SAP HANA on AWS

SAP HANA Guides

SAP HANA on AWS: SAP HANA Guides

The AWS Documentation website is getting a new look!

Try it now and let us know what you think. Switch to the new look >>

You can return to the original look by selecting English in the language selector above.

Table of Contents

SAP HANA Guides

This section of the SAP on AWS technical documentation provides information for SAP HANA users and partners, including information about planning, migrating, implementing, configuring, and operating your SAP HANA environment on the AWS Cloud. This section includes the following guides:

About this content set

SAP on AWS technical documentation provides detailed information on how to migrate, implement, configure, and operate SAP solutions on AWS.

Migrating SAP HANA to AWS: Patterns for AWS Migrations

SAP specialists, Amazon Web Services (AWS)

Last updated (p. 23): June 2019

This guide describes the most common scenarios, use cases, and options for migrating SAP HANA systems from on-premises or other cloud platforms to the Amazon Web Services (AWS) Cloud.

This guide is intended for SAP architects, SAP engineers, IT architects, and IT administrators who want to learn about the methodologies for migrating SAP HANA systems to AWS, or who want to have a better understanding of migration approaches to AWS in general.

This guide does not replace AWS and SAP documentation and is not intended to be a step-by-step, detailed migration guide. For a list of helpful resources, see the Additional Reading (p. 23) section. Information and recommendations regarding integrator and partner tools are also beyond the scope of this guide. Also, some of the migration scenarios may involve additional technology, expertise, and process changes, as discussed later in this guide (p. 13).

> **Note**
> To access the SAP notes and Knowledge Base articles (KBA) referenced in this guide, you must have an SAP ONE Support Launchpad user account. For more information, see the SAP Support website.

About this Guide

This guide is part of a content series that provides detailed information about hosting, configuring, and using SAP technologies in the AWS Cloud. For the other guides in the series, ranging from overviews to advanced topics, see https://aws.amazon.com/sap/docs/.

Migration Frameworks

Although this guide focuses on SAP HANA migrations to AWS, it is important to understand AWS migrations in a broader context. To help our customers conceptualize and understand AWS migrations in general, we have developed two major guidelines: 6 Rs and CAF.

6 Rs Framework

The 6 Rs migration strategy helps you understand and prioritize portfolio and application discovery, planning, change management, and the technical processes involved in migrating your applications to AWS. The 6 Rs represent six strategies listed in the following table that help you plan for your application migrations.

"R" migration strategy	Methodology
Rehosting	The application is migrated as is to AWS. This is also called a “lift-and-shift” approach.

"R" migration strategy	Methodology
Replatforming	The application is changed or transformed in some aspect as part of its migration to AWS.
Repurchasing	You move to a different application or solution on the cloud.
Refactoring / Re-architecting	The application is redesigned (for example, it's converted from a monolithic architecture to microservices) as part of the migration to AWS.
Retiring	The application is retired during migration to AWS.
Retaining	The application isn't migrated.

The decision tree diagram in Figure 1 will help you visualize the end-to-end process, starting from application discovery and moving through each 6 R strategy.

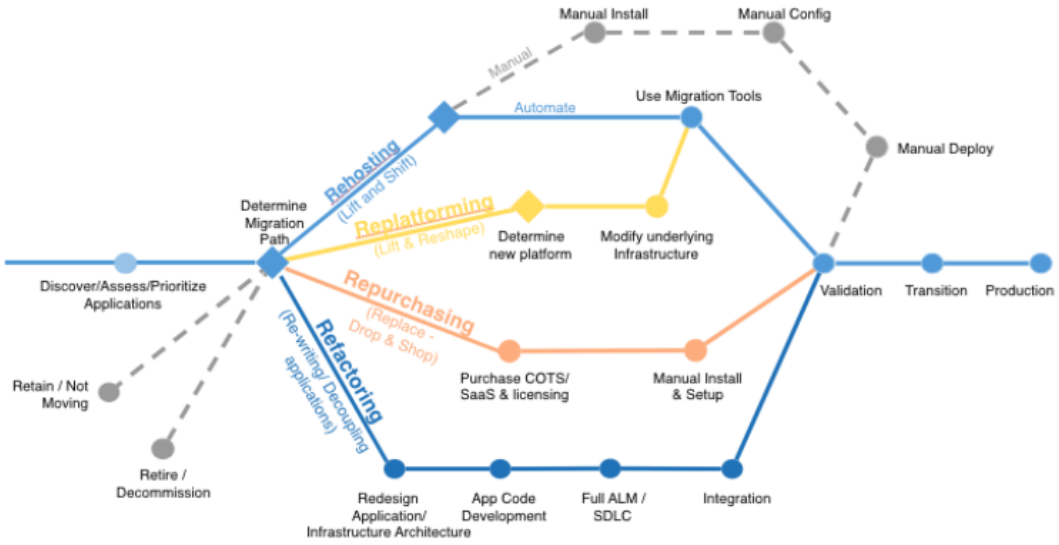

Figure 1: 6 Rs framework

The two strategies that are specifically applicable for SAP HANA migrations to AWS are rehosting and replatforming. Rehosting is applicable when you want to move your SAP HANA system as is to AWS. This type of migration involves minimal change and can be seen as a natural fit for customers who are already running some sort of SAP HANA system. Replatforming is applicable when you want to migrate from an *anyDB* source database (such as IBM DB2, Oracle Database, or SQL Server) to an SAP HANA database.

AWS CAF Framework

The second guideline is the AWS Cloud Adoption Framework (CAF). The AWS CAF breaks down the complex process of planning a move to the cloud into manageable pieces called *perspectives*. Perspectives represent essential areas of focus that span people, processes, and technology. Capabilities within each perspective identify the areas of your organization that require attention. From this

information, you can build an action plan organized into prescriptive work streams that support a successful cloud journey. Both the CAF and 6 Rs frameworks help you understand and plan the broader context of an AWS migration and what it means to you and your company.

Planning

Before you start migrating your SAP environment to AWS, there are some prerequisites that we recommend you go over, to ensure minimal interruptions or delays. For details, see the SAP on AWS overview. The following sections discuss additional considerations for planning your migration.

Understanding On-Premises Resource Utilization

If you are planning to rehost your on-premises SAP HANA environment on AWS, AWS Application Discovery Service can help you understand the utilization of resources as well as hardware configuration, performance data, and network connections in your on-premises SAP HANA environment. You can use this information to ensure that appropriate communication ports are enabled between SAP HANA and other systems in the security groups or virtual private clouds (VPCs) on AWS.

Application Discovery Service can be deployed in an agentless mode (for VMware environments) or with an agent-based mode (all VMs and physical servers). We recommend that you run Application Discovery Service for a few weeks to get a complete, initial assessment of how your on-premises environment is utilized, before you migrate to AWS.

Reviewing AWS Automation Tools for SAP

It is a good idea to review AWS automation tools and services that can help you migrate your SAP environment to AWS. For example, AWS Quick Starts are automated reference deployments for workloads such as SAP HANA and SAP NetWeaver application servers. For details, see the Migration Tools and Methodologies (p. 8) section later in this guide.

Data Tiering

If you are planning on replatforming your SAP HANA environment on AWS, you can also consider different services and options available to you for distributing your data into warm and cold SAP-certified storage solutions like SAP HANA dynamic tiering or Hadoop on AWS. Currently, SAP supports Cloudera, HortonWorks, and MapR as possible Hadoop distributions for SAP HANA. See the SAP HANA administration guide for details on how to connect SAP HANA systems with Hadoop distribution using smart data access.

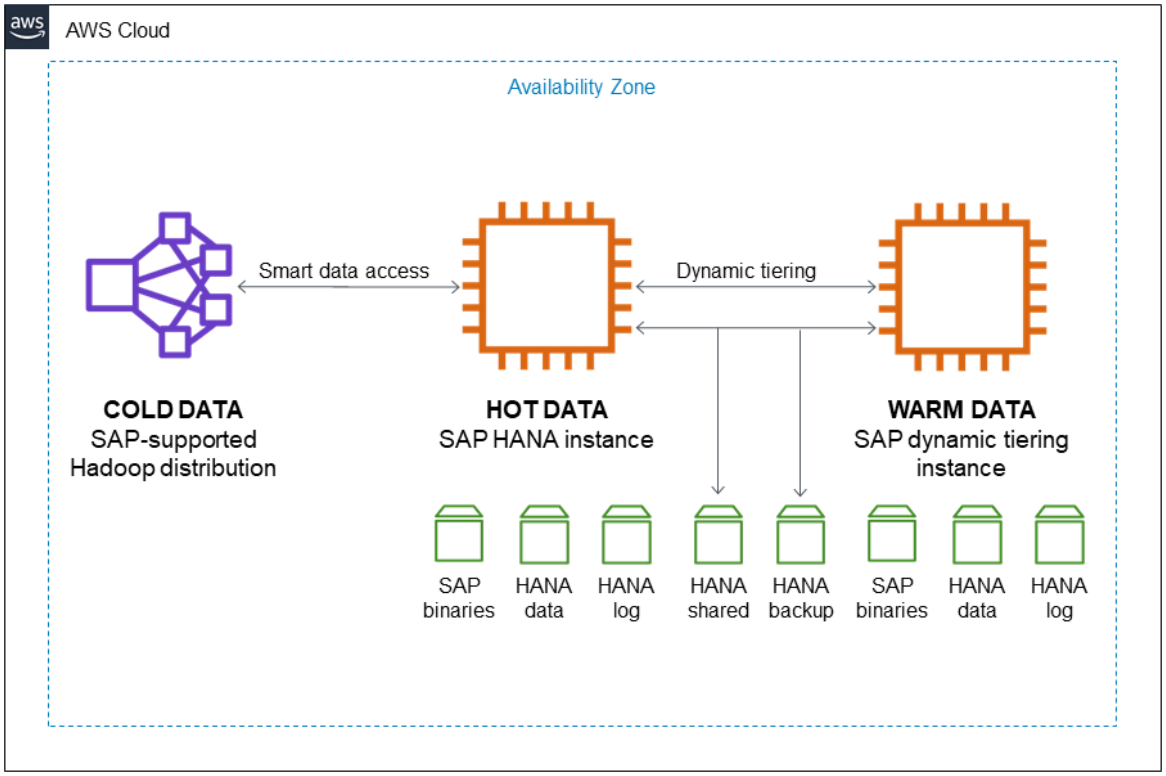

Figure 2: Data tiering

Migrating warm or cold data can further simplify your SAP environment and help reduce your total cost of ownership (TCO). For more information, see our web post for SAP dynamic tiering sizes and recommendations.

Prerequisites

SAP HANA system migration requires a moderate to high-level knowledge of the source and target IT technologies and environments. We recommend that you familiarize yourself with the following information:

AWS Cloud architecture and migration:

- AWS Well-Architected Framework
- An Overview of the AWS Cloud Adoption Framework
- Architecting for the Cloud: Best Practices
- Migrating Your Existing Applications to the AWS Cloud

AWS services:

- Amazon Virtual Private Cloud (Amazon VPC)
- Amazon Elastic Compute Cloud (Amazon EC2)
- Amazon Elastic Block Store (Amazon EBS)
- Amazon Simple Storage Service (Amazon S3)

SAP on AWS:

- Implementing SAP Solutions on Amazon Web Services
- SAP on AWS Implementation and Operations Guide
- SAP HANA Quick Start Reference Deployment
- SAP HANA Environment Setup on AWS
- SAP on Amazon Web Services High Availability Guide

SAP HANA Sizing

The size of the SAP HANA system required on the AWS Cloud depends on the migration scenario. As mentioned earlier, migrating SAP HANA to AWS involves two possible scenarios: rehosting or replatforming.

Memory Requirements for Rehosting

Because rehosting implies that you are already running SAP HANA, you can determine the size of the SAP HANA system you need on the AWS Cloud from the peak memory utilization of your existing SAP HANA system. You may have oversized your on-premises SAP HANA environment (for example, to support future growth), so measuring peak memory utilization is a better approach than measuring allocated memory. When you have determined the base memory requirement, you should choose the smallest SAP-certified EC2 instance that provides more memory than your base requirement.

There are three ways to determine peak memory utilization of your existing SAP HANA system:

- SAP HANA Studio: The overview tab of the SAP HANA Studio administration view provides a memory utilization summary.
- SAP EarlyWatch alerts: This is a free, automated service from SAP that helps you monitor major administrative areas of your SAP system. See the SAP portal for details.
- SQL statements: SAP provides SQL statements that you can use to determine peak memory utilization. For details, see SAP KBA 1999997 – FAQ: SAP HANA Memory and SAP Note 1969700 – SQL statement collection for SAP HANA.

> **Tip**
> We recommend determining peak memory utilization for a timeframe during which your system utilization is likely to be high (for example, during year-end processing or a major sales event).

Memory Requirements for Replatforming

The replatforming scenario involves two possibilities:

- You are already running SAP HANA but you want to change your operating system—for example, from Red Hat Enterprise Linux (RHEL) to SUSE Linux Enterprise Server (SLES) or the other way around—when you migrate to the AWS Cloud, or you are migrating from an IBM POWER system to the x86 platform. In this case, you should size SAP HANA as described for the rehosting scenario.
- You are migrating from *anyDB* to SAP HANA. There are multiple ways you can estimate your memory requirements:
  - SAP standard reports for estimation: This is the best possible approach and is based on standard sizing reports provided by SAP. For examples, see the following SAP Notes:
    - 1736976 – Sizing Report for BW on HANA
    - 1637145 – SAP BW on HANA: Sizing SAP In-Memory Database

  - 1872170 - Business Suite on HANA and S/4HANA sizing report
  - 1736976 – Sizing Report for BW on HANA
- SQL statements: SAP provides scripts that you can run in your existing environment to get high-level SAP HANA sizing estimates. These scripts run SQL statements against your existing database to estimate SAP HANA memory requirements. For more information, see SAP Note 1514966 - SAP HANA 1.0: Sizing SAP In-Memory Database.
- Rule of thumb: See the PDF attached to SAP Note 1514966 - SAP HANA 1.0: Sizing SAP In-Memory Database for instructions on estimating SAP HANA memory requirements manually. Note that this will be a very rough and generic estimate.

You should also consider the following SAP notes and Knowledge Base articles for SAP HANA sizing considerations:

- 706478 – Preventing Basis tables from increasing considerably
- 1855041 – Sizing Recommendation for Master Node in BW-on-HANA
- 1702409 – HANA DB: Optimal number of scale out nodes for BW on HANA

Instance Sizing for SAP HANA

AWS offers SAP-certified systems that are configured to meet the specific SAP HANA performance requirements (see SAP Note 1943937 – Hardware Configuration Check Tool - Central Note and the SAP Certified SAP HANA Hardware Directory). After you have determined your SAP HANA sizing, you can map your requirements to the EC2 instance family sizes. That is, you map the maximum amount of memory required for each of your SAP HANA instances to the maximum amount of memory available for your desired EC2 instance type. You should also consider appropriate storage volume types and sizes to ensure optimal performance of the SAP HANA database. For best practices and recommendations for volume types and file system layout, see the Planning the Deployment section of the SAP HANA Quick Start deployment guide.

> **Note**
> Only production SAP HANA systems need to run on certified configurations that meet SAP HANA key performance indicators (KPIs). SAP provides more flexibility when running SAP HANA non-production systems. For more information, see SAP HANA TDI – FAQ and OSS Note 2271345 on the SAP website.

Network Planning and Sizing

You will need to consider network planning and sizing for the amount of data you will be transferring to AWS. Data transfer time depends on network bandwidth available to AWS and influences total downtime. Higher bandwidth helps with faster data transfer and helps reduce overall migration time. For non-production systems where downtime isn't critical, you can use a smaller network pipe to reduce costs. Alternatively, to transfer extremely large data, you can use services like AWS Snowball for a physical (non-network) transport of data to AWS. We'll discuss AWS Snowball more extensively later in this guide.

