High Availability and Disaster Recovery Options for SAP HANA on AWS

November 2015

Notices

Contents

Abstract

Amazon Web Services (AWS) provides a suite of services and tools to deploy business-critical SAP HANA workloads on the AWS cloud infrastructure. This whitepaper discusses options that a customer or partner can choose to set up a highly available and disaster-tolerant SAP HANA system on AWS platform using AWS services and SAP HANA tool sets.

Introduction

Many enterprises are deploying SAP HANA systems on AWS to run their mission-critical business operations. High availability (HA) and disaster recovery (DR) for SAP HANA solutions are important considerations for business continuation. AWS provides many services that can be set up in a redundant way to achieve high availability for many applications, including SAP. In addition to these, SAP provides robust native tools to set up and run SAP HANA workloads in a highly available manner.

This whitepaper provides an overview of how a customer or partner could use AWS services, native SAP HANA HA tools, and third-party software to achieve HA/DR for SAP HANA systems on the AWS platform. The paper presents six high availability options that leverage different AWS management features, Availability Zones, and global regions. It discusses the pros, cons, and related costs of each option. We recommend that you select the option that best meets

your needs, and use the technical implementation guides we've referenced to implement its architecture.

Scope of the Document

The scope of this whitepaper is to provide you with different implementation options for setting up your SAP HANA system in a highly available and disaster-tolerant way on AWS platform while considering your own costs, recovery time objective (RTO), and recovery point objective (RPO). This whitepaper does not provide detailed implementation steps or discuss high availability or disaster recovery options for the SAP application tier. For that discussion, see the SAP on AWS High Availability Guide.[1]

Overview of AWS Services for SAP Solutions

AWS provides an extensive set of computing resources and services. The following sections provide a brief overview of the core AWS services that are most relevant for the implementation and operation of SAP solutions.

Amazon Global Infrastructure

AWS is built on a global infrastructure. AWS services are available to use in different geographical regions across the globe. AWS regions enable you to deploy your SAP environment in a location that is closer to your users, and to meet legal or other requirements. Regions are isolated from one another, and AWS does not replicate your resources or data across regions automatically.

Each region contains multiple isolated locations called *Availability Zones*. Each Availability Zone is engineered to be isolated from failures in other Availability Zones, and provides inexpensive, low-latency network connectivity to other Availability Zones in the same region. AWS is currently available in over 10 regions and in multiple Availability Zones within each region. For additional information about AWS regions and Availability Zones, see Global Infrastructure on the AWS website.[2]

Compute

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable compute capacity in the cloud. Its simple web service interface allows you to obtain and configure capacity with minimal effort. Amazon EC2 provides you with complete control over your computing resources and lets you run on Amazon's proven computing environment. Amazon EC2 helps reduce the time required to obtain and boot new server instances to minutes, allowing you to quickly scale capacity both up and down, as your computing requirements change. Amazon EC2 changes the economics of computing by allowing you to pay only for capacity that you actually use. It provides you with the tools to help build failure-resilient applications that are isolated from common failure scenarios. For more information about this service, see Amazon EC2 on the AWS website.[3]

Storage

Amazon Elastic Block Store (Amazon EBS) provides persistent block-level storage volumes for use with Amazon EC2 instances in the AWS cloud. Each Amazon EBS volume is automatically replicated within its Availability Zone to help protect you from component failure, offering high availability and durability. Amazon EBS volumes provide the consistent and low-latency performance needed to run your workloads. With Amazon EBS, you can scale your usage up or down within minutes – all while paying a low price for only what you provision. For SAP workloads, Amazon EBS is used to store database-related files like data and log files. See Amazon EBS on the AWS website to find out more about this service.[4]

Amazon Simple Storage Service (Amazon S3) provides developers and IT teams with secure, durable, highly scalable object storage. Amazon S3 is easy to use, and provides a simple web service interface to store and retrieve any amount of data from anywhere on the web. With Amazon S3, you pay only for the storage you actually use. Amazon S3 provides cost-effective object storage for a wide variety of use cases, including cloud applications, content distribution, backup and archiving, disaster recovery, and big data analytics. Additionally, you can take advantage of the S3 cross-region replication feature, which enables automatic, asynchronous copying of objects across buckets in different AWS regions. For SAP workloads, Amazon S3 can be used to store database backups, SAP archive

files, and other miscellaneous files like SAP media. See Amazon S3 on the AWS website to find out more about this service.5

