

InfoScale agent for AWS Route 53

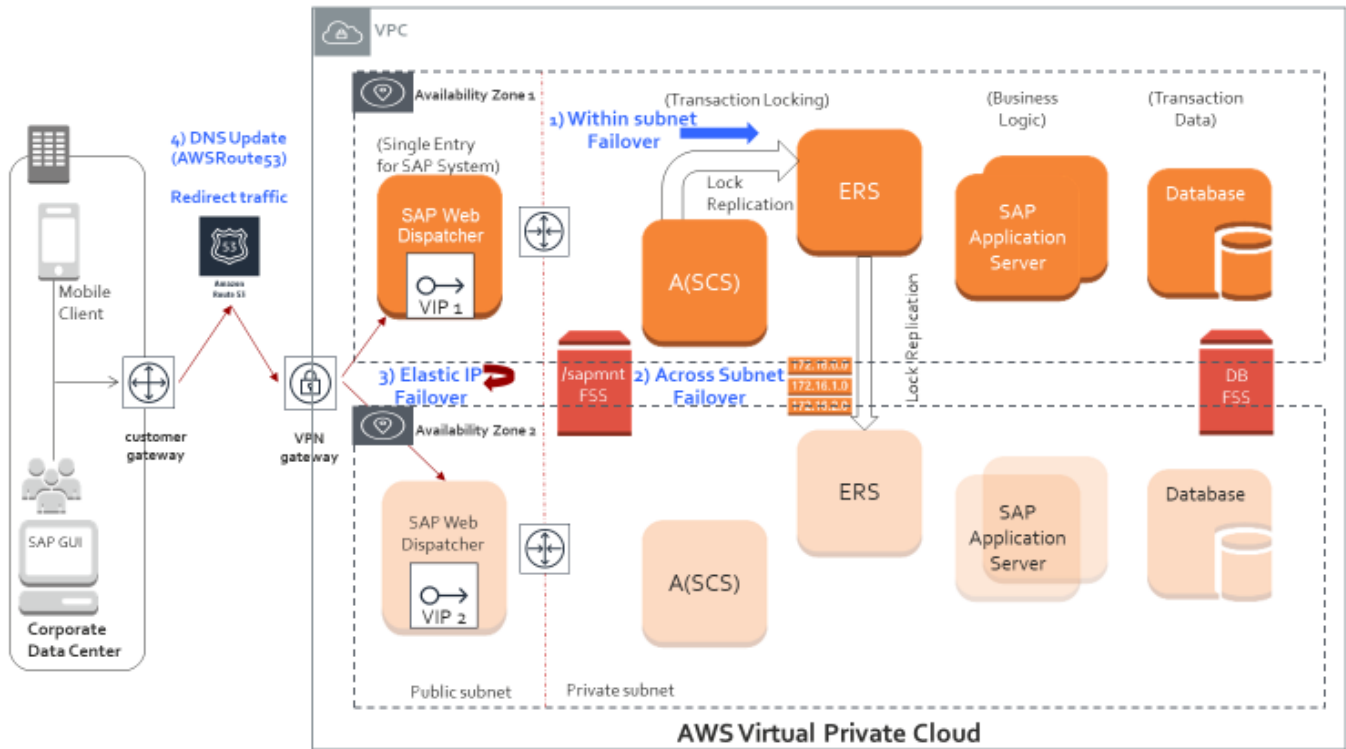
Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service. InfoScale Enterprise provides the AWSRoute53 agent to update and monitor the mapping between host names and IP addresses. The agent manages the mapping for the Amazon Route 53 domain when failing over nodes across subnets. When you create a hosted zone, Amazon Route 53 automatically creates a name server (NS) record and a Start of Authority (SOA) record for the zone.

If the resource records need to be dynamically added and deleted from the Amazon Route 53 domain during failover, you must use the AWSRoute53 agent. The agent updates the NS with the new resource record mappings during failover and allows the clients to connect to the failed-over instance of the application.

If you do not want to use the AWSRoute53 agent, you can continue to use the legacy InfoScale DNS agent for managing DNS records.

TYPICAL S/4HANA / NETWEAVER DEPLOYMENT ARCHITECTURE ON AWS

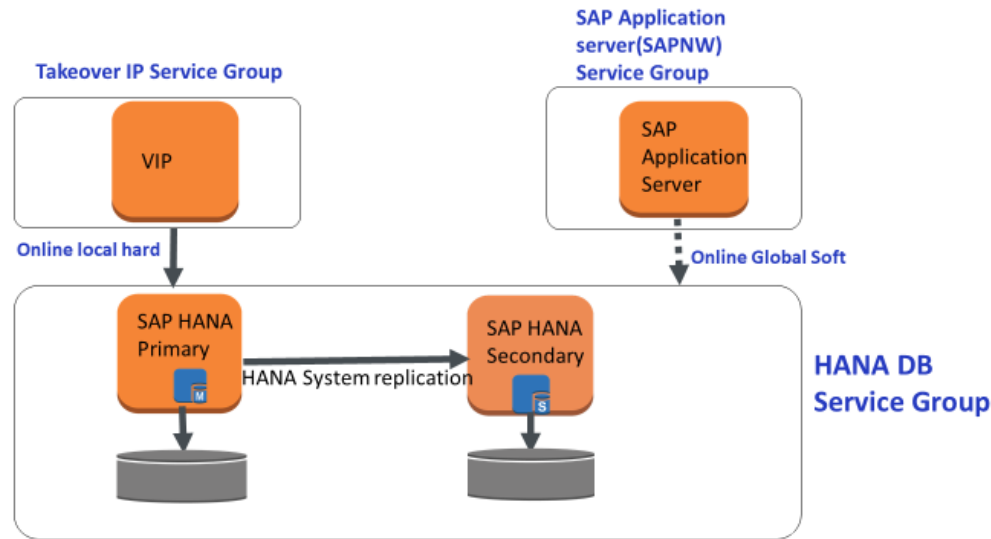
The following graphic describes the overall deployment of an SAP S/4HANA application server with an SAP HANA database on AWS:



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

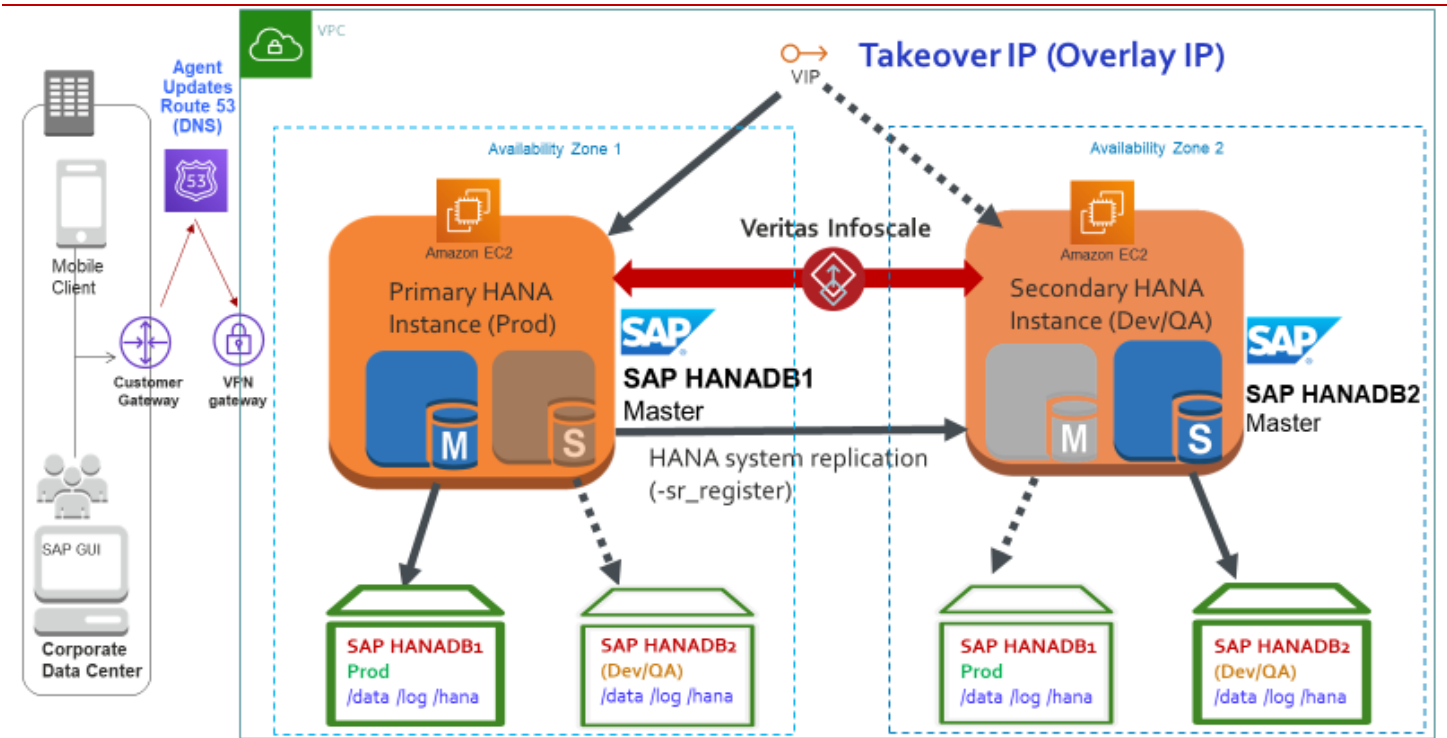
COMMON DEPENDENCY BETWEEN SAP S/4HANA AND SAP HANA DATABASE

The following graphic depicts the typical dependency between an SAP application and an SAP HANA database:



COST OPTIMIZATION OF SAP INSTANCES ON AWS

In AWS, you can 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. This helps reduce the Total Cost of Ownership (TCO) in case of a disruption or outage of SAP instances in AWS.



SUPPORTED USE CASES FOR SAP HANA ON AWS

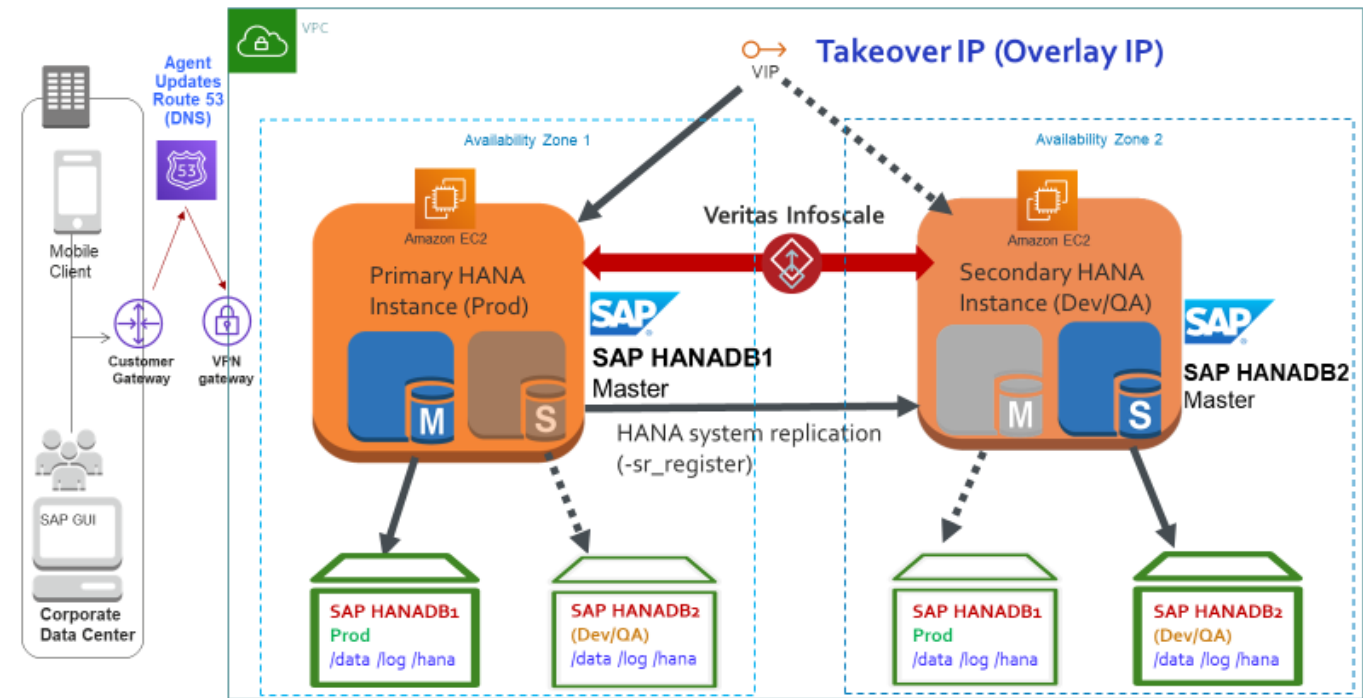
The following graphics represent the availability configuration scenarios supported by InfoScale for SAP HANA. If an SAP HDB instance in a production environment fails, InfoScale detects the failure and moves that instance to either a development or a testing node.