As a guideline, you can use this formula to help estimate how long your network data transfer might take:

(Total bytes to be transferred / Transfer rate per second) = Total transfer time in seconds

For example, for a 1 TB SAP HANA appliance, the total bytes to be transferred is usually 50% of the memory, which would be 512 GB. The transfer rate per second is your network transfer rate—if you had a 1 Gb AWS Direct Connect connection to AWS, you could transfer up to 125 MB per second, and your total data transfer time would be:

512 GB / 125 MB per second = 4,096 seconds (or 1.1 hours)

After you determine the amount of data you need to transfer and how much time you have available to transfer the files, you can determine the AWS connectivity options that best fit your cost, speed, and connectivity requirements. Presenting all available network connectivity options is beyond the scope of this document; see the Additional Reading (p. 23) section of this document for more detailed references.

SAP HANA Scale-up and Scale-out

AWS provides several types of EC2 instances for SAP HANA workloads. This gives you options for your SAP HANA scale-up and scale-out deployments. In a scale-up scenario, you utilize the compute, memory, network, and I/O capacity of a single EC2 instance. If you require more capacity, you can resize your instances to a different EC2 instance type. For example, if you're using an R4 instance type and it becomes too small for your workload, you can change it to an R5, X1, or X1e instance type. The limitation is the maximum capacity of a single EC2 instance. In AWS, scale-up enables you to start with the smallest EC2 instance type that meets your requirements and grow as needed. If your requirements change or new requirements surface, you can easily scale up to meet the changing requirements.

In a scale-out scenario, you add capacity to your SAP HANA system by adding new EC2 instances to the SAP HANA cluster. For example, once you reach the maximum memory capacity of a single EC2 instance, you can scale out your SAP HANA cluster and add more instances. AWS has certified SAP HANA scale-out clusters that support up to 100 TiB of memory. Please note that the minimum number of recommended nodes in an SAP HANA scale-out cluster can be as low as two nodes; for more information, see SAP Note 1702409 - HANA DB: Optimal number of scale out nodes for BW on HANA. It's likely that your sizing estimates will reveal the need to plan for a scale-out configuration before you start your SAP HANA migration. AWS gives you the ability to easily deploy SAP HANA scale-out configurations when you use the SAP HANA Quick Start.

The following table illustrates example scale-up and scale-out sizing.

Scenario	Source configuration	Target configuration
Scale-up	r4.8xlarge	r4.16xlarge
Scale-up	r4.16xlarge	x1.16xlarge
Scale-up	x1.32xlarge	x1e.32xlarge
Scale-out	3 nodes of x1.16xlarge	4 nodes of x1.16xlarge
Scale-out	x1.32xlarge	3 nodes of x1.16xlarge

When you finalize your SAP sizing and SAP HANA deployment models, you can plan your migration strategy.

In addition to SAP HANA sizing, you may also need to size your SAP application tier. To find the SAP Application Performance Standard (SAPS) ratings of SAP-certified EC2 instances, see SAP Standard Application Benchmarks and the SAP on AWS support note on the SAP website (SAP login required).

Migration Tools and Methodologies

This section provides an introduction to the tools and methodologies available to you for your SAP system migration.

AWS Quick Starts

AWS Quick Starts are automated reference deployments designed by AWS solutions architects and AWS partners. These reference deployments implement key technologies automatically on the AWS Cloud, often with a single click and in less than an hour. You can build your test or production environment in a few steps, and start using it immediately. For SAP HANA migrations, you can use either the SAP HANA or the SAP NetWeaver Quick Starts to automatically provision, deploy, configure, and install your SAP HANA and SAP NetWeaver system in the AWS Cloud. Using AWS Quick Starts saves you time and ensures repeatability, because you don't have to develop custom deployment scripts or manually deploy, configure, and install your SAP HANA systems. As a result, you can often migrate your SAP systems faster.

Migration Using DMO with System Move

SAP has enhanced the database migration option (DMO) of their Software Update Manager (SUM) tool to accelerate the testing of SAP application migrations (see SAP Note 2377305). DMO with System Move enables you to migrate your SAP system from your on-premises environment to AWS by using a DMO tool and a special export and import process. You can use AWS services such as Amazon S3, Amazon EFS (over AWS Direct Connect), AWS Storage Gateway file interface, and AWS Snowball to transfer your SAP export files to AWS.

You can then use the AWS Quick Start for SAP HANA to rapidly provision SAP HANA instances and build your SAP application servers on AWS, when you are ready to trigger the import process of the DMO tool.

The SUM DMO tool can convert data from *anyDB* to SAP HANA or SAP ASE, with OS migrations, release/enhancement pack upgrades, and Unicode conversions occurring at the same time. Results are written to flat files, which are transferred to the target SAP HANA system on AWS. The second phase of DMO with System Move imports the flat files and builds the migrated SAP application with the extracted data, code, and configuration. Here's a conceptual flow of the major steps involved:

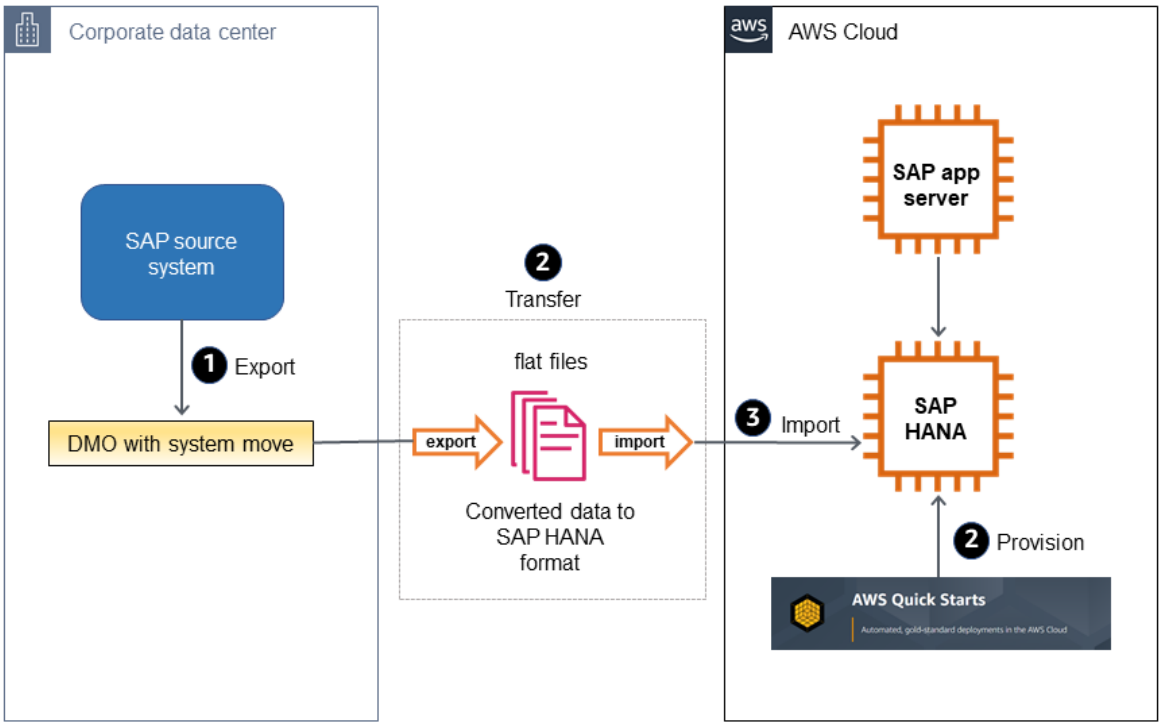

Figure 3: DMO with System Move

SAP HANA Classical Migration

SAP offers the SAP HANA classical migration option for migrating from other database systems to SAP HANA. This option uses the SAP heterogeneous system copy process and tools. To copy the exported files, you can use the options described in the Backup/Restore Tools (p. 10) section later in this guide. For details on the classical migration approach, see the classical migration overview on the SAP website.

SAP Software SUM DMO

SAP offers the standard SUM DMO approach as a one-step migration option from other database systems to HANA. This option uses the SAP DMO process and tool to automate multiple required migration steps. This is a preferred option if you are already running SAP on *anyDB* on AWS, as it will improve your migration times to SAP HANA, since there is no need for data export/import at a file system level. For details, see the DMO of SUM overview on the SAP website.

SAP HANA HSR

SAP HANA System Replication (HSR) is a tool for replicating the SAP HANA database to a secondary database or location. The secondary database is an exact copy of the primary database and can be used as the new primary database in the event of a takeover. The advantage of HSR is that it replicates the data directly from source to target. For details, see SAP HANA Disaster Recovery Support in the *SAP HANA Administration Guide* and the High Availability and Disaster Recovery Options for SAP HANA on AWS whitepaper.

SAP HANA HSR with Initialization via Backup and Restore

SAP supports the option of initializing the HSR target system with a backup and restore process. Using backup and restore can be useful if the network connection between your source SAP HANA system and the target system does not have enough bandwidth to replicate the data in a timely manner. Additionally, you may not want the data replication to consume part of your network traffic bandwith. For details, see SAP Note 1999880 – FAQ: SAP HANA System Replication.

Backup/Restore Tools

Backup and restore options are tried-and-true mechanisms for saving data on a source system and restoring it to another destination. AWS has various storage options available to help facilitate data transfer to AWS. Some of those are explained in this section. We recommend that you discuss which option would work best for your specific workload with your systems integrator (SI) partner or with an AWS solutions architect.

- AWS Storage Gateway: This is a virtual appliance installed in your on-premises data center that helps you replicate files, block storage, or tape libraries by integrating with AWS storage services such as Amazon S3 and by using standard protocols like Network File system (NFS) or Internet Small Computer System Interface (iSCSI). AWS Storage Gateway offers file-based, volume-based, and tape-based storage solutions. For SAP systems, we will focus on file replication using a file gateway and block storage replication using a volume gateway. For scenarios where multiple backups or logs need to be continuously copied to AWS, you can copy these files to the locally mounted storage and they will be replicated to AWS.

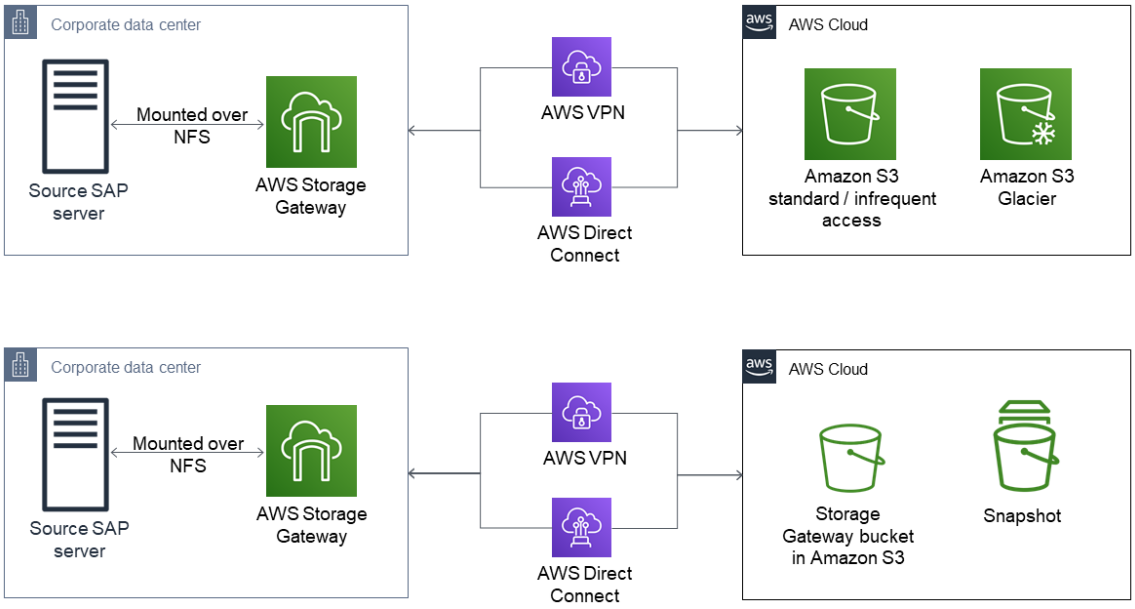

Figure 4: SAP file replication with AWS Storage Gateway

See the SAP ASE Cloud Backup to Amazon S3 using AWS File Gateway whitepaper on the SAP website to learn how to use a file gateway to manage backup files of SAP ASE on AWS with Amazon S3, with the STANDARD-IA (infrequent access) and Amazon S3 Glacier storage classes. For more information about these storage classes, see the Amazon S3 documentation.

- Amazon EFS file transfer: AWS provides options to copy data from an on-premises environment to AWS by using Amazon Elastic File System (Amazon EFS). Amazon EFS is a fully managed service, and you pay only for the storage that you use. You can mount an Amazon EFS file share on your on-premises server, as long as you have AWS Direct Connect set up between your corporate data center and AWS. This is illustrated in Figure 5.

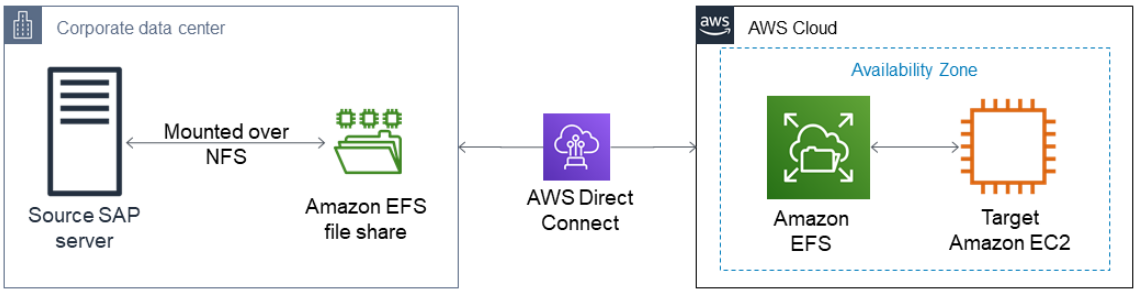

Figure 5: Transferring SAP files with Amazon EFS

AWS Snowball

With AWS Snowball, you can copy large amounts of data from your on-premises environment to AWS, when it's not practical or possible to copy the data over the network. AWS Snowball is a storage appliance that is shipped to your data center. You plug it into your local network to copy large volumes of data at high speed. When your data has been copied to the appliance, you can ship it back to AWS, and your data will be copied to Amazon S3 based on the desired target storage destination that you specify. AWS Snowball is very useful when you're planning very large, multi-TB SAP system migrations.

For more information, see *When should I consider using Snowball instead of the Internet* in the AWS Snowball FAQ.

Amazon S3 Transfer Acceleration

Amazon S3 Transfer Acceleration provides a faster way to copy data from your on-premises environment to AWS by copying data first to Amazon CloudFront edge locations that are closest to the source, and then using an optimized network path to copy data to Amazon S3. There is a network charge associated with this type of transfer. You can run an AWS-provided test tool to compare the speed of Amazon S3 Transfer Acceleration to standard Amazon S3 data transfer. For SAP workloads, you can copy backups or DB logs at regular intervals over Amazon S3 Transfer Acceleration to reduce the transfer time, if your regular network connection is slow—for example, if your SAP environment is hosted in a location that doesn't have very strong internet connectivity. For more information, see the Amazon S3 documentation.

Amazon EC2 Instance Resize

Amazon EC2 provides you with the ability to easily change your instance type in minutes, from the Amazon EC2 console, the AWS Command Line Interface (AWS CLI), or the Amazon EC2 API. You can start with an instance type that meets your current needs and size your instance up or down, when your requirements change. When you change your EC2 instance type, all instance metadata, including the IP address, instance ID, and hostname, remains the same. This enables you to migrate your SAP HANA to a new instance type seamlessly, without incurring a longer downtime. For details, see the Changing the Instance Type in the Amazon EC2 documentation.

AMIs

You can use an Amazon Machine Image (AMI) to launch any EC2 instance. You can create an AMI of an EC2 instance that hosts SAP HANA, including the attached EBS volumes, through the Amazon EC2 console, the AWS CLI, or the Amazon EC2 API. You can then use the AMI to launch a new EC2 instance with SAP HANA in any Availability Zone within the AWS Region where the AMI was created. You can also copy your AMI to another AWS Region and use it to launch a new instance. You can use this feature to move your SAP HANA instance to another Availability Zone or AWS Region, or to change the tenancy type of your EC2 instance. For example, you can create an AMI of your EC2 instance with default tenancy and use it to launch a new EC2 instance with host or dedicated tenancy and vice versa. For details, see the Amazon Machine Images (AMIs) in the Amazon EC2 documentation.