Automatic Recovery

Automatic recovery is a feature of Amazon EC2 that is designed to increase instance availability. You can enable automatic recovery for an instance by creating an Amazon CloudWatch alarm that monitors an Amazon EC2 instance and automatically recovers the instance if it becomes impaired due to an underlying hardware failure or a problem that requires AWS involvement to repair. A recovered instance is identical to the original instance, including the instance ID, private IP addresses, Elastic IP addresses, and all instance metadata. For more information about how to configure an Amazon CloudWatch alarm to enable automatic recovery, see the AWS documentation.6 The following diagram illustrates this feature.
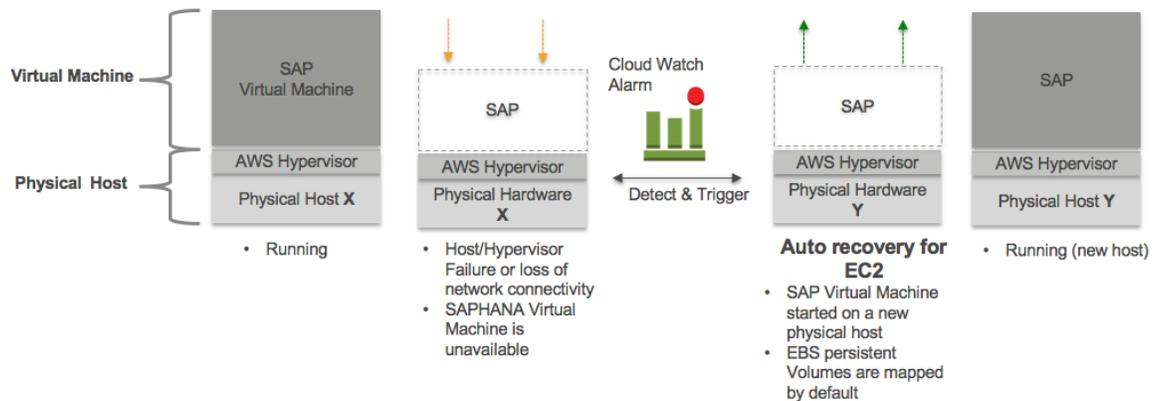
Quick Starts

The AWS Quick Start reference deployment for SAP HANA helps you rapidly deploy fully functional SAP HANA systems on the AWS cloud, following best practices from AWS and SAP. The AWS Quick Start ensures that the Amazon EC2 instance, Amazon EBS, and the operating system (SUSE Linux Enterprise Server or Red Hat Enterprise Linux) are optimally configured to achieve best performance for your SAP HANA system right out of the box. For more information, see the SAP HANA on AWS Quick Start Reference Deployment Guide.7

SAP HANA High Availability on AWS

SAP provides several native tool sets to recover your SAP HANA system from a wide range of failures, starting from simple software malfunctions to major disasters like complete data center outage. This section provides an overview of various high availability (HA) tool sets that SAP provides for SAP HANA and their use on the AWS platform. For details about these tool sets, see the SAP HANA High Availability whitepaper on the SAP website.[8]

HANA Service Auto-Restart

HANA Service Auto-Restart is a fault recovery solution provided out of the box by SAP. 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 service is automatically restarted with the HANA Service Auto-Restart watchdog function. When the service is restarted, it loads all the necessary data back into memory and resumes its operation.

Use on AWS – HANA Service Auto-Restart works the same way on AWS as it does on any other platform.

Host Auto-Failover

SAP HANA Host Auto-Failover is a fault recovery solution provided by SAP. To achieve high availability within a data center, you can add and configure standby nodes to take over in case one or more SAP HANA nodes become impaired or unavailable. When an existing node fails, a standby node can automatically take over by gaining access to the persistent layer (data and log volumes) of the failed host. It does this either by using shared networked storage (NFS) or with any storage connector API.

Use on AWS – Instead of having to purchase and maintain a standby node for your SAP HANA deployment, we recommend that you use the native AWS automatic recovery feature previously discussed to recover your SAP HANA instance on a different piece of hardware if an underlying failure occurs. This essentially provides you with standby hardware without incurring the additional cost for a typical standby node. A recovered instance is identical to the original instance, including your existing storage volumes as well as configurations such as hostname, IP address, and AWS instance IDs. In addition, we recommend that

you configure SAP HANA services to start up automatically after a system restart for a fully automated recovery.

HANA System Replication

HANA System Replication (HSR) is a high availability and disaster recovery solution provided by SAP. HSR follows the "N+N" approach, where you configure the same number of nodes for both the primary and the secondary system. Each service in the primary system communicates with its counterpart in the secondary system to replicate the data. With HSR, you can choose many replication options, including synchronous, synchronous in-memory, and asynchronous, depending on your recovery time objective (RTO) and recovery point objective (RPO). To learn more, see [How to Perform System Replication for SAP HANA](#) on the SAP website.[9]

Use on AWS – SAP HANA System Replication is fully supported on the AWS platform. You can use it in combination with AWS Availability Zones to help protect your SAP HANA installations, even from catastrophic data center failures.

HANA Backup/Restore

Although SAP HANA is an in-memory database, it persists all changes in persistent storage to recover and resume from power outages without any loss of data. Persisted SAP HANA data and log files can be regularly backed up to a remote location for disaster recovery purposes. To learn more, see [Backup and Recovery - SAP HANA](#) on the SAP website.[10]

Use on AWS – There are no AWS-specific restrictions for using the HANA Backup/Restore functionality. It works the same way on AWS as it does on any other platform. In addition, you can take advantage of secure, durable, highly scalable and cost-effective Amazon S3 object storage, either by copying your backup files to Amazon S3 or by using Amazon EBS snapshots to help recover your HANA system from a disaster.

Storage Replication

SAP HANA hardware partners offer a storage-level replication solution, which replicates the SAP HANA storage volumes (data, logs, etc.) or file system on a remote, networked storage system to recover the SAP HANA system with low RTO during a disaster.