InfoScale Enterprise lets you monitor and control SAP HANA database instances in the following use cases.

SAP HANA instances in the same AZ

In this scenario, a master instance (primary for replication) and a standby instance (secondary for replication) 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.

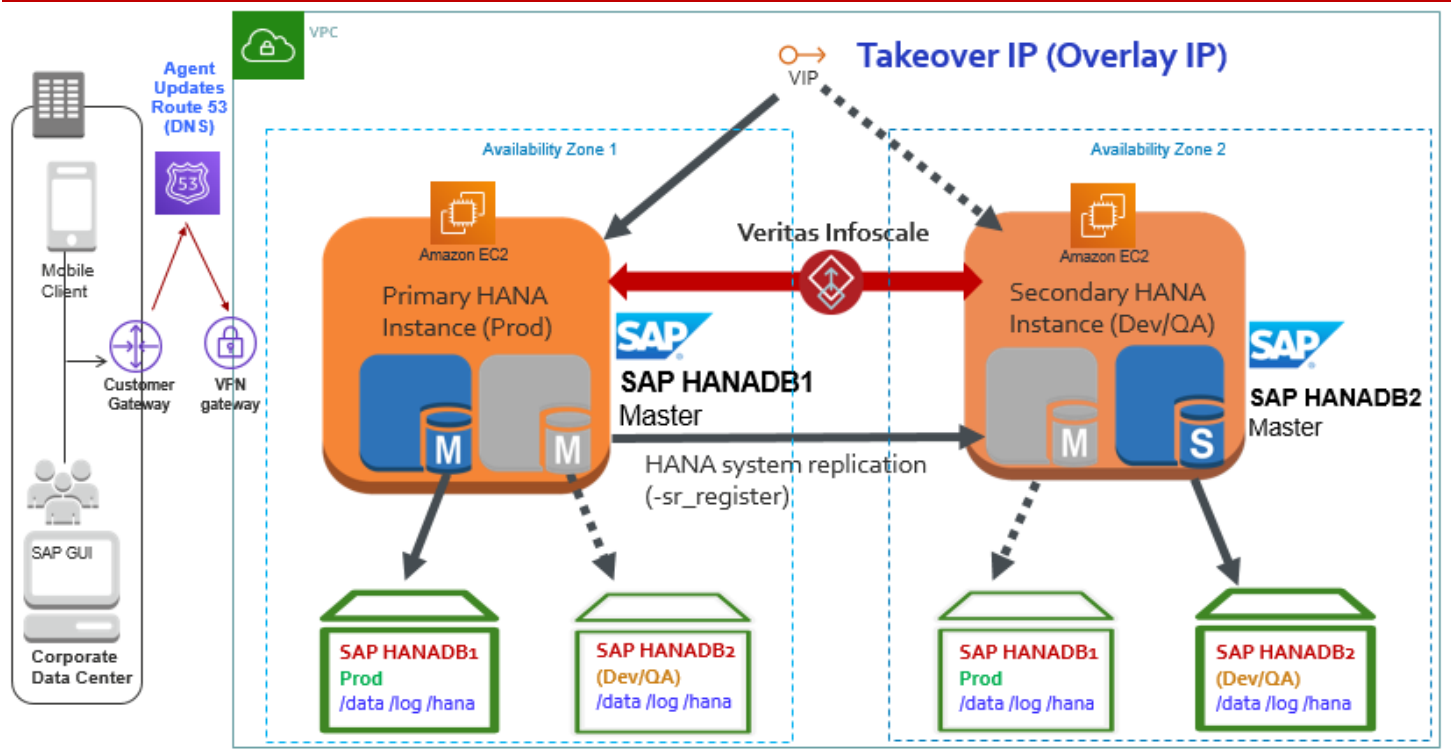


Note: The failover operations are performed according to the HANA HA takeover guidelines and the SAPHDB agent adheres to the HANA Takeover rules.