Migration Scenarios

The following table lists the migration scenarios that we will cover in detail in this guide. The tools and methodologies listed in the table were discussed in the previous section.

Migration scenario	Source database	Target database	Migration tool or methodology
Migration of *anyDB* from other platforms to AWS*	*anyDB* (any non-SAP HANA database such as IBM DB2, Oracle Database, or SQL Server)	SAP HANA	[✓] SAP HANA classical migration [✓] SAP DMO with System Move
Migration of SAP HANA from other platforms to AWS*	SAP HANA (scale-up and scale-out considerations apply here as well)	SAP HANA	[✓] SAP HANA backup and restore [✓] SAP HANA classical migration (considered

Migration scenario	Source database	Target database	Migration tool or methodology
			a homogeneous system copy in this scenario)** [✓] SAP HANA HSR [✓] SAP HANA HSR with initialization via backup and restore
Migration of SAP HANA from an existing EC2 instance to an EC2 High Memory instance	SAP HANA	SAP HANA	[✓] Instance resize [✓] Amazon Machine Image (AMI) [✓] SAP HANA backup and restore [✓] SAP HANA HSR

***** Other platforms include on-premises infrastructures and other cloud infrastructures outside of AWS.

****** See SAP Note 1844468 – Homogeneous system copy on SAP HANA.

Migrating *AnyDB* to SAP HANA on AWS

Migrating from *anyDB* to HANA typically involves changes to the database platform and sometimes includes operating system changes. However, migration might also involve additional technical changes and impacts, such as the following:

- SAP ABAP code changes. For example, you might have custom code that has database or operating system dependencies, such as database hints coded for the *anyDB* platform. You might also need to change custom ABAP code so it performs optimally on SAP HANA. See SAP's recommendations and guidance for these SAP HANA-specific optimizations. For details and guidance, see Considerations for Custom ABAP Code During a Migration to SAP HANA and SAP Notes 1885926 – ABAP SQL monitor and 1912445 – ABAP custom code migration for SAP HANA on the SAP website.
- Operating system-specific dependencies such as custom file shares and scripts that would need to be re-created or moved to a different solution.
- Operating system tunings (for example, kernel parameters) that would need to be accounted for. Note that the AWS Quick Start for SAP HANA incorporates best practices from operating system partners like SUSE and Red Hat for SAP HANA.
- Technology expertise such as Linux administration and support, if your organization doesn't already have experience with Linux.

SAP provides tools and methodologies such as classical migration and SUM DMO to help its customers with the migration process for this scenario. (For more information, see the section Migration Tools and Methodologies (p. 8).) AWS customers can use the SAP SUM DMO tool (p. 10) to migrate their database to SAP HANA on AWS. Some considerations for the SAP SUM DMO method are network bandwidth, amount of data to be transferred, and the amount of time available for the data to be transferred.

Implementing SAP HANA on AWS enables quick provisioning of scale-up and scale-out SAP HANA configurations and enables you to have your SAP HANA system available in minutes. In addition to fast

provisioning, AWS lets you quickly scale up by changing your EC2 instance type, as discussed earlier in the SAP HANA Sizing (p. 6) section. With this capability, you can react to changing requirements promptly and focus less on getting your sizing absolutely perfect. This means that you can spend less time sizing (that is, you can move through your project's planning and sizing phase faster) knowing that you can scale up later, if needed.

Migrating SAP HANA from Other Platforms to AWS

This scenario is more straightforward than migrating from *anyDB*, because you're already using SAP HANA. For this migration, you need to map your existing SAP HANA systems and sizing that are on a different platform to SAP HANA solutions on AWS.

EC2 instance memory capabilities give you the option to consolidate multiple SAP HANA databases on a single EC2 instance (scale-up) or multiple EC2 instances (scale-out). SAP calls these options HANA and ABAP One Server, Multiple Components in One Database (MCOD), Multiple Components in One System (MCOS), and Multitenant Database Containers (MDC). It is beyond the scope of this guide to recommend specific consolidation combinations; for possible combinations, see SAP Note 1661202 – Support for multiple applications on SAP HANA.

This migration scenario involves provisioning your SAP HANA system on AWS, backing up your source database, transferring your data to AWS, and installing your SAP application servers. If you are resizing your HANA environment from scale-up to scale-out, please follow the process highlighted in SAP Note 2130603. If you are resizing your HANA environment from scale-out to scale-up, refer to SAP Note 2093572. Depending on your specific scenario, you can use standard backup and restore, SAP HANA classical migration, SAP HANA HSR, AWS Server Migration Service (AWS SMS), or third-party continuous data protection (CDP) tools; see the following sections for details on each option.

Option 1: SAP HANA Backup and Restore

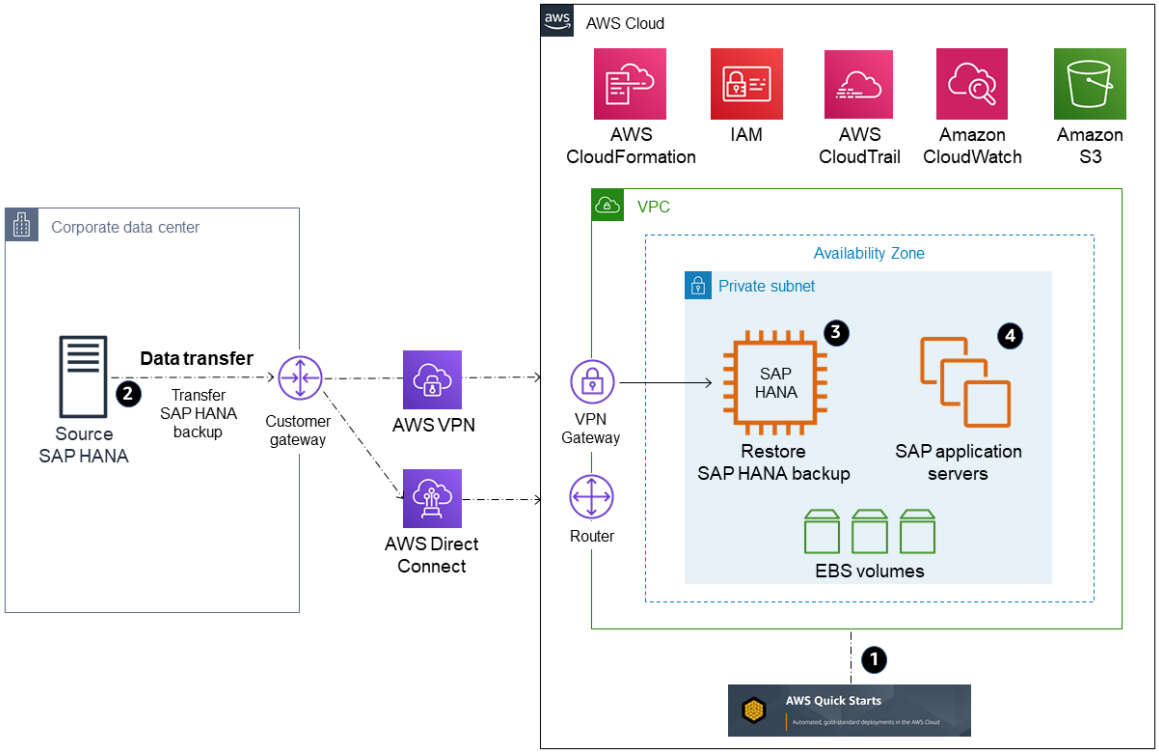

Figure 6: Backup and restore

1. Provision your SAP HANA system and landscape on AWS. (The AWS Quick Start for SAP NetWeaver can help expedite and automate this process for you.)
2. Transfer (**sftp** or **rsync**) a full SAP HANA backup, making sure to transfer any necessary SAP HANA logs for point-in-time recovery, from your source system to your target EC2 instance on AWS. A general tip here is to compress your files and split your files into smaller chunks to parallelize the transfer. If your transfer destination is Amazon S3, using the **aws s3 cp** command will automatically parallelize the file upload for you. For other options for transferring your data to AWS, see the AWS services listed previously in the Backup/Restore Tools (p. 10) section.
3. Recover your SAP HANA database.
4. Install your SAP application servers. (Skip this step if you used the AWS Quick Start for SAP NetWeaver in step 1.)
5. Depending on your application architecture, you might need to reconnect your applications to the newly migrated SAP HANA system.

Option 2: SAP HANA Classical Migration

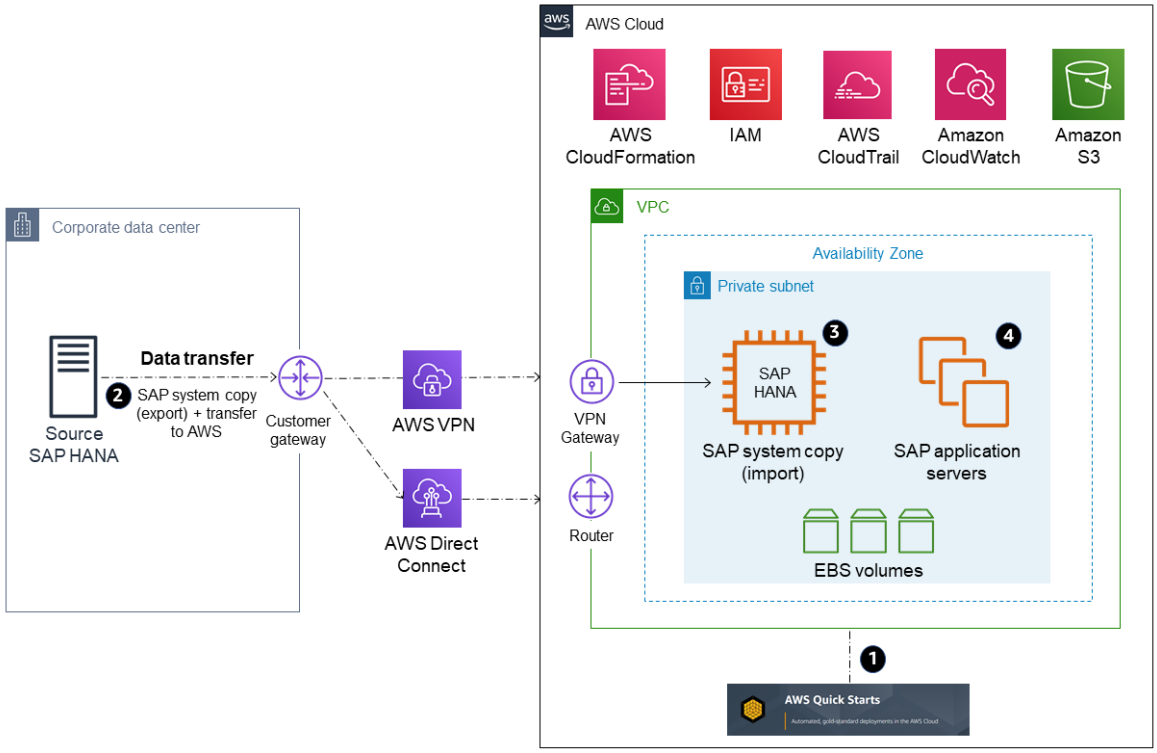

Figure 7: SAP HANA classical migration

1. Provision your SAP HANA system and landscape on AWS. (The AWS Quick Start for SAP NetWeaver can help expedite and automate this process for you.)
2. Perform an SAP homogeneous system copy to export your source SAP HANA database. You may also choose to use a database backup as the export; see SAP Note 1844468 – Homogeneous system copy on SAP HANA. When export is complete, transfer your data into AWS.
3. Continue the SAP system copy process on your SAP HANA system on AWS to import the data you exported in step 2.
4. Install your SAP application servers. (Skip this step if you used the AWS Quick Start for SAP NetWeaver in step 1.)

5. Depending on your application architecture, you might need to reconnect your applications to the newly migrated SAP HANA system.

Option 3: SAP HANA HSR

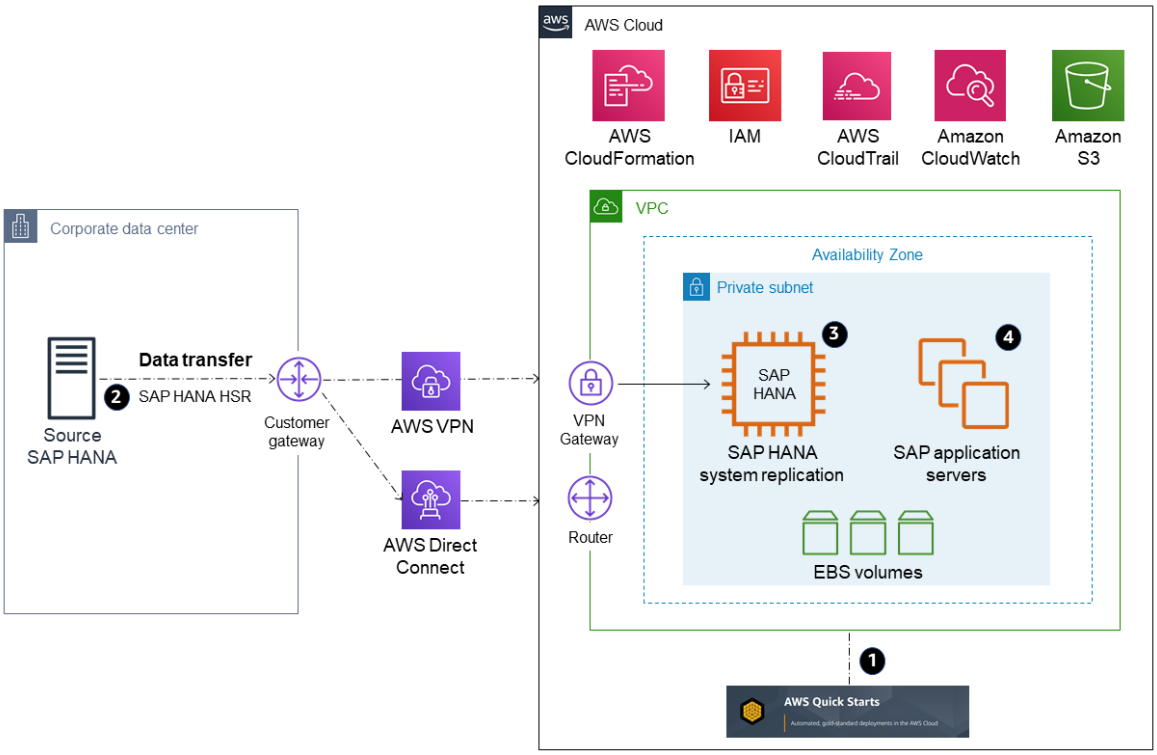

Figure 8: SAP HANA system replication

1. Provision your SAP HANA system and landscape on AWS. (The AWS Quick Start for SAP NetWeaver can help expedite and automate this process for you.) To save costs, you might choose to stand up a smaller EC2 instance type.
2. Establish asynchronous SAP HANA system replication from your source database to your standby SAP HANA database on AWS.
3. Perform an SAP HANA takeover on your standby database.
4. Install your SAP application servers. (Skip this step if you used the AWS Quick Start for SAP NetWeaver in step 1.)
5. Depending on your application architecture, you might need to reconnect your applications to the newly migrated SAP HANA system.