Use on AWS – Storage replication is not currently supported on AWS.

SAP HANA HA/DR Options on AWS

Combining the AWS global infrastructure, Amazon S3, Amazon CloudWatch, and automatic recovery services with native SAP HANA tool sets like HANA System Replication and Backup/Restore provides you with many options for architecting a highly available and disaster-tolerant HANA system on AWS to meet your specific RTO/RPO requirements. In this section, we will discuss some common HA/DR options for SAP HANA on AWS. For a comparison of the RPO, RTO, and costs associated with each option, see the summary.

Option 1 – Automatic Recovery & HANA Backup/Restore
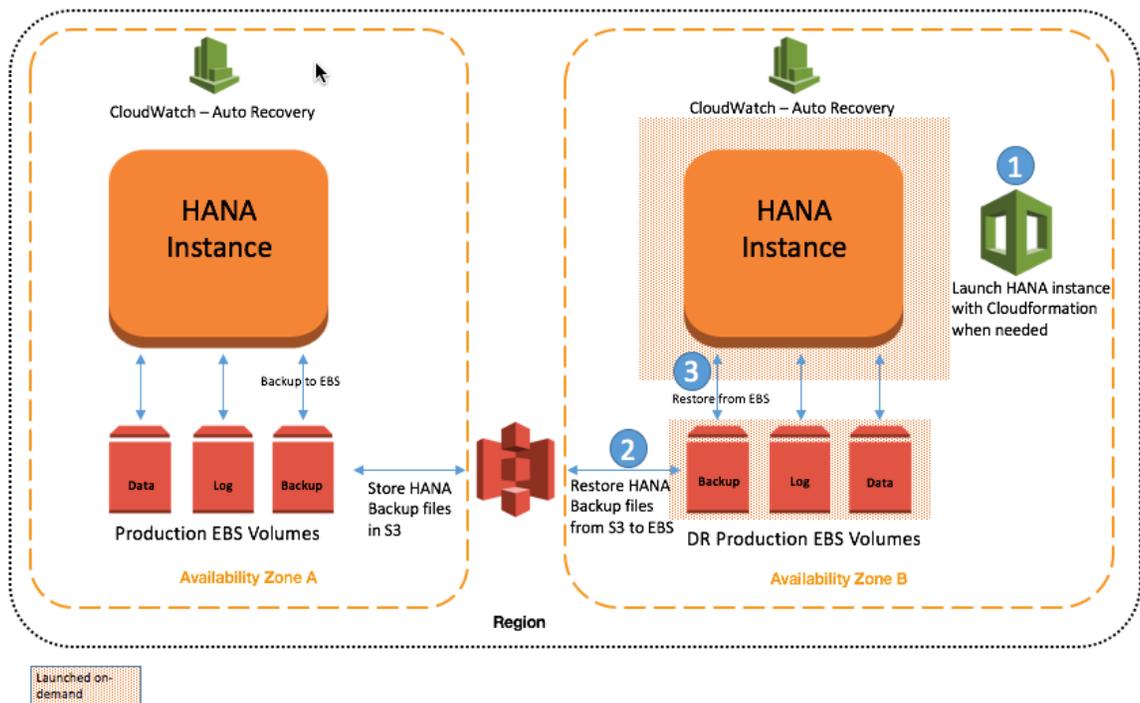
In this option, you can choose to run your HANA instance in Availability Zone A and set up a CloudWatch alarm to automatically recover your SAP HANA EC2 instance when a hardware impairment is detected. Automatic recovery helps recover your instance during a hardware failure within an Availability Zone, but you still have the risk of not being able to access your SAP HANA EC2 instance when the Availability Zone becomes unavailable. In this case, you can launch a new SAP HANA instance in another Availability Zone (for example, Availability

Zone B) by using an AWS CloudFormation template, and restore your SAP HANA backup from Amazon S3 into an Amazon EBS staging volume.

Another way to store your backup files in Amazon S3 is to take periodic snapshots of your backup volumes. Snapshots stored in Amazon S3 are available across Availability Zones within a region, and you can instantly create an EBS volume based on a snapshot to restore your SAP HANA instance immediately. Using snapshots helps you avoid manually downloading and staging backup files from Amazon S3 to Amazon EBS, thus decreasing your total RTO.

Once the SAP HANA backup is staged on an EBS volume, you can use standard SAP HANA backup/recovery procedures to restore and recover your SAP HANA database. When the database becomes operational, you can redirect your client traffic to your new SAP HANA EC2 instance.

With this model, you avoid the cost of a standby node. Instead, you take advantage of the AWS Multi-AZ infrastructure and Amazon S3 to stand up a new SAP HANA instance in a different Availability Zone and restore your SAP HANA backup so you can quickly resume your operation. This model is suitable when you can afford a longer RTO, and your business can afford an RPO that is greater than zero. Your RPO depends on how frequently you store your SAP HANA backup files in Amazon S3. You could implement custom scripts to copy or snapshot your SAP HANA backup files to Amazon S3 as soon as they are written to the Amazon EBS volume, to achieve a better RPO in the event of an Availability Zone failure.