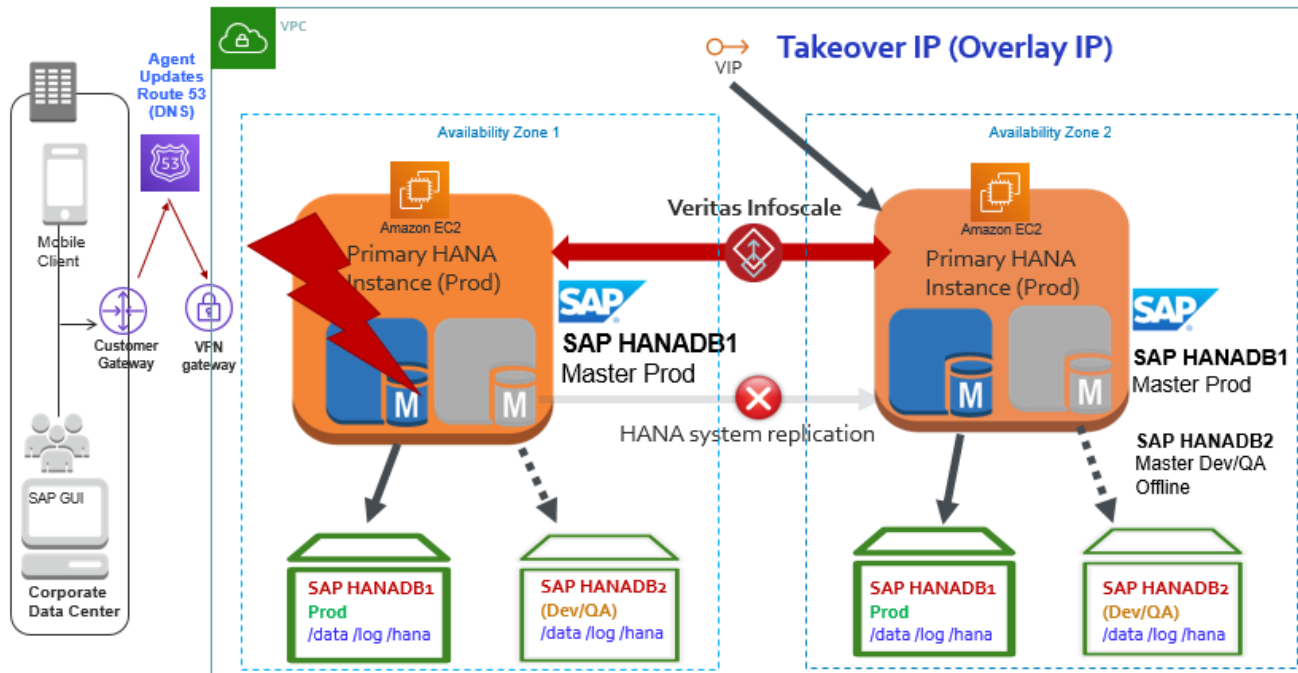
SAP HANA instances across AZs in the same AWS region

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, the data and logs are replicated to the secondary instance.

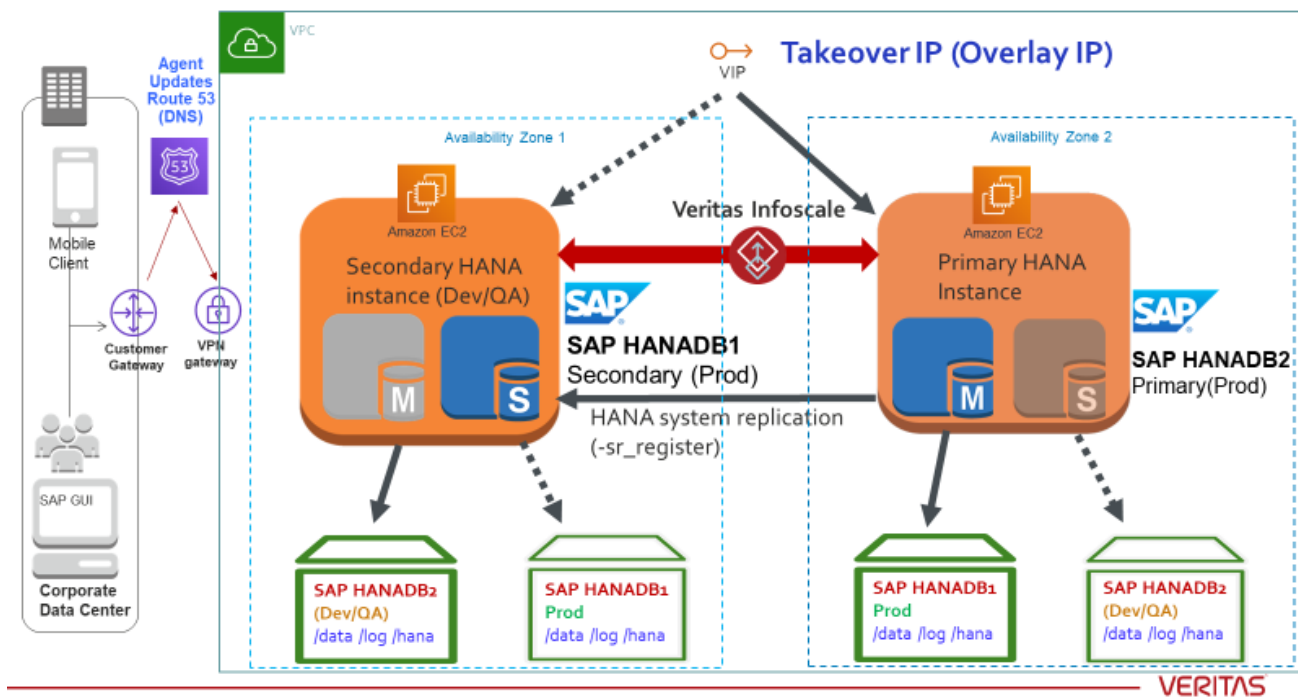
In this example, the primary (master) and the secondary (worker) instances are configured in different AZs with local clustering.