Option 4: SAP HANA HSR (with Initialization via Backup and Restore)

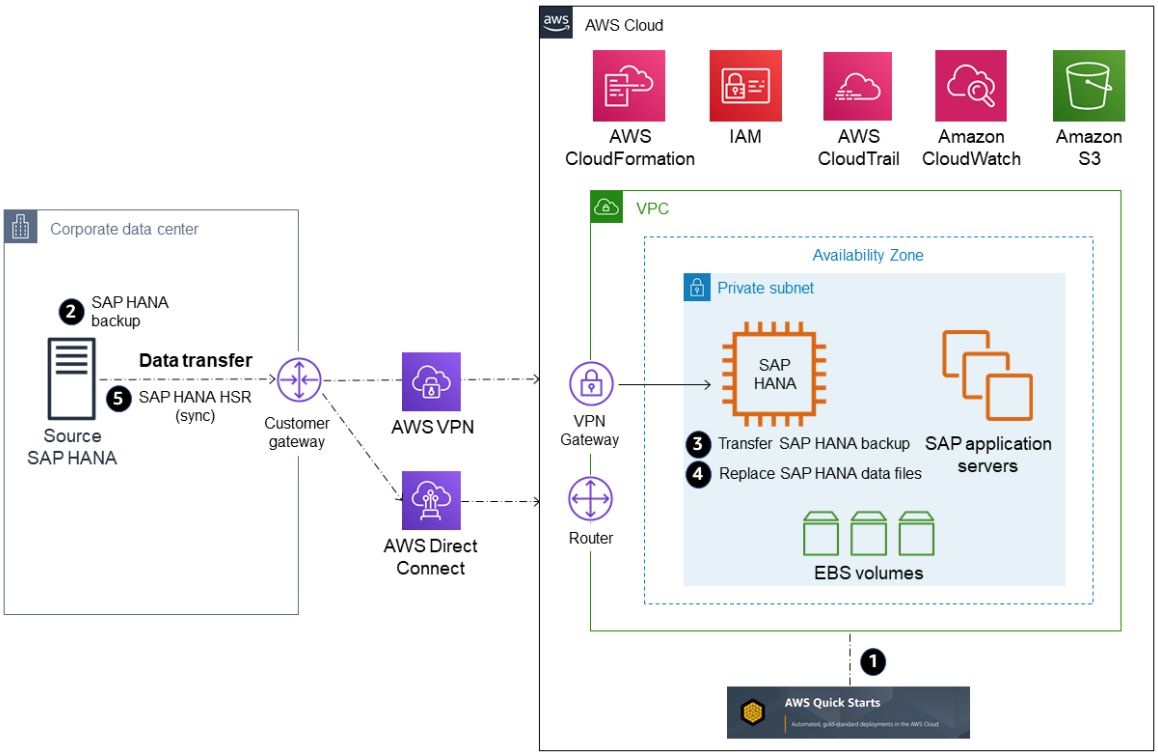

Figure 9: SAP HANA system replication (with initialization via backup and restore)

1. Provision your SAP HANA system and landscape on AWS. (The AWS Quick Start for SAP NetWeaver can help expedite and automate this process for you.) To save costs, you might choose to stand up a smaller EC2 instance type.
2. Stop the source SAP HANA database and obtain a copy of the data files (this is essentially a cold backup). After the files have been saved, you may start up your SAP HANA database again.
3. Transfer the SAP HANA data files to AWS, to the SAP HANA server you provisioned in step 1. (For example, you can store the data files in the /backup directory or in Amazon S3 during the transfer process.)
4. Stop the SAP HANA database on the target system in AWS. Replace the SAP HANA data files (on the target server) with the SAP HANA data files you transferred in step 3.
5. Start the SAP HANA system on the target system and establish asynchronous SAP HANA system replication from your source system to your target SAP HANA system in AWS.
6. Perform an SAP HANA takeover on your standby database.
7. Install your SAP application servers. (Skip this step if you used the AWS Quick Start for SAP NetWeaver in step 1.)
8. Depending on your application architecture, you might need to reconnect your applications to the newly migrated SAP HANA system.

Migrating SAP HANA on AWS to an EC2 High Memory Instance

EC2 High Memory instances provide up to 12 TiB of memory to support very large SAP HANA databases. EC2 High Memory instances support SUSE Linux Enterprise Server for SAP Applications (SLES for SAP) 12 SP3 or above, and Red Hat Enterprise Linux for SAP Solutions (RHEL for SAP) 7.4 or above. See the SAP HANA hardware directory to see the list of supported operating systems for your instance type. EC2 High Memory instances can be launched only as Amazon EC2 Dedicated Hosts with host tenancy.

Before you start your migration, make sure that an EC2 High Memory instance is allocated to your target account, Availability Zone, and AWS Region.

You have several options for migrating your existing SAP HANA workload on AWS to an EC2 High Memory instance, as discussed in the following sections.

Option 1: Resizing an Existing EC2 Instance with Host or Dedicated Tenancy

If your existing EC2 instance is running with host or dedicated tenancy, you can follow the steps in this section to migrate it to an EC2 High Memory instance. With this option, all your instance properties, including IP addresses, hostnames, and EBS volumes, remain the same after migration.

Figure 10 provides a high-level illustration of this method.

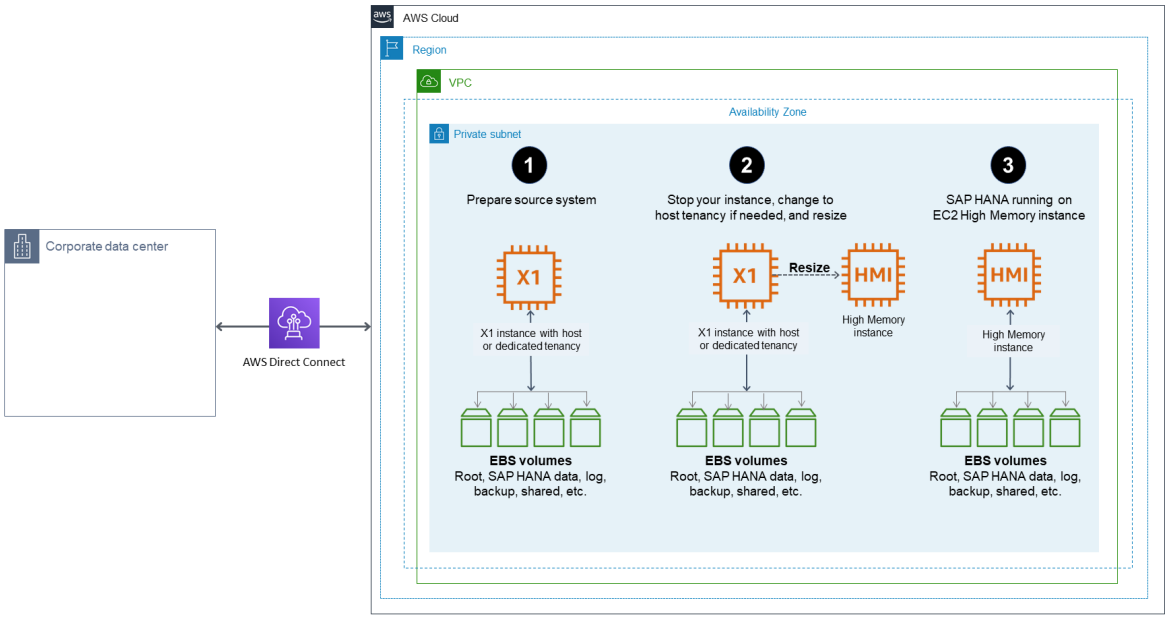

Figure 10: Migrating to an EC2 High Memory instance with resize option

1. Make sure that your source system is running on a supported operating system version. If not, you might have to upgrade your operating system before resizing to an EC2 High Memory instance.
2. If you are using SUSE Linux, make sure that your source system is configured to load the required non-volatile memory express (NVMe) drivers during boot. This step is not required if you are using a SUSE image that has a date stamp later than 20180810. For details, see Changing instance Types in AWS For SLES 12 SP2 and SLES 12 SP3 Based HVM instances on the SUSE blog.

3. EC2 High Memory instances are based on the Nitro system. On Nitro-based instances, EBS volumes are presented as NVMe block devices. If your source system has any mount point entries in `/etc/fstab` with reference to block devices such as `/dev/xvd<x>`, you need to create a label for these devices and mount them by label before migrating to EC2 High Memory instances. Otherwise, you will face issues when you start SAP HANA on an EC2 High Memory instance.
4. Make sure that you don't have more than 13 EBS volumes attached to your instance (excluding root volume). An EC2 High Memory instance currently supports only up to 13 EBS volumes. If needed, consolidate some of the EBS volumes to reduce the total number of attached EBS volumes. For details, see Instance Type Limits in the AWS documentation.
5. When you are ready to migrate, make sure that you have a good backup of your source system.
6. Stop the source instance in the Amazon EC2 console or by using the AWS CLI.
7. If your source EC2 instance is running with dedicated tenancy, modify the instance placement to host tenancy. For instructions, see Modifying instance Tenancy and Affinity in the AWS documentation. Skip this step if your instance is running with host tenancy.
8. Modify the instance placement of your existing instance to your target EC2 High Memory Dedicated Host through the Amazon EC2 console or the AWS CLI. For details, see modify-instance-placement in the AWS documentation.
9. Change your instance type to the desired EC2 High Memory instance type (for example, `u-12tb1.metal`) through the AWS CLI.

   Note
   You can change the instance type to an EC2 High Memory instance only through the AWS CLI or Amazon EC2 API.
10. Start your instance in the Amazon EC2 console or by using the AWS CLI.
11. When you increase the memory of your SAP HANA system, you might need to adjust the storage size of SAP HANA data, log, shared, and backup volumes as well to accommodate data growth and to get improved performance. For details, see SAP HANA on AWS Operations Guide.
12. Start your SAP HANA database and perform your validation.
13. Complete any SAP HANA-specific post-migration activities.
14. Complete any AWS-specific post-migration activities, such as setting up Amazon CloudWatch, AWS Config, and AWS CloudTrail.
15. Configure your SAP HANA system for high availability on the EC2 High Memory instance with SAP HANA HSR and clustering software, and test it.

Option 2: Migrating from an Existing EC2 Instance with Default Tenancy

If your existing EC2 instance is running with default tenancy, you have two options to migrate it to an EC2 High Memory instance: You can use an Amazon Machine Image (AMI) to launch your EC2 High Memory instance with host tenancy, or you can set up a new SAP HANA on EC2 High Memory instance and then copy the data over from your source system.

Option 2(a): Migrating Using an AMI

In this option, you launch a new EC2 High Memory instance based on the AMI that you created from your source system for the migration.

Figure 11 provides a high-level illustration of this method.

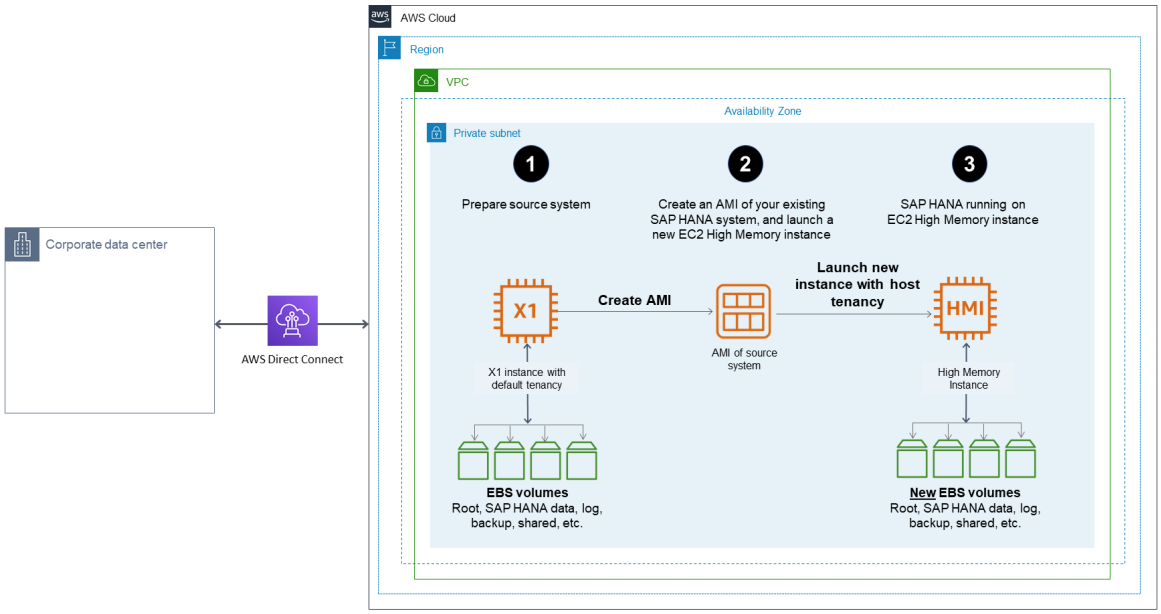

Figure 11: Migrating to an EC2 High Memory instance using an AMI

1. Make sure that your source system is running on a supported operating system version. If not, you might have to upgrade your operating system before resizing to an EC2 High Memory instance.
2. If you are using SUSE Linux, make sure that your source system is configured to load the required NVMe drivers during boot. This step is not required if you are using a SUSE image that has a date stamp later than 20180810. For details, see Changing instance Types in AWS For SLES 12 SP2 and SLES 12 SP3 Based HVM instances on the SUSE blog.
3. EC2 High Memory instances are based on the Nitro system. On Nitro-based instances, EBS volumes are presented as NVMe block devices. If your source system has any mount point entries in `/etc/fstab` with reference to block devices such as `/dev/xvd<x>`, you need to create a label for these devices and mount them by label before migrating to EC2 High Memory instances. Otherwise, you will face issues when you start SAP HANA on an EC2 High Memory instance.
4. When you are ready to migrate, make sure that you have a good backup of your source system.
5. Stop the source instance in the Amazon EC2 console or by using the AWS CLI.
6. Create an AMI of your source instance. For details, see Creating an Amazon EBS-Backed Linux AMI in the AWS documentation.

   Tip
   Creating an AMI for the first time with the attached EBS volumes could take a long time, depending on your data size. To expedite this process, we recommend that you take snapshots of EBS volumes attached to the instance ahead of time.
7. Launch a new EC2 High Memory instance.
8. The new instance will have a new IP address. Update all references to the IP address of the source system, including the `/etc/hosts` file for the operating system and DNS entries, to reflect the new IP address. The hostname and storage layout will remain the same as on the source system.
9. When you increase the memory of your SAP HANA system, you might need to adjust the storage size of SAP HANA data, log, shared, and backup volumes as well to accommodate data growth and to get improved performance. For details, see SAP HANA on AWS Operations Guide.
10. Start your SAP HANA database and perform your validation.

    Note
    You might notice that SAP HANA is slow when loading data into memory for the first time after you create your instance with an AMI. This is expected behavior when EBS volumes

associated with SAP HANA data are created from a snapshot. You will not experience the slowness after the initial hydration.

11. Complete any SAP HANA-specific post-migration activities.
12. Check connectivity between your SAP application servers and the new SAP HANA instance.
13. Complete any AWS-specific post-migration activities, such as setting up Amazon CloudWatch, AWS Config, and AWS CloudTrail.
14. Configure your SAP HANA system for high availability on the EC2 High Memory instance with SAP HANA HSR and clustering software, and test it.

Option 2(b): Migrating Using SAP HANA HSR or SAP HANA backup and restore

In this option, you launch a new EC2 High Memory instance, install and configure SAP HANA on the instance, and then copy the data over from your source system to complete the migration.

1. Launch a new SAP HANA EC2 High Memory instance. You can use the SAP HANA Quick Start to set up your instance automatically, or follow the SAP HANA Environment Setup on AWS guide to set up your instance manually. Make sure that you are using an operating system that supports EC2 High Memory instances.
2. Complete any AWS-specific post-migration activities, such as setting up Amazon CloudWatch, AWS Config, and AWS CloudTrail, ahead of time.
3. Migrate the data from your existing SAP HANA instance by using SAP HANA HSR or SAP HANA backup and restore tools.

    - If you plan to use SAP HANA HSR for data migration, configure HSR to move data from your source system to your target system. This is illustrated in Figure 12. For details, see the SAP HANA Administration Guide from SAP.