Finally, since HANA ensures that your committed transactions are persisted in persistent storage, you will not lose any data if you choose to wait for Availability Zone A to come back online, rather than recovering your instance in Availability Zone B. In this case, your recovery time depends on the time it takes to bring Availability Zone A back online.

Option 2 – Automatic Recovery & HSR without Data Preload (Warm Standby)
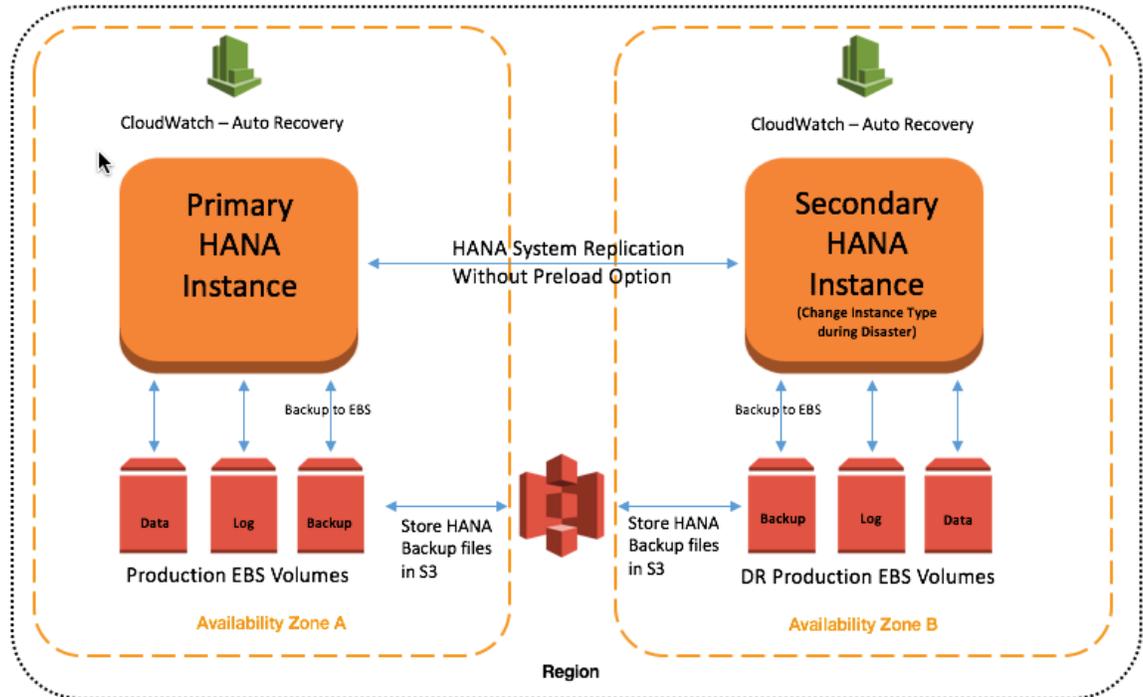
In this option, you host the primary SAP HANA instance in Availability Zone A and the secondary SAP HANA instance in Availability Zone B, and you set up HANA System Replication between these two instances. With HANA System Replication in place, data changes in your primary SAP HANA instance are continuously replicated to your secondary HANA instance, so you can immediately recover your HANA instance in Availability Zone B when Availability Zone A becomes unavailable. Both the primary and the secondary HANA instances can be protected from hardware impairment by configuring a CloudWatch alarm for automatic recovery.

With HANA System Replication, you can also set up the replication between your primary and secondary instances with the **preload** option turned off. When this option is turned off, replicated data is not loaded into memory in your secondary HANA instance, so your secondary instance doesn't need the same memory capacity as your primary instance. This helps lower your costs while still protecting your system against failure in Availability Zone A.

In this case, you can choose a smaller instance type with at least 64 GiB of memory, or row store size + 20 GiB (whichever is higher) for your secondary HANA instance in Availability Zone B. With HANA System Replication in place, you will be able to recover your HANA database within a short period of time in Availability Zone B. During a failover, you can change the secondary HANA instance type to match the primary HANA instance type. Once your secondary HANA instance is operational, you can start redirecting traffic to your secondary HANA instance.

This architecture enables you to implement your HANA solution across multiple Availability Zones to minimize your recovery time during an Availability Zone failure. You also take advantage of a unique flexibility in AWS that eliminates the need for a full-sized secondary HANA instance during replication. You need to change your secondary HANA instance type only when disaster strikes. While synchronous HANA System Replication ensures zero RPO, your RTO with this model will be moderate because you will need to change the instance type during failover. This model is suitable when you need to lower costs and still need a better RTO than the previous option.

When you use this option, you need to ensure that the secondary HANA instance is big enough and has adequate network bandwidth to support the primary HANA instance's data change rate, especially when you are using synchronous replication. We recommend that you choose the R3 instance family for your secondary HANA instance to achieve better performance during replication.