If the primary instance fails or becomes unavailable, the SAPHDB agent identifies the fault and automatically triggers the takeover operation on the secondary instance. This is shown in the following graphic:

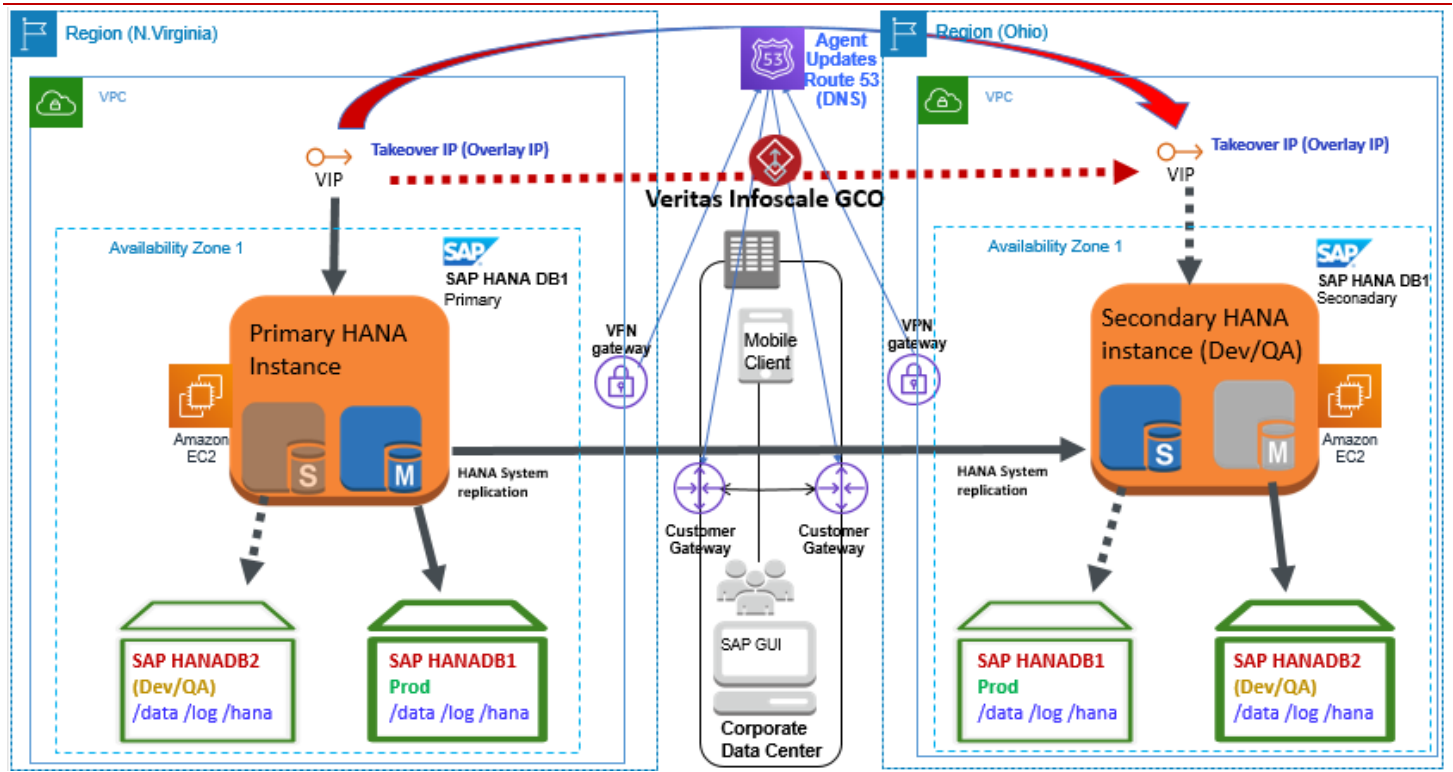


You may choose to clear the fault and perform any other necessary maintenance activities on the old primary instance. The SAPHDB agent can then automatically designate the original primary instance as the secondary instance by using the auto re-register feature. HANA System Replication will then continue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