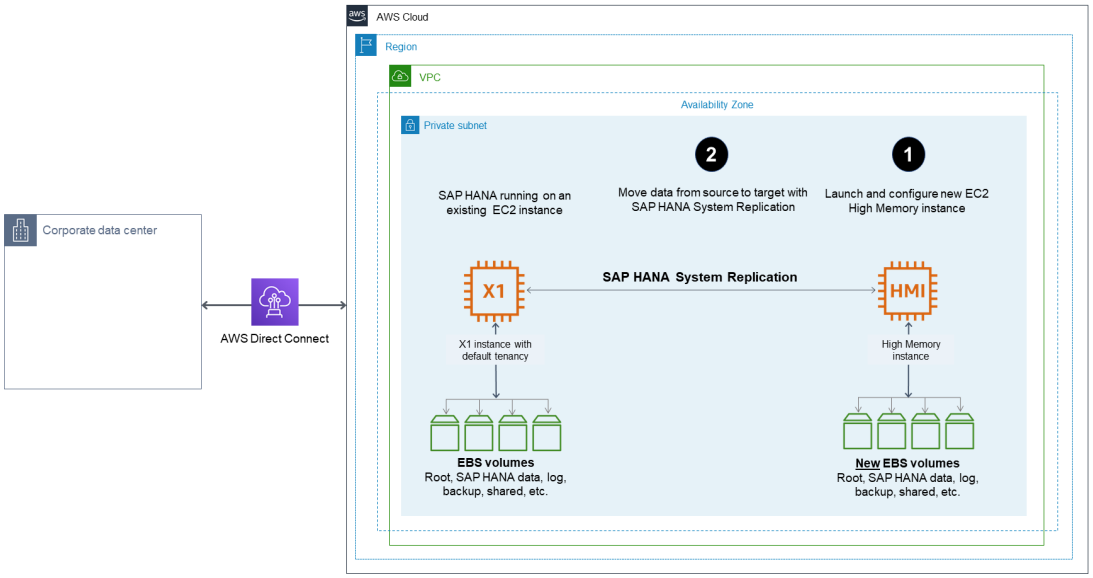

Figure 12: Migrating to an EC2 High Memory instance with HSR

- If you plan to use the SAP HANA backup and restore feature to migrate your data, back up your source SAP HANA system. When backup is complete, move the backup data to your target system

and perform a restore in your target system. If you back up your source SAP HANA system directly to Amazon S3 with `BACKINT`, you can directly restore it in the target system from Amazon S3. For details, see the `BACKINT` guide from your backup software vendor. This is illustrated in Figure 13.

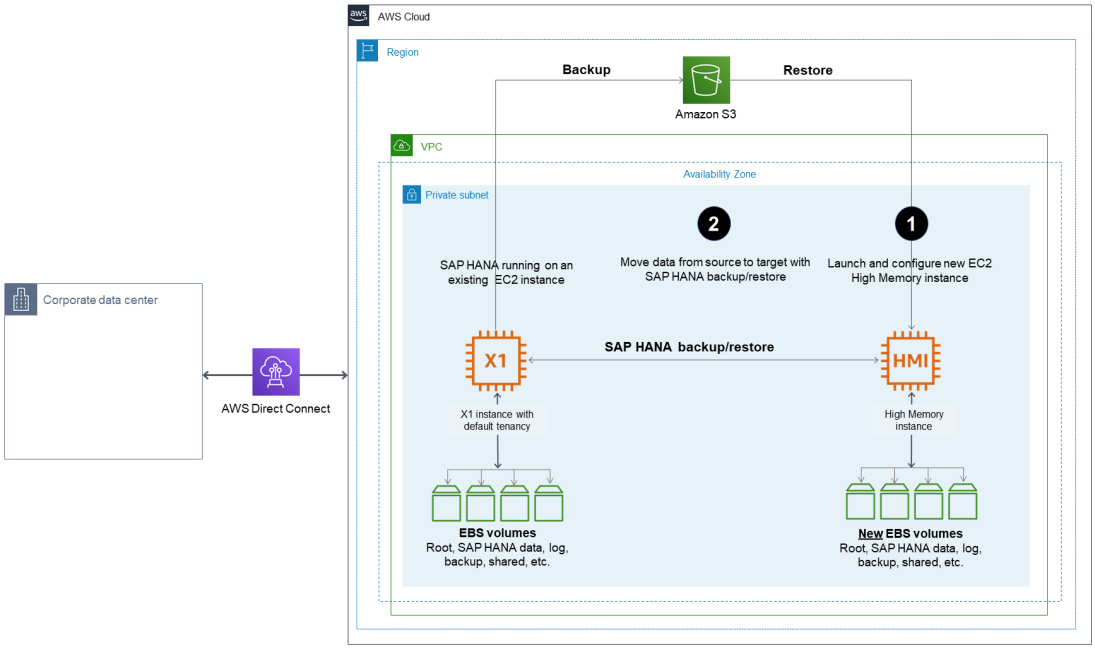

Figure 13: Migrating to an EC2 High Memory instance with SAP backup and restore

4. Stop your source system, complete any additional post-migration steps, like updating DNS and checking the connectivity between your SAP application servers and the new SAP HANA instance.
5. Configure your SAP HANA system for high availability on the EC2 High Memory instance with SAP HANA HSR and clustering software, and test it.

Third-Party Migration Tools

If you are interested in using the rehosting option (see the 6 Rs Framework (p. 2) section) for your on-premises SAP HANA environment, you can also leverage third-party continuous data protection (CDP) tools such as CloudEndure, Delphix, ATADATA, and Double-Take, which replicate the on-premises virtual machine, physical servers, and database on AWS. These tools provide an automated way to build your AWS environment, and migrate your source environment as is to AWS, including retaining host names and operating system configuration. These tools are application-agnostic and operate at the operating system and storage level, so they do not need to be SAP-certified for SAP migrations. There may be additional configuration steps needed to ensure that your SAP systems are running in the most optimized manner. For storage and instance requirements, see the Planning the Deployment section of the SAP HANA Quick Start deployment guide.

Security

In the AWS Cloud Adoption Framework (CAF), security is a perspective that focuses on subjects such as account governance, account ownership, control frameworks, change and access management, and other security best practices. We recommend that you become familiar with these security processes when planning any type of migration. In some cases, you might need to get sign-off from your internal IT audit

and security teams before you start your migration project or during migration. See the CAF security whitepaper for a deeper dive into each of these topic areas.

Additionally, there are AWS services that help you secure your systems in AWS. For example, AWS CloudTrail, Amazon CloudWatch, and AWS Config can help you secure your AWS environment.

See the following AWS blog posts for help analyzing and evaluating architectures and design patterns for the VPC setup and configuration of your SAP landscape.

- VPC Subnet Zoning Patterns for SAP on AWS, Part 1: Internal-Only Access
- VPC Subnet Zoning Patterns for SAP on AWS, Part 2: Network Zoning
- VPC Subnet Zoning Patterns for SAP on AWS, Part 3: Internal and External Access

Beyond VPC and network security, SAP HANA systems require routine maintenance to remain secure, reliable, and available; see the SAP HANA operations overview for specific recommendations in this topic area.

Additional Reading

- SAP FAST
- SAP HANA on the AWS Cloud: Quick Start Reference Deployment
- X1 Overview
- SAP and Amazon Web Services website
- SAP on AWS whitepapers
- AWS documentation

Document Revisions

Date	Change	Location
June 2019	Added new migration scenario for EC2 High Memory instances, and updated to reflect latest information	New: Amazon S3 Transfer Acceleration (p. 12), AMIs (p. 12), Migrating SAP HANA on AWS to an EC2 High Memory Instance (p. 18)
August 2018	Initial publication	–

SAP HANA Environment Setup on AWS

SAP specialists, Amazon Web Services (AWS)

Last updated: February 2019

This guide is part of a content series that provides detailed information about hosting, configuring, and using SAP technologies in the AWS Cloud. For the other guides in the series, ranging from overviews to advanced topics, see the SAP on AWS Technical Documentation home page.

This document provides guidance on how to set up AWS resources and configure SUSE Linux Enterprise Server (SLES) and Red Hat Enterprise Linux (RHEL) operating systems to deploy SAP HANA on Amazon Elastic Compute Cloud (Amazon EC2) instances in an existing virtual private cloud (VPC). It includes instructions for configuring storage for scale-up and scale-out workloads with Amazon Elastic Block Store (Amazon EBS) and Amazon Elastic File System (Amazon EFS).

This document follows AWS best practices to ensure that your system meets all key performance indicators (KPIs) that are required for Tailored Data Center Integration (TDI)–based SAP HANA implementations on AWS. In addition, this document also follows recommendations provided by SAP, SUSE, and Red Hat for SAP HANA in SAP OSS Notes 2205917, 1944799, 2292690 and 2009879. SAP regularly updates these OSS notes. Review the latest version of the OSS notes for up-to-date information before proceeding.

This guide is intended for users with a good understanding of AWS services, network concepts, the Linux operating system and SAP HANA administration to successfully launch and configure the resources that are required for SAP HANA.

AWS provides a Quick Start reference deployment for SAP HANA to fast-track your SAP HANA deployment on the AWS cloud. The Quick Start leverages AWS CloudFormation and scripts to quickly provision resources needed to deploy SAP HANA, and usually takes less than an hour to complete with minimal manual intervention. Refer to the SAP HANA on AWS Quick Start deployment guide if you want to use the automated deployment.

If your organization can't use the Quick Start reference deployment and you require additional customization to meet internal policies, you can follow the steps in this document to manually set up AWS resources such as Amazon EC2, Amazon EBS, Amazon EFS, by using the AWS Command Line Interface (AWS CLI) or the AWS Management Console.

Unlike the SAP HANA on AWS Quick Start, this document doesn't provide guidance on how to set up network and security constructs such as Amazon VPC, subnets, route tables, access control lists (ACLs), NAT Gateway, AWS Identity and Access Management (IAM) roles, security groups, etc. Instead, this document focuses on configuring compute, storage, and operating system resources for SAP HANA deployment on AWS.

Prerequisites

Specialized Knowledge

If you are new to AWS, see Getting Started with AWS.

Technical Requirements

1. If necessary, request a service limit increase for the instance type that you're planning to use for your SAP HANA system. If you already have an existing deployment that uses this instance type, and you think you might exceed the default limit with this deployment, you will need to request an increase. For details, see Amazon EC2 Service Limits in the AWS documentation.
2. Ensure that you have a key pair that you can use to launch your Amazon EC2 instance. If you need to create or import a key pair, refer to Amazon EC2 Key Pairs in the AWS documentation.
3. Ensure that you have the network details of the VPC, such as VPC ID and subnet ID, where you plan to launch the Amazon EC2 instance that will host SAP HANA.
4. Ensure that you have a security group to attach to the Amazon EC2 instance that will host SAP HANA and that the required ports are open. If needed, create a new security group that allows the traffic for SAP HANA ports. For a detailed list of ports, see Appendix C in the SAP HANA on AWS Quick Start guide.
5. If you intend to use AWS CLI to launch your instances, ensure that you have installed and configured AWS CLI with the necessary credentials. For details, see Installing the AWS Command Line Interface in the AWS documentation.
6. If you intend to use the console to launch your instances, ensure that you have credentials and permissions to launch and configure Amazon EC2, Amazon EBS, and other services. For details, see Access Management in the AWS documentation.

Architecture

This guide contains instructions for the following two environment setups:

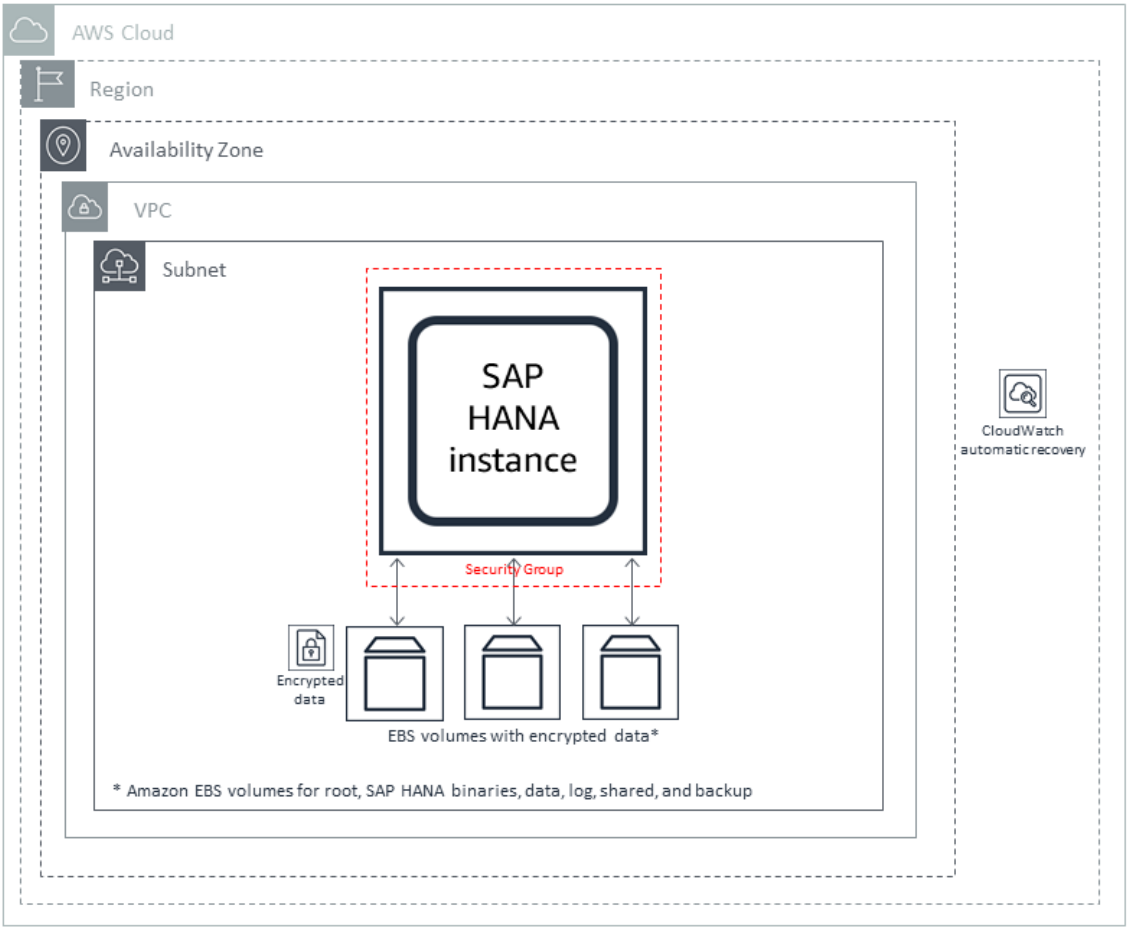

Figure 1: AWS configured for scale-up SAP HANA workloads

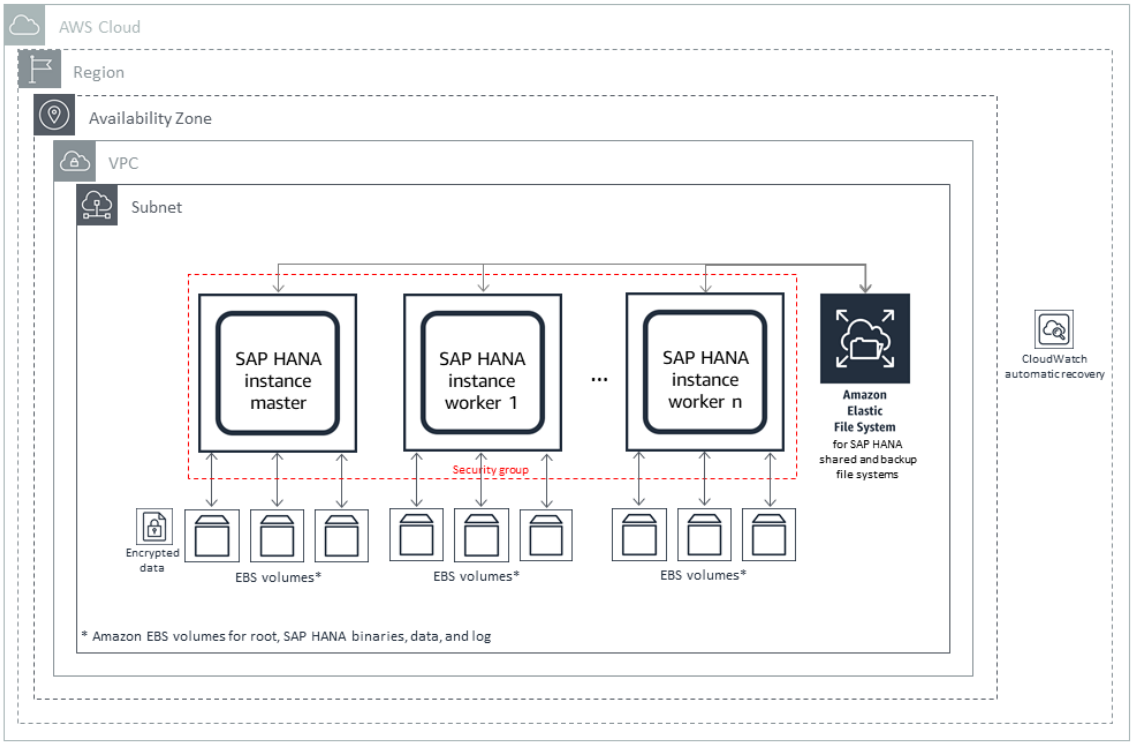

Figure 2: AWS configured for scale-out SAP HANA workloads

Planning the Deployment

Compute

AWS provides multiple instance families with different sizes to run SAP HANA workloads. See the SAP Certified and Supported SAP HANA Hardware Directory and the Amazon EC2 Instance Types for SAP page to find list of certified Amazon EC2 instances. For your production workloads, ensure that you choose an instance type that has been certified by SAP. You can run your non-production workloads on any size of a particular certified instance family to save costs.