Option 3 – Automatic Recovery & HSR without Data Preload (Warm Standby + Dev/QA)
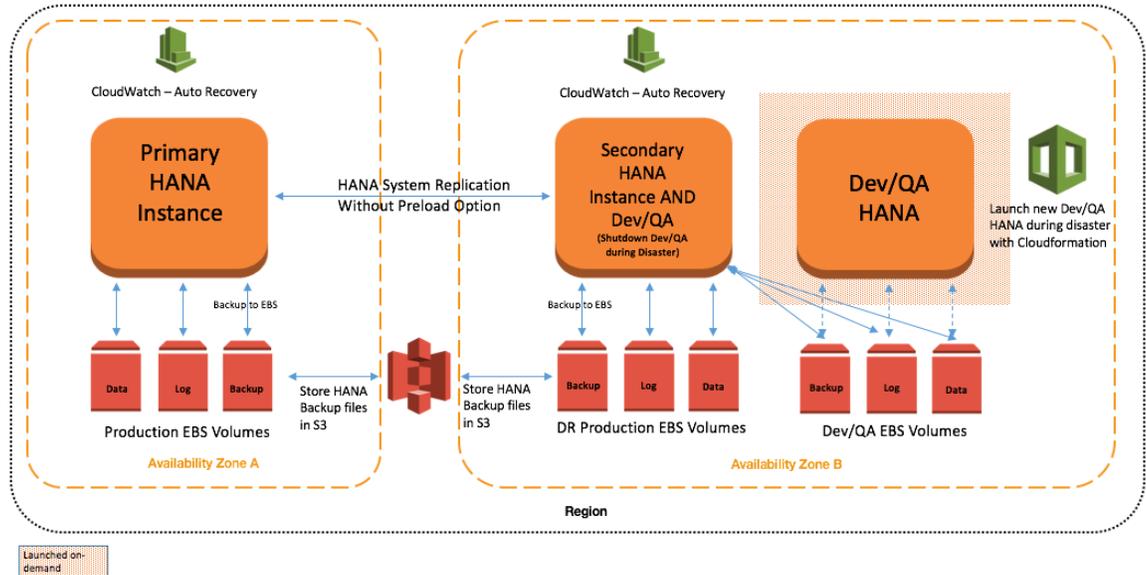
This option is very similar to option 2, except that your secondary HANA instance size is similar to the size of your primary HANA instance. In addition to hosting your production on the secondary HANA instance, you can co-host your Dev/QA HANA instances on the same EC2 instance to best utilize your resources and to lower your cost.

You need to ensure that you spare at least 64 GiB of memory, or row store size + 20 GiB (whichever is higher) for your replication workload while co-hosting your Dev/QA instances in your secondary HANA instance. You will need to shut down your Dev/QA systems on the secondary instance during a failover to free up resources for production. Once the failover is complete and your secondary instance is operational for production in Availability Zone B, you can start redirecting client traffic to your secondary HANA instance.

Although you have to shut down your HANA Dev/QA instances during the failover, you don't have to operate without Dev/QA after the failover. You can detach the Amazon EBS volumes for Dev/QA from the secondary HANA instance. You can launch a new HANA instance using AWS CloudFormation and reattach the Amazon EBS volumes for Dev/QA to the new instance to continue

your Dev/QA work. You can then redirect your Dev/QA client traffic to this newly launched instance.

This architecture provides you with the advantage of implementing your HANA solution across multiple Availability Zones to minimize your recovery time during an Availability Zone failure. It also enables you to share resources with HANA Dev/QA instances. You need to shut down your Dev/QA temporarily only when disaster strikes. While synchronous HANA System Replication ensures zero RPO, your RTO with this model will be moderate, because you will need to shut down the Dev/QA instances first during the failover. This model is suitable when you want to best utilize the resources of your secondary HANA instance and still need a better RTO than option 1.

Option 4 – Automatic Recovery & HSR with Data Preload (Hot Standby)
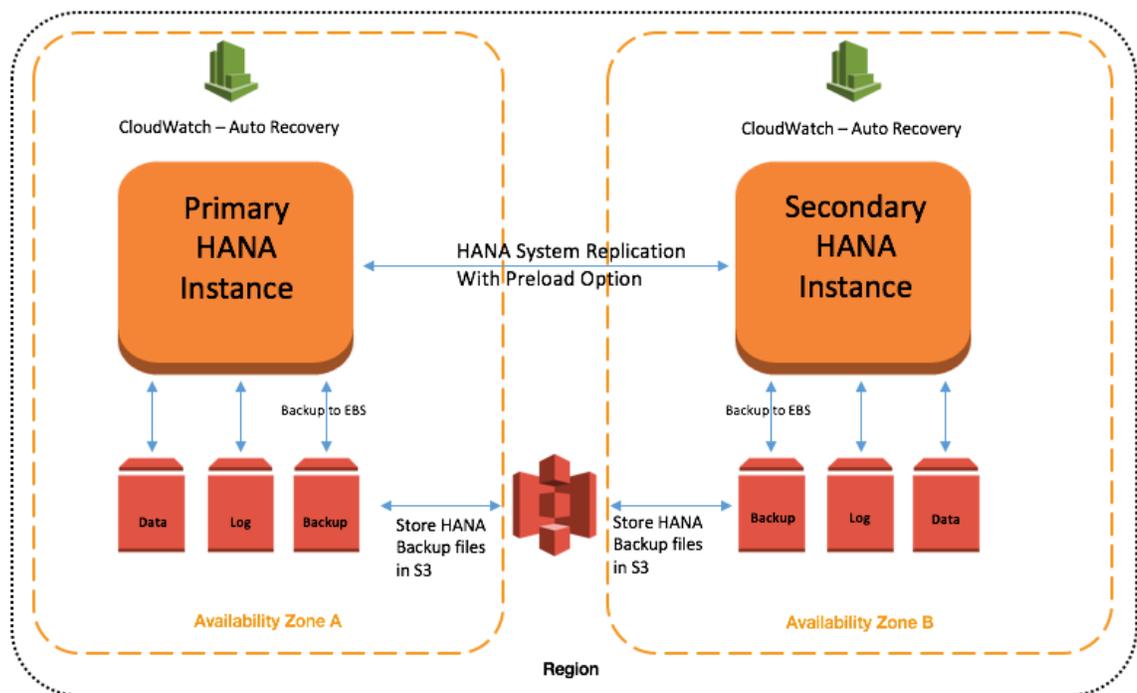