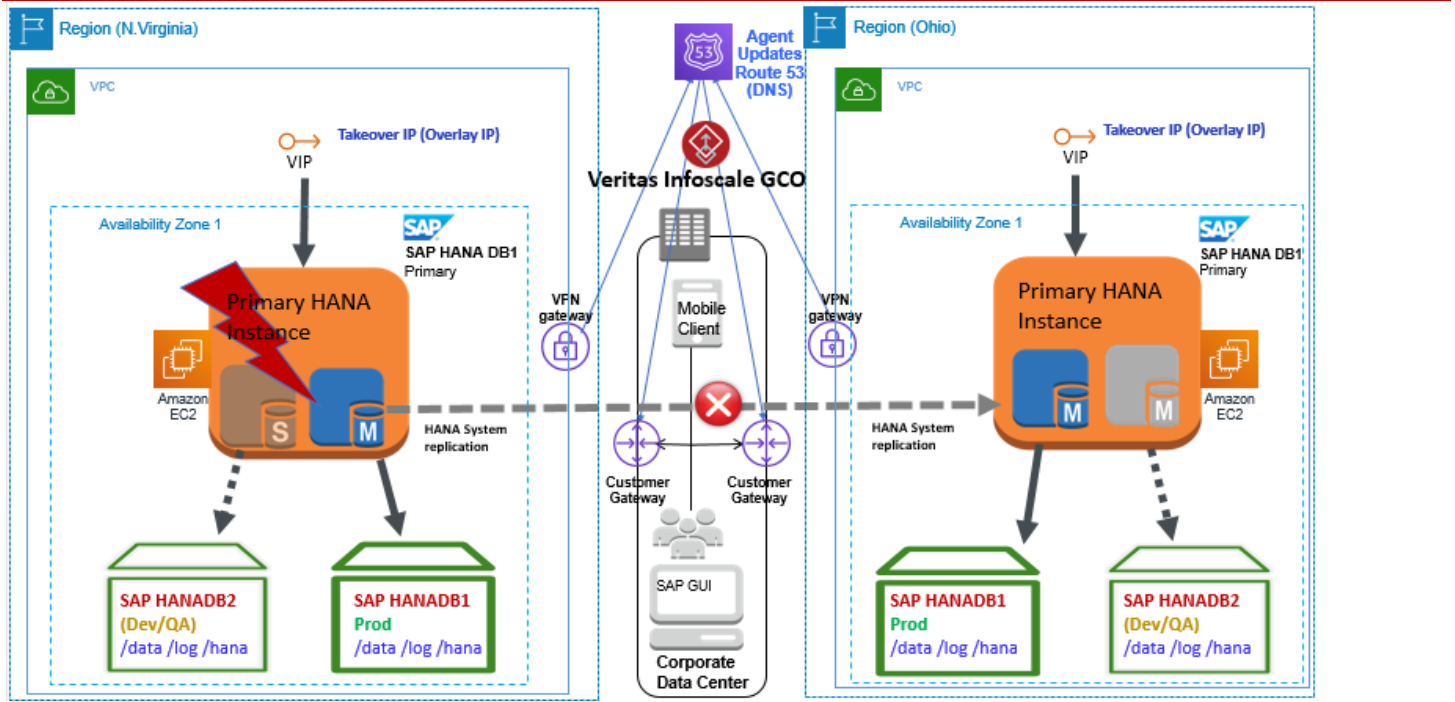


Single-instance SAP HANA databases across AWS regions

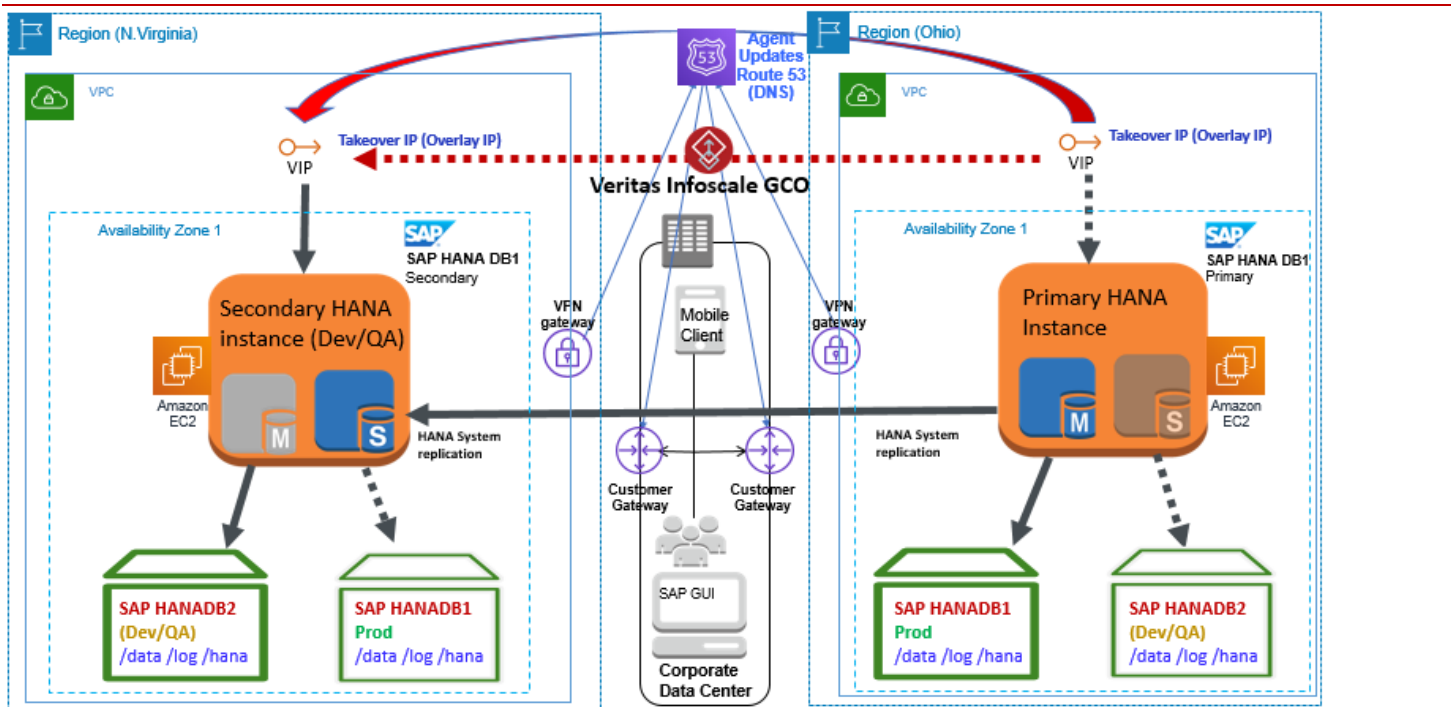
You can configure single HANA database instances across AWS regions, and they can be controlled and monitored by InfoScale agents using a GCO configuration. The is shown in the following graphic:



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. This is shown in the following graphic:



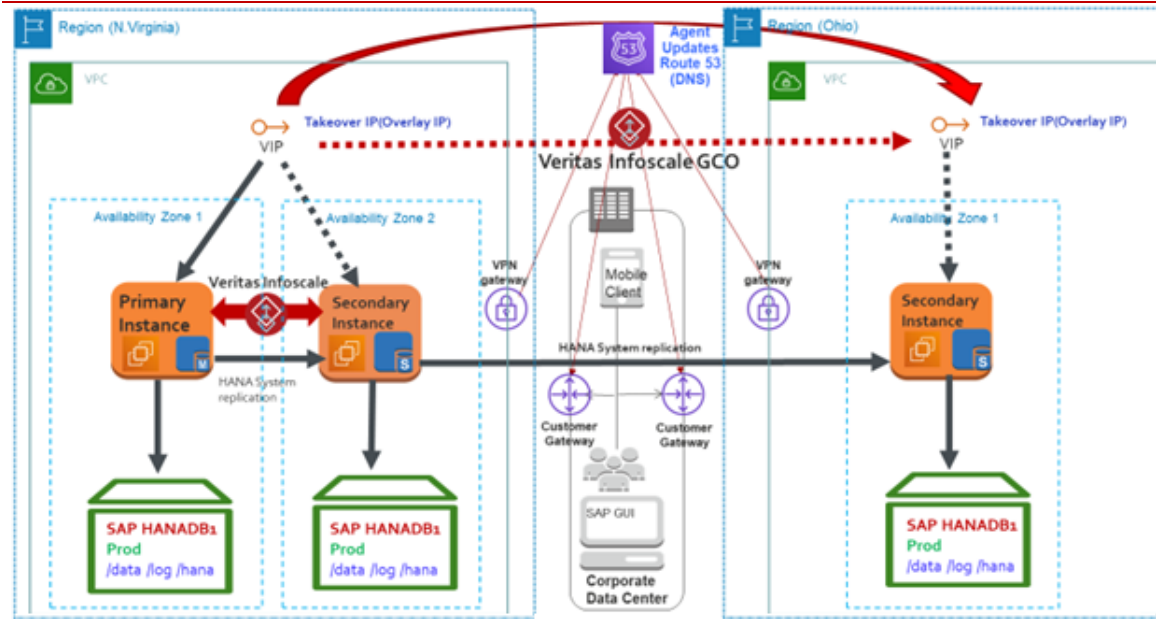
You'll need to clear the fault and perform any other necessary maintenance activities on the old primary instance. The SAPHDB agent can then automatically designate the original primary instance as the secondary instance by using the auto re-register feature. This is shown in the following graphic:



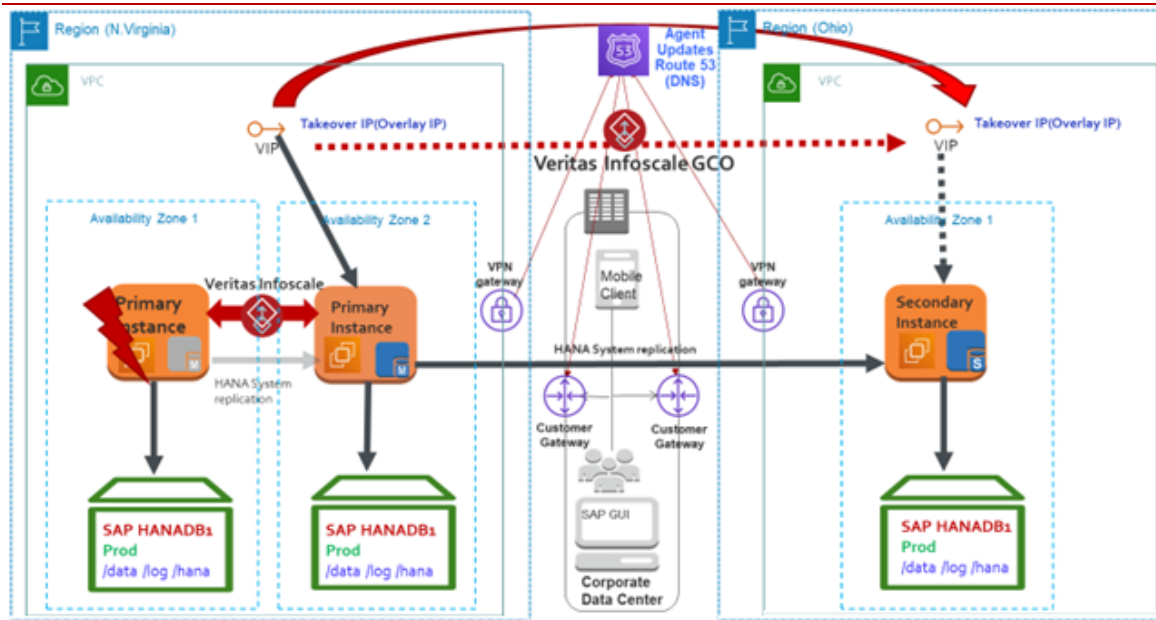
SAP HANA database instances across AWS regions (cascading scenario)

In this scenario, the HANA database primary and secondary are configured in AvailabilityZone1 and AvailabilityZone2 respectively in the N.Virginia region. The third HANA database instance is in the Ohio region.

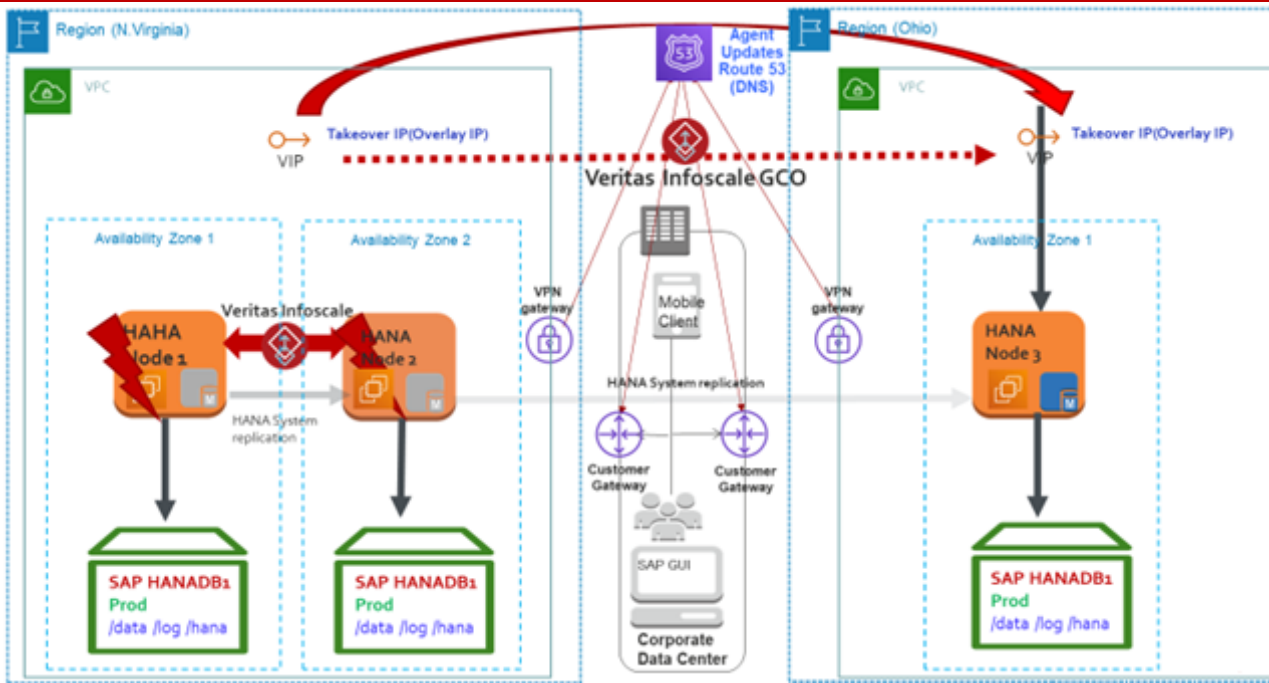
The is shown in the following graphic:



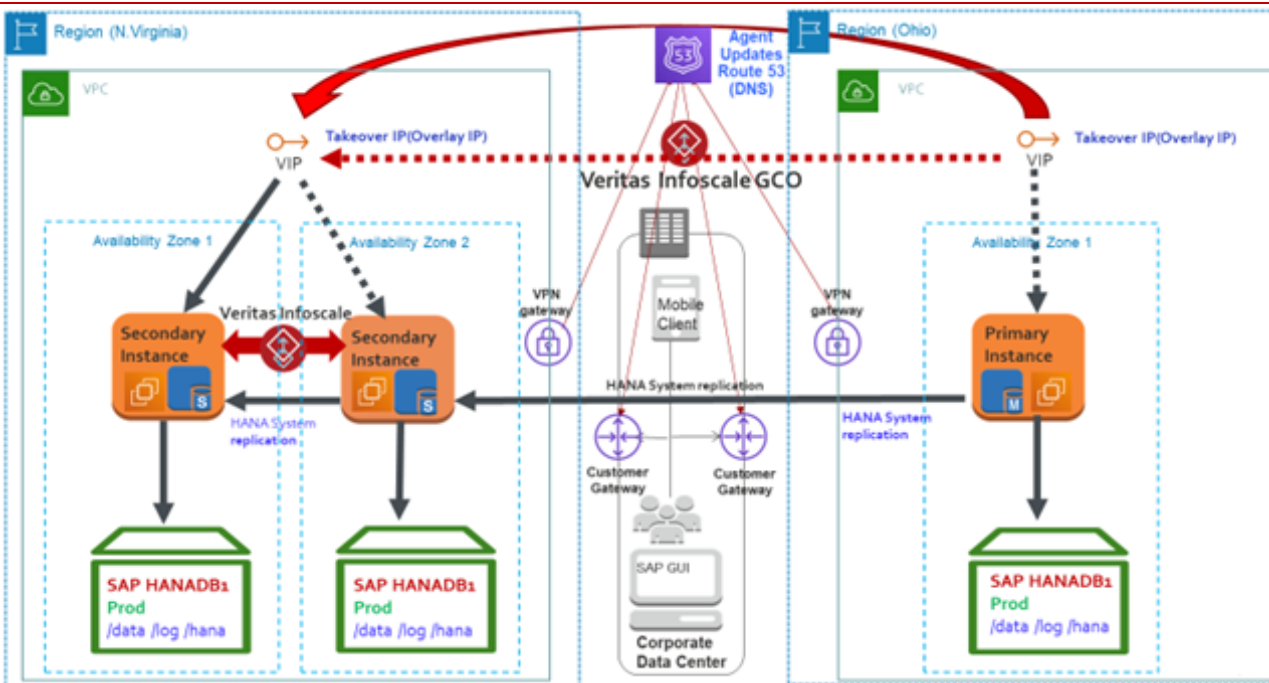
If the primary instance fails, the SAPHDB agent automatically triggers the takeover by failing over the IP resource from the primary to the secondary instance in the same region. Virtual IP failover operations are managed by the AWSIP and the IP agents. The following graphic shows how the secondary becomes the primary and is active in AvailabilityZone2:



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. This is shown in the following graphic where the secondary in the remote region becomes the primary:



Thereafter, you'll need to clear the faults on the instances in the N. Virginia region and re-register them with the current primary in Ohio. The following graphic shows the direction of the replication being reversed:



For details, refer to the InfoScale agent installation and configuration documentation at:

https://sort.veritas.com/agents/getting_started

Configuration Procedure

1. Configure a VPN to connect the corporate data center and the AWS virtual private cloud.
2. Create instances based on the planning document and the sizing of Instances recommended by AWS 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 Amazon EC2 instances.
 - b) Attach the required SSD or standard Amazon EBS volumes.
 - c) Mount the following SAP mount points using the 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 AWS volume types. VxVM, VxFS, and FSS 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

For specific instructions, see the Storage Foundation Cluster File System High Availability Administrator's Guide. Veritas used the gp2 (SSD) and the Standard (magnetic disks) EBS volume types for testing the SAP NetWeaver configurations in AWS.

7. Install and configure SAP HANA on the AWS instances on primary and secondary sites as per your disaster recovery plan.
8. Configure SAP HANA system replication between the sites on AWS. Ensure that all the required ports in AWS are enabled for replication.
9. Configure InfoScale cluster service groups and resources for the SAP HANA instances, the AWSIP agent, and the AWSRoute53 agent.

SUMMARY

With SAP HANA becoming 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 AWS. InfoScale has direct integration with both SAP HANA and AWS 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 for SAP HANA in AWS 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 AWS regions
- ✓ Flexible configuration options that support multiple usage scenarios within AWS AZs 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 AWS that provides the tools needed to run SAP applications with maximum uptime.

REFERENCES

SAP notes for SAP on AWS:

- [1588667](#): Overview of related SAP Notes and Web-Links
- [1656099](#): Supported SAP, DB/OS and AWS EC2 products
- [1656250](#): Support prerequisites
- [171356](#): SAP Software on Linux: Essential information

Amazon documentation:

- [SAP on AWS](#)
- [SAP on AWS High Availability Guide](#)
- [SAP on AWS Technical Documentation](#)

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 50,000 enterprises—including 99 of the Fortune 100—rely on us to abstract IT complexity and simplify data management. 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 supports more than 500 data sources and over 150 storage targets, including 60 clouds. Learn more at www.veritas.com. Follow us on Twitter at [@veritastechllc](https://twitter.com/veritastechllc).

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

For specific country offices
and contact numbers,
please visit our website.

VERITAS™