Operating System

You can deploy your SAP HANA workload on SLES, SLES for SAP, RHEL for SAP with High Availability and Update Services (RHEL for SAP with HA and US), or RHEL for SAP Solutions.

SLES for SAP and RHEL for SAP with HA and US products are available in AWS Marketplace under an hourly or an annual subscription model.

SLES for SAP

SLES for SAP provides additional benefits, including Extended Service Pack Overlap Support (ESPOS), cofiguration and tuning packages for SAP applications, and High Availability Extenstions (HAE). For details, see the SUSE SLES for SAP product page to learn more about the benefits of using SLES for SAP. We strongly recommend using SLES for SAP instead of SLES for all your SAP workloads.

If you plan to use Bring Your Own Subscription (BYOS) images provided by SUSE, ensure that you have the registration code required to register your instance with SUSE to access repositories for software updates.

RHEL for SAP

RHEL for SAP with HA and US provides access to Red Hat Pacemaker cluster software for High Availability, extended update support, and the libraries that are required to run SAP HANA. For details, see the RHEL for SAP Offerings on AWS FAQ in the Red Hat Knowledgebase.

If you plan to use the BYOS model with RHEL, either through the Red Hat Cloud Access program or another means, ensure that you have access to a RHEL for SAP Solutions subscription. For details, see Overview of Red Hat Enterprise Linux for SAP Solutions subscription in the Red Hat Knowledgebase.

Amazon Machine Image (AMI)

A base AMI is required to launch an Amazon EC2 instance. Depending on your choice of operating system, ensure that you have access to the appropriate AMI in your target region for the deployment.

If you plan to use the SLES for SAP or RHEL for SAP Amazon Machine Images (AMIs) offered in AWS Marketplace, ensure that you have completed the subscription process. For details on how to subscribe to one of these AMIs, see the Appendix sections of the SAP HANA on AWS Quick Start deployment guide.

If you are using AWS CLI, you will need to provide the AMI ID when you launch the instance.

Storage

Deploying SAP HANA on AWS requires specific storage size and performance to ensure that SAP HANA data and log volumes both meet the SAP KPIs and sizing recommendations. Refer to the Planning the deployment section of the SAP HANA on AWS Quick Start guide to understand the storage configuration details for different instance types. You need to configure your storage based on these recommendations during instance launch.

Network

Ensure that your network constructs are set up to deploy resources related to SAP HANA. If you haven't already set up network components such as Amazon VPC, subnets, route table, etc., you can use the AWS Modular and Scalable VPC Quick Start to easily deploy a scalable VPC architecture in minutes. For details, see the deployment guide.

Deployment steps using AWS CLI

Step 1. Prepare Storage Configuration for SAP HANA

Use the editor of your choice to create a .json file that contains block device mapping details similar to the following example, and save your file in a temporary directory. The example shows the block device mapping details for the x1.32xlarge instance type with io1 volumes for HANA data and log. Change the details depending on instance and storage type that you intend to use for your deployment. For more information about the storage details for different instance types, see the Planning the deployment section of the SAP HANA on AWS Quick Start guide.

```
[
```

```
  {"DeviceName":"/dev/sda1","Ebs":
{"VolumeSize":50,"VolumeType":"gp2","DeleteOnTermination":false}},
  {"DeviceName":"/dev/sdb","Ebs":
{"VolumeSize":800,"VolumeType":"io1","Iops":3000,"Encrypted":true,"DeleteOnTermination":false}},
  {"DeviceName":"/dev/sdc","Ebs":
{"VolumeSize":800,"VolumeType":"io1","Iops":3000,"Encrypted":true,"DeleteOnTermination":false}},
  {"DeviceName":"/dev/sdd","Ebs":
{"VolumeSize":800,"VolumeType":"io1","Iops":3000,"Encrypted":true,"DeleteOnTermination":false}},
  {"DeviceName":"/dev/sde","Ebs":
{"VolumeSize":1024,"VolumeType":"gp2","Encrypted":true,"DeleteOnTermination":false}},
  {"DeviceName":"/dev/sdf","Ebs":
{"VolumeSize":4096,"VolumeType":"st1","Encrypted":true,"DeleteOnTermination":false}},
  {"DeviceName":"/dev/sdh","Ebs":
{"VolumeSize":525,"VolumeType":"io1","Iops":2000,"Encrypted":true,"DeleteOnTermination":false}},
  {"DeviceName":"/dev/sdr","Ebs":
{"VolumeSize":50,"VolumeType":"gp2","Encrypted":true,"DeleteOnTermination":false}}
  ]
```

Important
If the `DeleteOnTermination` flag is set to false, Amazon EBS volumes are not deleted when you terminate your Amazon EC2 instance. This helps preserve your data from accidental termination of your Amazon EC2 instance. When you terminate the instance, you need to manually delete the Amazon EBS volumes that are associated with the terminated instance to stop incurring storage cost.

See Appendix A (p. 41) for more examples of block device mappings for other Amazon EC2 instance types and Amazon EBS volume types.

Note
If you plan to deploy scale-out workloads, you don't have to include Amazon EBS volumes for SAP HANA shared and backup volumes. You can use Amazon EFS and Network File System (NFS) to mount the SAP HANA shared and backup volumes to your master and worker nodes.

Step 2. Launch the Amazon EC2 instance

Use AWS CLI to launch the Amazon EC2 instance for SAP HANA, including Amazon EBS storage, in the VPC in your target AWS Region by using the information you gathered during the preparation steps; for example:

Important
Be sure to enter the command on a single line.

```
$ aws ec2 run-instances
--image-id ami-xxxxxxxx
--count 1
--instance-type x1.32xlarge
--region us-west-2
--key-name=my_key
--security-group-ids sg-xxxxxxxx
--subnet-id subnet-xxxxxxxx
--placement GroupName=My-PlacementGroup,Tenancy=default,HostId=My-DedicatedHostId
--block-device-mappings file:///tmp/ebs_hana.json
--tag-specifications 'ResourceType=instance,Tags=[{Key=Name,Value=MyHANA}]'
'ResourceType=volume,Tags=[{Key=Name,Value=MyHANAVolumes}]'
```

Notes

- The `--placement` parameter is optional and needed only when you use a dedicated host with host tenancy or you want to place all your Amazon EC2 instances in close proximity. You may also pass additional parameters like `private-ip-address`, `disable-api-termination`, etc., as needed for your environment. For additional details, see run-instances in the AWS CLI Command Reference.

- After the instance and volumes are created, you can adjust the values of Amazon EBS volume tags to be more specific for ease of management. You can also add any additional tags that you need.
- For scale-out workloads, you can use the `--count` parameter to specify the total number of required nodes.
- Amazon EC2 High Memory Instances can be launched only through AWS CLI or APIs. After launch, however, you can manage them by using the console, AWS CLI, or APIs.

Deployment steps using the AWS Management Console

1. Log in to the console with appropriate permissions and ensure that you have the right Region selected.
2. Choose **Services**, and then choose **EC2** (under **Compute**).
3. Choose **Launch Instance**.
4. Search for the image that you want to use:

   - Choose **AWS Marketplace** to search for RHEL for SAP and SLES for SAP images.
   - Choose **My AMIs** to search for your BYOS or custom AMI ID.

   When you find the image, choose **Select**, and then confirm to continue.
5. On the **Choose an Instance Type** page, select the instance type that you identified when planning the deployment (p. 27), and choose **Configure Instance Details** to proceed with instance configuration.
6. On the **Configure Instance Details** page, do the following:

   a. Enter the number of instances (typically 1). For scale-out workloads, specify the number of nodes.
   b. Select the VPC ID and subnet for the network.
   c. Turn off the **Auto-assign Public IP** option.
   d. Select **Add instance to placement group** if needed (recommended for scale-out workloads; for details, see the AWS documentation).
   e. Select any IAM role that you want to assign to the instance to access AWS services from the instance.
   f. Select **Stop** for **Shutdown behavior**.
   g. Enable termination protection if needed (strongly recommended).
   h. Enable Amazon CloudWatch detailed monitoring (strongly recommended; for details, see the AWS documentation).
   i. Select the **Tenancy** or proceed with the default (**Shared**). For dedicated hosts, select the **Dedicated host** option.
   j. Choose **Add Storage** to proceed with storage configuration.
7. On the **Add Storage** page, choose **Add New Volume** to add volumes required for SAP HANA with the appropriate device, size, volume type, IOPS (for io1 only), and the **Delete on Termination** flag. Ensure that you follow the storage guidance (p. 28) discussed earlier in this document. Add volumes for SAP HANA data, log, shared, backup, and binaries.

   Figure 3 shows the storage configuration for x1.32xlarge instance type with io1 volume type for SAP HANA data and log.

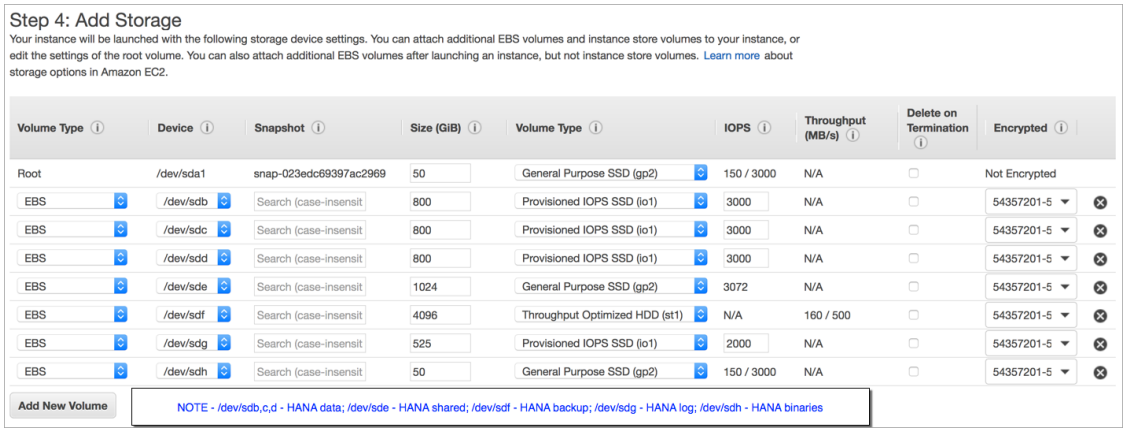

Figure 3: SAP HANA Storage Configuration with the console

> **Note**
> If you are planning to deploy scale-out workloads, you don't have to include Amazon EBS volumes for SAP HANA shared and backup volumes. You can use Amazon EFS with NFS to mount the HANA shared and backup volumes to your master and worker nodes.

Choose **Add Tags** to proceed with configuring tags

8. Choose **Add Tag** and add the key-value pair to track and manage your resources. We recommend adding `Name` as a minimum key to easily identify your resources.

   Next, choose **Configure Security Group**.

9. Choose **Select an existing security group** and select a security group, if you have one, to attach to your instance. Otherwise, choose **Create a new security group** and configure the **Type**, **Protocol**, **Port Range**, and the **Source IP address** from where you want to allow traffic to your SAP HANA instance. Refer to Appendix C of the SAP HANA on AWS Quick Start guide for a list of ports that we recommend. You can change the port as needed to meet your security requirements.
10. Choose **Review and Launch** to review your selections, and then choose **Launch**.
11. Select an existing key pair if you have one. Otherwise, create a new key pair, acknowledge it, and choose **Launch Instances**.
12. Your instance should be launching now with the selected configuration. After the instance is launched, you can proceed with the operating system and storage configuration steps.

> **Note**
> Amazon EBS volumes are presented as NVME block devices on Nitro-based instances. You need to perform additional mapping at the operating system level when you configure these volumes.

Operating System and Storage Configuration

Use the instructions for your operating system:

- SLES for SAP 12.x (p. 32)
- RHEL for SAP 7.x (p. 34)

> **Note**
> For scale-out workloads, repeat these steps for every node in the cluster.

Configure Operating System – SLES for SAP 12.x

Important
In the following steps, you need to update several configuration files. We recommend taking a backup of the files before you modify them. This will help you to revert to the previous configuration if needed.

1. After your instance is up and running, connect to the instance by using Secure Shell (SSH) and the key pair that you used to launch the instance.

   Note
   Depending on your network and security settings, you might have to first connect by using SSH to a bastion host before accessing your SAP HANA instance, or you might have to add IP addresses or ports to the security group to allow SSH access.

2. Switch to root user.

   Alternatively, you can use `sudo` to execute the following commands as ec2-user.

3. Set a hostname and fully qualified domain name (FQDN) for your instance by executing the `hostnamectl` command and updating the `/etc/hostname` file.

   ```
   # hostnamectl set-hostname --static your_hostname
   # echo your_hostname.example.com > /etc/hostname
   ```

   Open a new session to verify the hostname change.

4. Ensure that the `DHCLIENT_SET_HOSTNAME` parameter is set to **no** to prevent DHCP from changing the hostname during restart.

   ```
   # grep DHCLIENT_SET_HOSTNAME /etc/sysconfig/network/dhcp
   ```

5. Set the `preserve_hostname` parameter to true to ensure your hostname is preserved during restart.

   ```
   # sed -i '/preserve_hostname/ c\preserve_hostname: true' /etc/cloud/cloud.cfg
   ```

6. Add an entry to the `/etc/hosts` file with the new hostname and IP address.

   ```
   ip_address hostname.example.com hostname
   ```

7. If you are using a BYOS SLES for SAP image, register your instance with SUSE. Ensure that your subscription is for SLES for SAP.

   ```
   # SUSEConnect -r Your_Registration_Code
   # SUSEConnect -s
   ```

8. Ensure that the following packages are installed:

   `systemd`, `tuned`, `saptune`, `libgcc_s1`, `libstdc++6`, `cpupower`, `autofs`, `nvme-cli`

   You can use the `rpm` command to check whether a package is installed.

   ```
   # rpm -qi package_name
   ```

   You can then use the zypper install command to install the missing packages.

```
# zypper install package_name
```

Note
If you are importing your own SLES image, additional packages might be required to ensure that your instance is optimally setup. For the latest information, refer to the Package List section in the SLES for SAP Application Configuration Guide for SAP HANA, which is attached to SAP OSS Note 1944799

9. Ensure that your instance is running on a kernel version that is recommended in SAP OSS Note 2205917. If needed, update your system to meet the minimum kernel version. You can check the version of the kernel and other packages by using the following command:

```
# rpm -qi kernel*
```

10. Start `saptune daemon` and use the following command to set it to automatically start when the system reboots.

```
# saptune daemon start
```

11. Check whether the `force_latency` parameter is set in the `saptune` configuration file.

```
# grep force_latency /usr/lib/tuned/saptune/tuned.conf
```

If the parameter is set, skip the next step and proceed with activating the HANA profile with `saptune`.

12. Update the `saptune HANA` profile according to SAP OSS Note 2205917, and then run the following commands to create a custom profile for SAP HANA. This step is not required if the `force_latency` parameter is already set.

```
# mkdir /etc/tuned/saptune
# cp /usr/lib/tuned/saptune/tuned.conf /etc/tuned/saptune/tuned.conf
# sed -i "/\[cpu\]/ a force_latency=70" /etc/tuned/saptune/tuned.conf
# sed -i "s/script.sh/\/usr\/lib\/tuned\/saptune\/script.sh/"
```

13. Switch the `tuned` profile to HANA and verify that all settings are configured appropriately.

```
# saptune solution apply HANA
# saptune solution verify HANA
```

14. Configure and start the Network Time Protocol (NTP) service. You can adjust the NTP server pool based on your requirements; for example:

Note
Remove any existing invalid NTP server pools from `/etc/ntp.conf` before adding the following.

```
# echo "server 0.pool.ntp.org" >> /etc/ntp.conf
# echo "server 1.pool.ntp.org" >> /etc/ntp.conf
# echo "server 2.pool.ntp.org" >> /etc/ntp.conf
# echo "server 3.pool.ntp.org" >> /etc/ntp.conf
# systemctl enable ntpd.service
# systemctl start ntpd.service
```

Tip
Instead of connecting to the global NTP server pool, you can connect to your internal NTP server if needed. Or you can use Amazon Time Sync Service to keep your system time in sync.