This option is identical to the previous two options, except that replication between your primary and secondary HANA instance is configured with the **preload** option turned on. When this option is turned on, replicated data is always loaded into the memory of your secondary HANA instance for instant

failover. Therefore, your secondary HANA instance must be sized like your primary HANA instance. You need to ensure that your client traffic is redirected to your secondary HANA instance after the failover.

This architecture provides you with the advantage of implementing your HANA solution across multiple Availability Zones with the ability to fail over instantly and continue your business operations during an Availability Zone failure. You get the best RTO and RPO with this option, but it will cost you the most. This model is suitable when you need the lowest possible RTO.

Option 5 – Automatic Recovery & Multi-Tier HSR (Hot Standby + Out-of-Region DR)
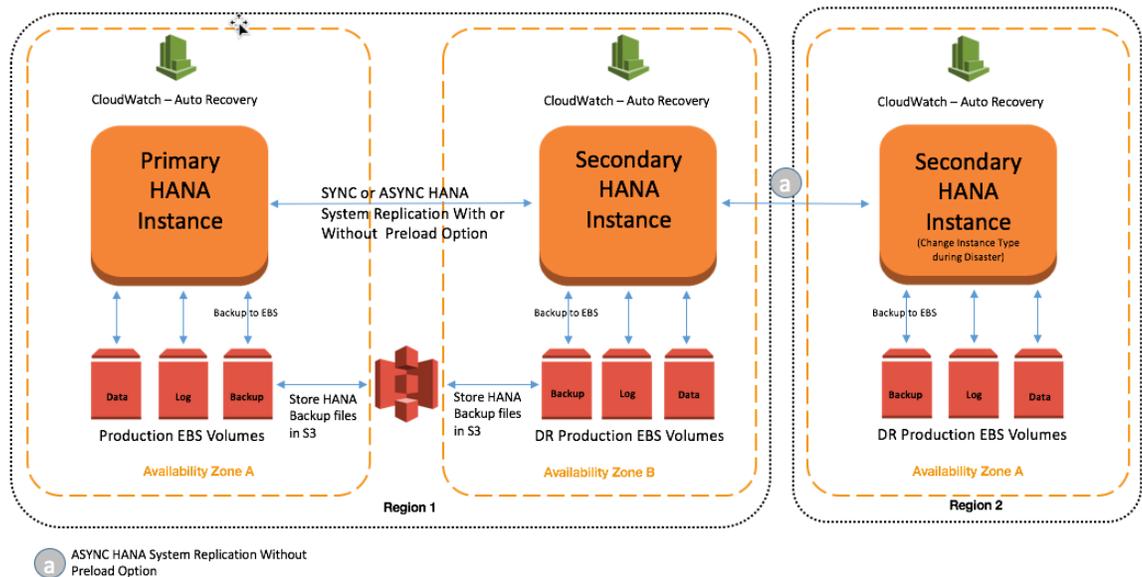
This option is identical to option 4, except that you add a third level of protection by using SAP HANA Multitier System Replication.

Your HANA instance becomes highly available when you deploy it in a Multi-AZ setup. However, if you are worried about region-level failure and need to set up disaster recovery for your HANA instance in another region, you can use asynchronous replication between your secondary HANA instance in Availability Zone B of region 1, and a third HANA instance in Availability Zone A in region 2. You could choose a smaller instance type for your third HANA instance, as you

did in option 2, and change the instance type to the desired size only when you fail over to the third instance.

This architecture provides you with the advantage of implementing your HANA solution across multiple Availability Zones with the ability to fail over instantly during an Availability Zone failure. It also provides protection when the whole region becomes unavailable. You get the best RTO and RPO, and an additional third-level, out-of-region protection with this option, but it will cost you the most. You need to be aware of costs associated with data transfers between regions as well. This model is suitable when you need out-of-region disaster recovery with the lowest possible RTO and RPO.

Note that your out-of-region disaster recovery RPO is constrained by the network latency between your secondary HANA instance in Availability Zone B of region 1, and your third HANA instance in Availability Zone A of region 2.

Option 6 – Automatic Recovery & HSR with Amazon S3 Cross-Region Replication (Hot Standby + Out-of-Region DR)

This option is similar to option 4, except that the third level of protection is provided by Amazon S3 cross-region replication.
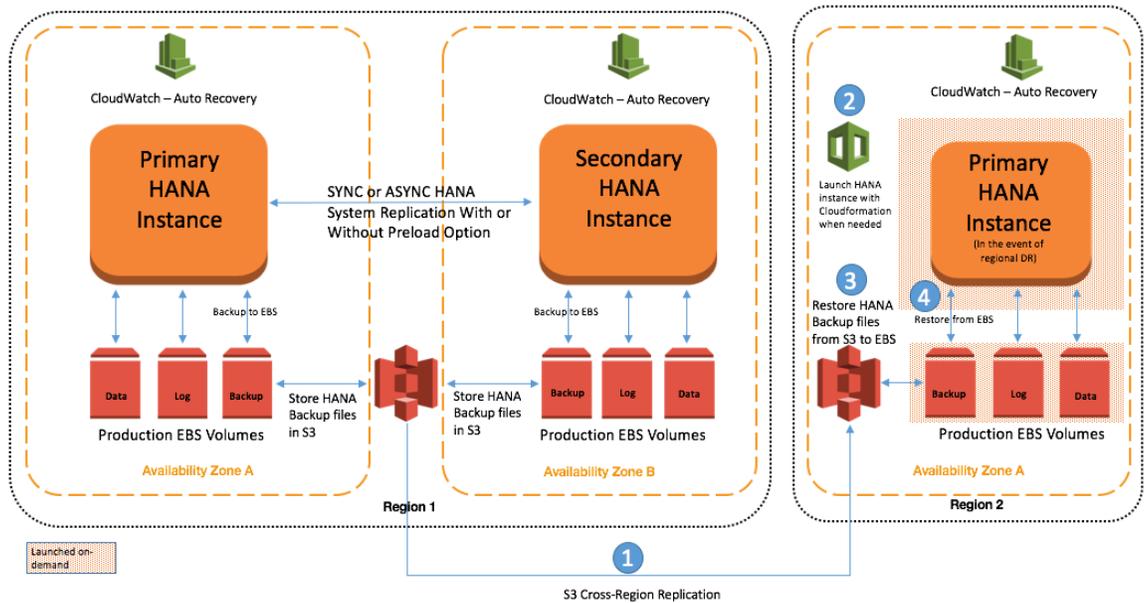
Your HANA instance becomes highly available when you deploy it in a Multi-AZ setup. However, if you are worried about region-level failure and want to set up disaster recovery for your HANA instance in another region in a cost-effective way, you can use cross-region replication for the HANA backup files stored in Amazon S3. When configured, cross-region replication in Amazon S3 ensures that the data stored in a bucket is automatically copied (asynchronously) to the target region.

When you need to recover your HANA instance in another region, you can use AWS CloudFormation templates to automate the launch of a new HANA instance, and copy the HANA backup files from Amazon S3 into a staging EBS volume. When the HANA backup is staged on an EBS volume, you can use standard HANA backup/recovery procedures to restore and recover your HANA database, and redirect your client traffic to your new HANA EC2 instance.

This architecture provides you with the advantage of implementing your HANA solution across multiple Availability Zones with the ability to fail over instantly during an Availability Zone failure. It also provides protection when a whole region becomes unavailable. You get the best RTO and RPO within one region and can recover your HANA instance in another region when needed. This model is suitable when you need the lowest possible RTO and RPO within a region, and an out-of-region disaster recovery option with moderate RTO and RPO.

Note that your out-of-region disaster recovery RPO is constrained by how often you store your HANA backup files in the S3 bucket and the time it takes to replicate your S3 bucket to your target region.

Summary

The following table summarizes all the options we discussed, and shows the RPO, RTO, and cost (of AWS infrastructure) associated with each option. This will help you determine the right solution to meet your needs.

Scenarios	Multi-AZ			Out-of-Region DR		
	RPO	RTO	Cost	RPO	RTO	Cost
Option 1 – Automatic recovery & HANA Backup/Restore	Medium	High	Low	n/a	n/a	n/a
Option 2 – Automatic recovery & HSR without data preload (warm standby)	Low	Medium	Medium	n/a	n/a	n/a
Option 3 – Automatic recovery & HSR without data preload (warm standby + Dev/QA)	Low	Medium	High**	n/a	n/a	n/a
Option 4 – Automatic recovery & HSR with data preload (hot standby)	Low	Low	High	n/a	n/a	n/a
Option 5 – Automatic Recovery & Multitier HSR (hot standby + out-of-region DR)	Low	Low	High	Low	Medium	High
Option 6 – Automatic Recovery & HSR with Amazon S3 Cross-Region Replication (hot standby + out-of-region DR)	Low	Low	High	Medium	High	Low

** Cost savings achieved by co-hosting Dev/QA

Triggering HANA System Replication Takeover

When there is a need for HANA System Replication takeover, you need to trigger the takeover in your secondary system. To trigger the takeover, follow the standard SAP HANA takeover process. If you have enabled automatic recovery, decide whether it is worth waiting for your system to be recovered in the primary Availability Zone before performing the takeover in the secondary Availability Zone. SAP OSS Note 2063657 (requires SAP Support Portal access) provides guidelines that you can follow to decide whether takeover is the best choice.[11]

Redirecting Client Traffic

In almost all scenarios, failover of the SAP HANA system alone does not guarantee business continuity. You need to ensure that your client applications (for example, the NetWeaver application server, JDBC, ODBC connection, etc.) can re-establish their connection with SAP HANA system after the failover. You can achieve this either by performing network-based IP redirection or network-

based DNS redirection of your SAP HANA system. For details about redirecting client traffic after a SAP HANA failover, see the "Client Connection Recovery" section in the [HANA Administration Guide](#).[12]

Automating HANA System Replication Failover

When a failure is detected at the hardware or software level, you can manually perform the failover process using SAP HANA Studio or the SAP HANA command-line interface. However, this could cause extended outage to your critical business processes due to the manual efforts required to complete all the needed steps. To reduce your down time, you can choose to implement a third-party cluster solution like SUSE's [SAP HANA Replication Failover](#)[13] or NEC's [NEC Express Cluster](#)[14] to automate the failover steps.

Starting with SAP HANA SPS09, you can also use the Python-based API included with SAP HANA to create your own HA/DR provider and integrate it with the SAP HANA System Replication takeover process to automate tasks such as restarting the primary node, IP redirection, DNS redirection, and shutdown of Dev/QA systems in the secondary node. To learn more about the HA/DR provider implementation, see [Implementing a HA/DR Provider](#) on the SAP website.[15]

Testing Your Solution

We highly recommend that you test your HA/DR solution periodically to ensure that the failover process that you have in place is accurate and up to date. Testing becomes even more critical whenever you upgrade or update the components of your system, such as SAP HANA, the operating system, and the HA cluster software. You need to ensure that your updates are replicated to the HA/DR environments as well. With AWS, you can periodically test your HA/DR solutions without incurring any long-term investment. You can spawn your HA/DR resources in your target Availability Zone or region during the test and then tear them down when testing is complete.

Conclusion

This whitepaper provided you an overview of the AWS global infrastructure and services that are relevant to implementing SAP HANA on AWS. By integrating the AWS global infrastructure and features such as automatic recovery and Amazon S3 with SAP HANA tool sets for high availability, you can deploy a robust, highly available, and disaster-tolerant SAP HANA system on the AWS platform in a cost-effective way.

Contributors

The following individuals and organizations contributed to this document:

- Sabareesan Radhakrishnan, Partner Solution Architect, AWS

Further Reading

For additional information, please consult the following sources:

- SAP HANA on AWS Quick Start Reference Deployment Guide

- High Availability for SAP HANA[16]

- How to Perform System Replication for SAP HANA

- Introduction: High Availability for SAP HANA

- SAP HANA on AWS Operations Guide[17]

- Useful SAP OSS Notes (these require SAP Support Portal access)

  - 1999880 - FAQ: SAP HANA System Replication[18]

  - 2057595 - FAQ: SAP HANA High Availability[19]

  - 2063657 - HANA System Replication takeover decision guideline[20]

  - 1913302 - HANA: Suspend DB connections for short maintenance tasks[21]

  - 2063657 - HANA System Replication takeover decision guideline[22]

Notes

[1] https://d0.awsstatic.com/enterprise-marketing/SAP/SAP_on_AWS_High_Availability_Guide_v3.2.pdf

[2] https://aws.amazon.com/about-aws/global-infrastructure/

[3] https://aws.amazon.com/ec2/

[4] https://aws.amazon.com/ebs/

[5] https://aws.amazon.com/s3/

[6] https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-recover.html

[7] https://docs.aws.amazon.com/quickstart/latest/sap-hana/

[8] http://scn.sap.com/docs/DOC-65585

[9] http://scn.sap.com/docs/DOC-47702

[10] http://scn.sap.com/docs/DOC-60361

[11] http://service.sap.com/sap/support/notes/2063657

[12] https://help.sap.com/hana/SAP_HANA_Administration_Guide_en.pdf

[13] https://www.suse.com/promo/sap/hana/replication.html

[14] http://www.nec.com/en/global/prod/expresscluster/en/collaterals/pdf/EXPRESSCLUSTER_x33_newfeatures_nec_en.pdf

[15] http://help.sap.com/saphelp_hanaplatform/helpdata/en/13/67c8fdefaa4808a7485b09815ae0f3/content.htm

[16] http://help.sap.com/saphelp_hanaplatform/helpdata/en/6d/252db7cdd044d19ad85b46e6c294a4/content.htm

[17] http://d0.awsstatic.com/enterprise-marketing/SAP/SAP_HANA_on_AWS_Implementation_and_Operations_Guide.pdf

[18] http://service.sap.com/sap/support/notes/1999880

[19] http://service.sap.com/sap/support/notes/2057595

[20] http://service.sap.com/sap/support/notes/2063657

[21] http://service.sap.com/sap/support/notes/1913302

[22] http://service.sap.com/sap/support/notes/2063657