15. Set the clocksource to `tsc` by updating the `current_clocksource` file and the GRUB2 boot loader.

```
# echo "tsc" > /sys/devices/system/clocksource/*/current_clocksource
# cp /etc/default/grub /etc/default/grub.backup
# sed -i '/GRUB_CMDLINE_LINUX/ s|"| clocksource=tsc"|2' /etc/default/grub
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

16. Reboot your system for the changes to take effect.
17. Continue with storage configuration for SAP HANA (p. 36).

Configure Operating System – RHEL for SAP 7.x

Important
In the following steps, you need to update several configuration files. We recommend taking a backup of the files before you modify them. This will help you to revert to the previous configuration if needed.

1. After your instance is up and running, connect to the instance by using Secure Shell (SSH) and the key pair that you used to launch the instance.

   Note
   Depending on your network and security settings, you might have to first connect by using SSH to a bastion host before accessing your SAP HANA instance, or you might have to add IP addresses or ports to the security group to allow SSH access.

2. Switch to root user.

   Alternatively, you can use sudo to execute the following commands as ec2-user.

3. Set a hostname for your instance by executing the `hostnamectl` command and update the `/etc/cloud/cloud.cfg` file to ensure that your hostname is preserved during system reboots.

```
# hostnamectl set-hostname --static your_hostname
# echo "preserve_hostname: true" >> /etc/cloud/cloud.cfg
```

   Open a new session to verify the hostname change.

4. Add an entry to the `/etc/hosts` file with the new hostname and IP address.

```
ip address hostname.example.com hostname
```

   Ensure that the following packages are installed:

   `xfsprogs`, `gcc`, `compat-sap-c++-5`, `compat-sap-c++-6`, `tuned-profiles-sap-hana`, `glibc.x86_64`, `autofs`, and `nvme-cli`

   Note that your instance should have access to the SAP HANA channel to install libraries requires for SAP HANA installations.

   You can use the `rpm` command to check whether a package is installed:

```
# rpm -qi package_name
```

You can then install any missing packages by using the `yum -y install` command.

```
# yum -y install package name
```

Note
Depending on your base RHEL image, additional packages might be required to ensure that your instance is optimally setup. (You can skip this step if you are using the RHEL for SAP with HA & US image.) For the latest information, refer to the RHEL configuration guide that is attached to SAP OSS Note 2009879. Review the packages in the Install Additional Required Packages section and the Appendix–Required Packages for SAP HANA on RHEL 7 section.

5. Ensure that your instance is running on a kernel version that is recommended in SAP OSS Note 2292690. If needed, update your system to meet the minimum kernel version. You can check the version of the kernel and other packages using the following command.

```
# rpm -qi kernel*
```

6. Start `tuned daemon` and use the following commands to set it to automatically start when the system reboots.

```
# systemctl start tuned

# systemctl enable tuned
```

7. Configure the `tuned HANA` profile to optimize your instance for SAP HANA workloads.

   Check whether the `force_latency` parameter is already set in the `/usr/lib/tuned/sap-hana/tuned.conf` file. If the parameter is set, execute the following commands to apply and activate the `sap-hana` profile.

```
# tuned-adm profile sap-hana
# tuned-adm active
```

   If the `force_latency` parameter is not set, execute the following steps to modify and activate the `sap-hana` profile.

```
# mkdir /etc/tuned/sap-hana
# cp /usr/lib/tuned/sap-hana/tuned.conf /etc/tuned/sap-hana/tuned.conf
# sed -i '/force_latency/ c\force_latency=70' /etc/tuned/sap-hana/tuned.conf
# tuned-adm profile sap-hana
# tuned-adm active
```

8. Disable Security-Enhanced Linux (SELinux) by running the following command. (Skip this step if you are using the RHEL for SAP with HA & US image.)

```
# sed -i 's/\(SELINUX=enforcing\|SELINUX=permissive\)/SELINUX=disabled/g' \/etc/
selinux/config
```

9. Disable Transparent Hugepages (THP) at boot time by adding the following to the line that starts with GRUB_CMDLINE_LINUX in the `/etc/default/grub` file. Execute the following commands to add the required parameter and to re-configure grub (Skip this step if you are using the RHEL for SAP with HA & US image.)

```
# sed -i '/GRUB_CMDLINE_LINUX/ s|"| transparent_hugepage=never"|2' /etc/default/grub
# cat /etc/default/grub
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

10. Add symbolic links by executing following commands. (Skip this step if you are using the RHEL for SAP with HA & US image.)

```
# ln -s /usr/lib64/libssl.so.10 /usr/lib64/libssl.so.1.0.1
# ln -s /usr/lib64/libcrypto.so.10 /usr/lib64/libcrypto.so.1.0.1
```

11. Configure and start the Network Time Protocol (NTP) service. You can adjust the NTP server pool based on your requirements. The following is just an example.

    Note
    Remove any existing invalid NTP server pools from `/etc/ntp.conf` before adding the following.

```
# echo "server 0.pool.ntp.org" >> /etc/ntp.conf
# echo "server 1.pool.ntp.org" >> /etc/ntp.conf
# echo "server 2.pool.ntp.org" >> /etc/ntp.conf
# echo "server 3.pool.ntp.org" >> /etc/ntp.conf
# systemctl enable ntpd.service
# systemctl start ntpd.service
# systemctl restart systemd-timedated.service
```

    Tip
    Instead of connecting to the global NTP server pool, you can connect to your internal NTP server if needed. Alternatively, you can also use Amazon Time Sync Service to keep your system time in sync.

12. Set clocksource to `tsc` by the updating the `current_clocksource` file and the GRUB2 boot loader.

```
# echo "tsc" > /sys/devices/system/clocksource/*/current_clocksource
# cp /etc/default/grub /etc/default/grub.backup
# sed -i '/GRUB_CMDLINE_LINUX/ s|"| clocksource=tsc"|2' /etc/default/grub
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

13. Reboot your system for the changes to take effect.
14. After the reboot, log in as root and execute the `tuned-adm` command to verify that all SAP recommended settings are in place.

```
# tuned-adm verify

“tuned-adm verify” creates a log file under /var/log/tuned/tuned.log Review this log
file and ensure that  all checks have passed.
```

15. Continue with storage configuration.

Configure Storage for SAP HANA

1. Amazon EBS volumes should have been created and attached when you launched the Amazon EC2 instance. Confirm that all the required volumes are attached to the instance by running the `lsblk` command, which returns a list of the storage devices that are attached to the instance.

Note
On Nitro-based instances, Amazon EBS volumes are presented as NVME block devices. You need to perform additional mapping when configuring these volumes.

Depending on the instance and storage volume types, your block device mapping will look similar to the following examples.

Example from a non-Nitro instance

```
# lsblk
NAME     MAJ:MIN  RM  SIZE RO TYPE MOUNTPOINT
xvda     202:0     0   50G  0 disk
##xvda1 202:1     0    1M  0 part
##xvda2 202:2     0   50G  0 part /
xvdb     202:16    0  800G  0 disk
xvdc     202:32    0  800G  0 disk
xvdd     202:48    0  800G  0 disk
xvde     202:64    0    1T  0 disk
xvdf     202:80    0    4T  0 disk
xvdh     202:112   0  525G  0 disk
xvdr     202:4352  0   50G  0 disk
#
```

Example from a Nitro instance

```
## lsblk
NAME          MAJ:MIN  RM  SIZE RO TYPE MOUNTPOINT
nvme0n1       259:0     0   50G 0  disk
##nvme0n1p1  259:1     0   50G 0  part /
nvme1n1       259:2     0    4T 0  disk
nvme2n1       259:3     0  800G 0  disk
nvme3n1       259:4     0  800G 0  disk
nvme4n1       259:5     0  800G 0  disk
nvme5n1       259:6     0  525G 0  disk
nvme6n1       259:7     0    1T 0  disk
nvme7n1       259:8     0   50G 0  disk
#
```

2. Initialize the volumes of SAP HANA data, log, and backup to use with Linux Logical Volume Manager (LVM).

   Note
   Ensure you are choosing the devices that are associated with the SAP HANA data, log, and backup volumes. The device names might be different in your environment.

   Example from a non-Nitro instance

```
# pvcreate /dev/xvdb /dev/xvdc /dev/xvdd /dev/xvdf /dev/xvdh
  Physical volume "/dev/xvdb" successfully created.
  Physical volume "/dev/xvdc" successfully created.
  Physical volume "/dev/xvdd" successfully created.
  Physical volume "/dev/xvdf" successfully created.
  Physical volume "/dev/xvdh" successfully created.
#
```

   Example from a Nitro instance

```
# pvcreate /dev/nvme2n1 /dev/nvme3n1 /dev/nvme4n1 /dev/nvme5n1 /dev/nvme1n1
  Physical volume "/dev/nvme2n1" successfully created.
  Physical volume "/dev/nvme3n1" successfully created.
```

```
  Physical volume "/dev/nvme4n1" successfully created.
  Physical volume "/dev/nvme5n1" successfully created.
  Physical volume "/dev/nvme1n1" successfully created.
#
```

3. Create volume groups for SAP HANA data, log, and backup. Ensure that device IDs are associated correctly with the appropriate volume group.

   Example from a non-Nitro instance

```
# vgcreate vghanadata /dev/xvdb /dev/xvdc /dev/xvdd
  Volume group "vghanadata" successfully created
# vgcreate vghanalog /dev/xvdh
  Volume group "vghanalog" successfully created
# vgcreate vghanaback /dev/xvdf
  Volume group "vghanaback" successfully created
#
```

   Example from a Nitro instance

```
# vgcreate vghanadata /dev/nvme2n1 /dev/nvme3n1 /dev/nvme4n1
  Volume group "vghanadata" successfully created
# vgcreate vghanalog /dev/nvme5n1
  Volume group "vghanalog" successfully created
# vgcreate vghanaback /dev/nvme1n1
  Volume group "vghanaback" successfully created
#
```

4. Create a logical volume for SAP HANA data.

   In the following command, `-i 3` represents stripes based on the number of volumes that are used for a HANA data volume group. Adjust the number depending on the number of volumes that are allocated to the HANA data volume group, based on instance and storage type.

```
# lvcreate -n lvhanadata -i 3 -I 256 -L 2350G vghanadata
  Rounding size 2.29 TiB (601600 extents) up to stripe boundary size 2.29 TiB (601602
extents).
  Logical volume "lvhanadata" created.
#
```

5. Create a logical volume for SAP HANA log.

   In the following command, `-i 1` represents stripes based on the number of volumes that are used for a HANA log volume group. Adjust the number depending on the number of volumes that are allocated to the HANA log volume group, based on instance and storage type.

```
# lvcreate -n lvhanalog -i 1 -I 256 -L 512G vghanalog
  Ignoring stripesize argument with single stripe.
  Logical volume "lvhanalog" created.
#
```

6. Create a logical volume for SAP HANA backup.

```
# lvcreate -n lvhanaback -i 1 -I 256 -L 4095G vghanaback
  Ignoring stripesize argument with single stripe.
  Logical volume "lvhanaback" created.
#
```

7. Construct XFS file systems with the newly created logical volumes for HANA data, log, and backup by using the following commands:

```
# mkfs.xfs -f /dev/mapper/vghanadata-lvhanadata
# mkfs.xfs -f /dev/mapper/vghanalog-lvhanalog
# mkfs.xfs -f /dev/mapper/vghanaback-lvhanaback
```

8. Construct XFS file systems for HANA shared and HANA binaries.

```
# mkfs.xfs -f /dev/xvde -L HANA_SHARE
# mkfs.xfs -f /dev/xvdr -L USR_SAP
```

Note
On Nitro-based instance types, device names can change during instance restarts. To prevent file system mount issues, it is important to create labels for devices that aren't part of logical volumes so that the devices can be mounted by using labels instead of the actual device names.

9. Create directories for HANA data, log, backup, shared, and binaries.

```
# mkdir /hana /hana/data /hana/log /hana/shared /backup /usr/sap
```

10. Use the `echo` command to add entries to the `/etc/fstab` file with the following mount options to automatically mount these file systems during restart.

```
 # echo "/dev/mapper/vghanadata-lvhanadata /hana/data xfs
nobarrier,noatime,nodiratime,logbsize=256k 0 0" >> /etc/fstab
 # echo "/dev/mapper/vghanalog-lvhanalog /hana/log xfs
nobarrier,noatime,nodiratime,logbsize=256k 0 0" >> /etc/fstab
 # echo "/dev/mapper/vghanaback-lvhanaback /backup xfs
nobarrier,noatime,nodiratime,logbsize=256k 0 0" >> /etc/fstab
 # echo "/dev/disk/by-label/HANA_SHARE /hana/shared xfs
nobarrier,noatime,nodiratime,logbsize=256k 0 0" >> /etc/fstab
 # echo "/dev/disk/by-label/USR_SAP /usr/sap xfs
nobarrier,noatime,nodiratime,logbsize=256k 0 0" >> /etc/fstab
```

11. Mount the file systems.

```
# mount -a
```

12. Check to make sure that all file systems are mounted appropriately; for example, here is the output from an x1.32xlarge system:

```
# df -h
Filesystem                          Size  Used Avail Use% Mounted on
/dev/xvda2                           50G  1.8G  49G    4% /
devtmpfs                            961G     0 961G    0% /dev
tmpfs                               960G     0 960G    0% /dev/shm
tmpfs                               960G   17M 960G    1% /run
tmpfs                               960G     0 960G    0% /sys/fs/cgroup
tmpfs                               192G     0 192G    0% /run/user/1000
/dev/mapper/vghanadata-lvhanadata 2.3T   34M 2.3T    1% /hana/data
/dev/mapper/vghanalog-lvhanalog   512G   33M 512G    1% /hana/log
/dev/mapper/vghanaback-lvhanaback 4.0T   33M 4.0T    1% /backup
/dev/xvde                           1.0T   33M 1.0T    1% /hana/shared
/dev/xvdr                            50G   33M 50G     1% /usr/sap
#
```

13. At this time, we recommend rebooting the system and confirming that all the file systems mount automatically after the restart.
14. If you are deploying a scale-out workload, follow the steps specified in Configure NFS for scale-out workloads (p. 40) to set up SAP HANA shared and backup NFS file systems with Amazon EFS.

If you are not deploying a scale-out workload, you can now proceed with your SAP HANA software installation.

Configure NFS for scale-out workloads

Amazon EFS provides easy-to-set-up, scalable, and highly available shared file systems that can be mounted with the NFSv4 client. For scale-out workloads, we recommend using Amazon EFS for SAP HANA shared and backup volumes. You can choose between different performance options for your file systems depending on your requirements. We recommend starting with the General Purpose and Provisioned Throughput options, with approximately 100 MiB/s to 200 MiB/s throughput. To set up your file systems, do the following:

1. Install the `nfs-utils` package in all the nodes in your scale-out cluster.
   - For RHEL, use `yum install nfs-utils`.
   - For SLES, use `zypper install nfs-utils`.
2. Create two Amazon EFS file systems and target mounts for SAP HANA shared and backup in your target VPC and subnet. For detailed steps, follow the instructions specified in the AWS documentation.
3. After the file systems are created, mount the newly created file systems in all the nodes by using the following commands:

```
  # mount -t nfs -o
 nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 EFS DNS Name:/ /hana/
shared

  # mount -t nfs -o
 nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 EFS DNS Name:/ /
backup
```

   Note
   If you have trouble mounting the NFS file systems, you might need to adjust your security groups to allow access to port 2049. For details, see Security Groups for Amazon EC2 Instances and Mount Targets in the AWS documentation.

4. Add NFS mount entries to the `/etc/fstab` file in all the nodes to automatically mount these file systems during system restart; for example:

```
  # echo “nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 EFS DNS
 Name:/ /hana/shared” >> /etc/fstab
  # echo “nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2 EFS DNS
 Name:/ /backup” >> /etc/fstab
```

5. Set appropriate permissions and ownership for your target mount points.

Post Deployment Steps

1. Complete the steps required to connect your instance to your corporate directory service, such as Microsoft Active Directory, if needed.
2. Set up any monitoring required for your environment.

3. Set up a CloudWatch alarm and Amazon EC2 automatic recovery to automatically recover your instance from hardware failures. For details, see Recover Your Instance in the AWS documentation. You can also refer to the Knowledge Center video for detailed instructions.

   Note
   Automatic recovery is not supported for Amazon EC2 instances running in dedicated hosts.
4. Create an AMI of your newly deployed system to take a full backup of your instance. For details, see Create an AMI from an Amazon EC2 Instance in the AWS documentation.
5. If you have deployed an SAP HANA scale-out cluster, consider adding additional elastic network interfaces and security groups to logically separate network traffic for client, inter-node, and optional SAP HANA System Replication (HSR) communications. For details, see the SAP HANA Operations Guide.

Additional Reading

AWS services

- Amazon EC2
- Amazon EBS
- Amazon VPC
- Amazon EFS

SAP document reference

- SAP OSS Note 2292690 - SAP HANA DB: Recommended OS settings for RHEL 7
- SAP OSS Note 2009879 - SAP HANA Guidelines for Red Hat Enterprise Linux (RHEL) Operating System
- SAP OSS Note 2205917 - SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12
- SAP OSS Note 1944799 - SAP HANA Guidelines for SLES Operating System Installation

Appendix A: Sample Block Device Mapping Configuration

Following are two block device mapping examples for your reference. You can find details of the recommended storage configuration for different types in the SAP HANA on AWS Quick Start deployment guide.

Example with x1.16xlarge instance type, GP2 storage type

```
  [
  {"DeviceName":"/dev/sda1","Ebs":
{"VolumeSize":50,"VolumeType":"gp2","DeleteOnTermination":false}},
  {"DeviceName":"/dev/sdb","Ebs":
{"VolumeSize":400,"VolumeType":"gp2,"Encrypted":true,"DeleteOnTermination":false}},
  {"DeviceName":"/dev/sdc","Ebs":
{"VolumeSize":400,"VolumeType":"gp2","Encrypted":true,"DeleteOnTermination":false}},
  {"DeviceName":"/dev/sdd","Ebs":
{"VolumeSize":400,"VolumeType":"gp2", ,"Encrypted":true,"DeleteOnTermination":false}},
```

```
  {"DeviceName":"/dev/sde","Ebs":
{"VolumeSize":1024,"VolumeType":"gp2","Encrypted":true,"DeleteOnTermination":false}},
  {"DeviceName":"/dev/sdf","Ebs":
{"VolumeSize":2048,"VolumeType":"st1","Encrypted":true,"DeleteOnTermination":false}},
  {"DeviceName":"/dev/sdh","Ebs":{"VolumeSize":300,"VolumeType":"gp2",
 "Encrypted":true,"DeleteOnTermination":false}},
  {"DeviceName":"/dev/sdh","Ebs":
{"VolumeSize":300,"VolumeType":"gp2","Encrypted":true,"DeleteOnTermination":false}},
  {"DeviceName":"/dev/sdr","Ebs":
{"VolumeSize":50,"VolumeType":"gp2","Encrypted":true,"DeleteOnTermination":false}}
  ]
```

Example with r4.16xlarge instance type, io1 storage type

```
  [
  {"DeviceName":"/dev/sda1","Ebs":
{"VolumeSize":50,"VolumeType":"gp2","DeleteOnTermination":false}},
  {"DeviceName":"/dev/sdb","Ebs":
{"VolumeSize":600,"VolumeType":"io1","Iops":7500,"Encrypted":true,"DeleteOnTermination":false}},
  {"DeviceName":"/dev/sde","Ebs":
{"VolumeSize":512,"VolumeType":"gp2","Encrypted":true,"DeleteOnTermination":false}},
  {"DeviceName":"/dev/sdf","Ebs":
{"VolumeSize":1024,"VolumeType":"st1","Encrypted":true,"DeleteOnTermination":false}},
  {"DeviceName":"/dev/sdh","Ebs":
{"VolumeSize":260,"VolumeType":"io1","Iops":2000,"Encrypted":true,"DeleteOnTermination":false}},
  {"DeviceName":"/dev/sdr","Ebs":
{"VolumeSize":50,"VolumeType":"gp2","Encrypted":true,"DeleteOnTermination":false}}
  ]
```

Document Revisions

Date	Change	In sections
February 2019	Initial publication	—

SAP HANA on AWS Operations Guide

SAP specialists, Amazon Web Services (AWS)

Last updated: December 2017

Amazon Web Services (AWS) offers you the ability to run your SAP HANA systems of various sizes and operating systems. Running SAP systems on AWS is very similar to running SAP systems in your data center. To a SAP Basis or NetWeaver administrator, there are minimal differences between the two environments. There are a number of AWS Cloud considerations relating to security, storage, compute configurations, management, and monitoring that will help you get the most out of your SAP HANA implementation on AWS.

This technical article provides the best practices for deployment, operations, and management of SAP HANA systems on AWS. The target audience is SAP Basis and NetWeaver administrators who have experience running SAP HANA systems in an on-premises environment and want to run their SAP HANA systems on AWS.

> **Note**
> The SAP notes and Knowledge Base articles (KBA) referenced in this guide require an SAP ONE Support Launchpad user account. For more information, see the SAP Support website.

About this Guide

This guide is part of a content series that provides detailed information about hosting, configuring, and using SAP technologies in the AWS Cloud. For the other guides in the series, ranging from overviews to advanced topics, see the SAP on AWS Technical Documentation home page.

Introduction

This guide provides best practices for operating SAP HANA systems that have been deployed on AWS either by using the SAP HANA Quick Start reference deployment process or by manually following the instructions in Setting up AWS Resources and the SLES Operating System for SAP HANA Installation. This guide is not intended to replace any of the standard SAP documentation. See the following SAP guides and notes:

- SAP Library (help.sap.com) - SAP HANA Administration Guide
- SAP installation guides (these require an SAP One Support Launchpad user account)
- SAP notes (these require an SAP One Support Launchpad user account)

This guide assumes that you have a basic knowledge of AWS. If you are new to AWS, see the following on the AWS website before continuing:

- AWS Getting Started Resource Center
- What is Amazon EC2?

In addition, see the following SAP on AWS guides:

- SAP on AWS Implementation and Operations Guide provides best practices for achieving optimal performance, availability, and reliability, and lower total cost of ownership (TCO) while running SAP solutions on AWS.

- SAP on AWS High Availability Guide explains how to configure SAP systems on Amazon Elastic Compute Cloud (Amazon EC2) to protect your application from various single points of failure.
- SAP on AWS Backup and Recovery Guide explains how to back up SAP systems running on AWS, in contrast to backing up SAP systems on traditional infrastructure.

Administration

This section provides guidance on common administrative tasks required to operate an SAP HANA system, including information about starting, stopping, and cloning systems.

Starting and Stopping EC2 Instances Running SAP HANA Hosts

At any time, you can stop one or multiple SAP HANA hosts. Before stopping the EC2 instance of an SAP HANA host, first stop SAP HANA on that instance.

When you resume the instance, it will automatically start with the same IP address, network, and storage configuration as before. You also have the option of using the EC2 Scheduler to schedule starts and stops of your EC2 instances. The EC2 Scheduler relies on the native shutdown and start-up mechanisms of the operating system. These native mechanisms will invoke the orderly shutdown and startup of your SAP HANA instance. Here is an architectural diagram of how the EC2 Scheduler works:

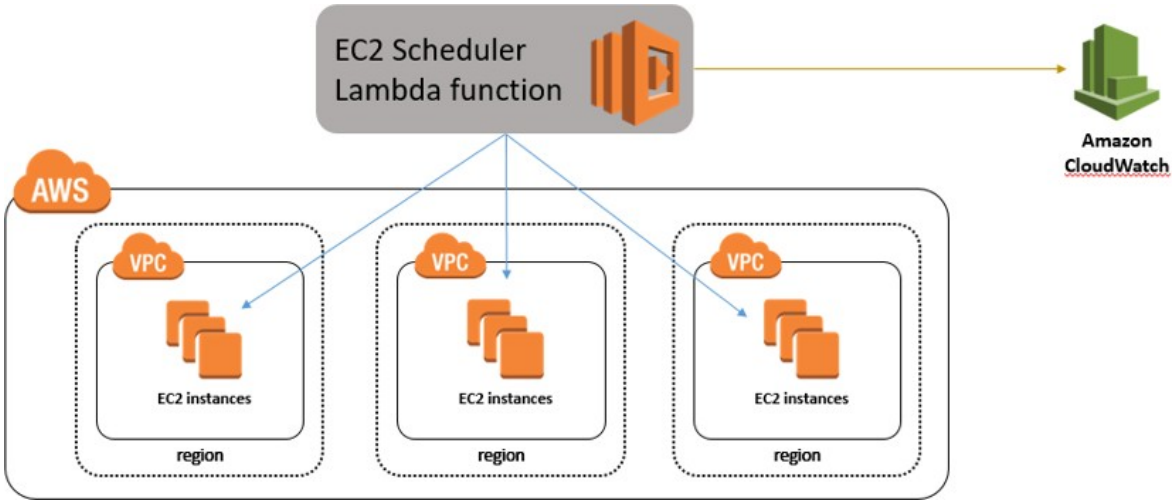

Figure 1: EC2 Scheduler

Tagging SAP Resources on AWS

Tagging your SAP resources on AWS can significantly simplify identification, security, manageability, and billing of those resources. You can tag your resources using the AWS Management Console or by using the `create-tags` functionality of the AWS Command Line Interface (AWS CLI). This table lists some example tag names and tag values:

Tag name	Tag value
Name	SAP server's virtual (host) name
Environment	SAP server's landscape role; for example: SBX, DEV, QAT, STG, PRD.

Tag name	Tag value
Application	SAP solution or product; for example: ECC, CRM, BW, PI, SCM, SRM, EP
Owner	SAP point of contact
Service level	Known uptime and downtime schedule

After you have tagged your resources, you can apply specific security restrictions such as access control, based on the tag values. Here is an example of such a policy from the AWS Security blog:

```
{
   "Version" : "2012-10-17",
   "Statement" : [
      {
         "Sid" : "LaunchEC2Instances", "Effect" : "Allow",
         "Action" : [
            "ec2:Describe*", "ec2:RunInstances"
         ],
         "Resource" : [
            "*"
         ]
      },
      {
         "Sid" : "AllowActionsIfYouAreTheOwner",
         "Effect" : "Allow",
         "Action" : [
            "ec2:StopInstances",
            "ec2:StartInstances",
            "ec2:RebootInstances",
            "ec2:TerminateInstances"
         ],
         "Condition" : {
            "StringEquals" : {
               "ec2:ResourceTag/PrincipalId" : "${aws:userid}"
            }
         },
         "Resource" : [
            "*"
         ]
      }
   ]
}
```

The AWS Identity and Access Management (IAM) policy allows only specific permissions based on the tag value. In this scenario, the current user ID must match the tag value in order for the user to be granted permissions. For more information on tagging, see the AWS documentation and AWS blog.

Monitoring

You can use various AWS, SAP, and third-party solutions to monitor your SAP workloads. Here are some of the core AWS monitoring services:

- Amazon CloudWatch – CloudWatch is a monitoring service for AWS resources. It's critical for SAP workloads where it's used to collect resource utilization logs and to create alarms to automatically react to changes in AWS resources.
- AWS CloudTrail – CloudTrail keeps track of all API calls made within your AWS account. It captures key metrics about the API calls and can be useful for automating trail creation for your SAP resources.

Configuring CloudWatch detailed monitoring for SAP resources is mandatory for getting AWS and SAP support. You can use native AWS monitoring services in a complementary fashion with the SAP Solution Manager. You can find third-party monitoring tools in AWS Marketplace.

Automation

AWS offers multiple options for programmatically scripting your resources to operate or scale them in a predictable and repeatable manner. You can use AWS CloudFormation to automate and operate SAP systems on AWS. Here are some examples for automating your SAP environment on AWS:

Area	Activities	AWS services
Infrastructure deployment	Provision new SAP environment SAP system cloning	AWS CloudFormation AWS CLI
Capacity management	Automate scale-up/scale-out of SAP application servers	AWS Lambda AWS CloudFormation
Operations	SAP backup automation (see the backup (p. 51) example (p. 51)) Performing monitoring and visualization	Amazon CloudWatch AWS Systems Manager

Patching

There are two ways for you to patch your SAP HANA database, with options for minimizing cost and/or downtime. With AWS, you can provision additional servers as needed to minimize downtime for patching in a cost-effective manner. You can also minimize risks by creating on-demand copies of your existing production SAP HANA databases for lifelike production readiness testing.

This table summarizes the tradeoffs of the two patching methods:

Patching method	Benefits	Technologies available
Patch an existing server	[x] Patch existing OS and DB [x] Longest downtime to existing server and DB [#] No costs for additional on-demand instances [✓] Lowest levels of relative complexity and setup tasks involved	Native OS patching tools Patch Manager Native SAP HANA patching tools
Provision and patch a new server	[✓] Leverage latest AMIs (only DB patch needed) [✓] Shortest downtime to existing server and DB [✓] Can patch and test OS and DB separately and together	Amazon Machine Image (AMI) AWS CLI AWS CloudFormation

Patching method	Benefits	Technologies available
	[x] More costs for additional on-demand instances [x] More complexity and setup tasks involved	SAP HANA System Replication SAP HANA System Cloning SAP HANA backups SAP Notes: 1984882 - Using HANA System Replication for Hardware Exchange with minimum/zero downtime 1913302 - HANA: Suspend DB connections for short maintenance tasks

The first method (patch an existing server) involves patching the operating system (OS) and database (DB) components of your SAP HANA server. The goal of this method is to minimize any additional server costs and to avoid any tasks needed to set up additional systems or tests. This method may be most appropriate if you have a well-defined patching process and are satisfied with your current downtime and costs. With this method you must use the correct operating system (OS) update process and tools for your Linux distribution. See this SUSE blog and Red Hat FAQ, or check each vendor's documentation for their specific processes and procedures.

In addition to patching tools provided by our Linux partners, AWS offers a free of charge patching service called Patch Manager. Patch Manager is an automated tool that helps you simplify your OS patching process. You can scan your EC2 instances for missing patches and automatically install them, select the timing for patch rollouts, control instance reboots, and many other tasks. You can also define auto-approval rules for patches with an added ability to black-list or white-list specific patches, control how the patches are deployed on the target instances (e.g., stop services before applying the patch), and schedule the automatic rollout through maintenance windows.

The second method (provision and patch a new server) involves provisioning a new EC2 instance that will receive a copy of your source system and database. The goal of the method is to minimize downtime, minimize risks (by having production data and executing production-like testing), and have repeatable processes. This method may be most appropriate if you are looking for higher degrees of automation to enable these goals and are comfortable with the trade- offs. This method is more complex and has a many more options to fit your requirements. Certain options are not exclusive and can be used together. For example, your AWS CloudFormation template can include the latest Amazon Machine Images (AMIs), which you can then use to automate the provisioning, set up, and configuration of a new SAP HANA server.

Here is an example of a process that can be used to automate OS/HANA patching/upgrade:

1. Download the AWS CloudFormation template offered in the SAP HANA Quick Start.
2. Update the CloudFormation template with the latest OS AMI ID and execute the updated template to provision a new SAP HANA server. The latest OS AMI ID has the specific security patches that your organization needs. As part of the provisioning process, you need to provide the latest SAP HANA installation binaries to get to the required version. This allows you to provision a new HANA system with the required OS version and security patches along with SAP HANA software versions.
3. After the new SAP HANA system is available, use one of the following methods to copy the data from the original SAP HANA instance to the newly created system:

   - Use SAP HANA native backup/restore.
   - Use SAP HANA System Replication (HSR) technology to replicate the data and then perform an HSR take-over.

- Take snapshots of the old system's Amazon Elastic Block Store (Amazon EBS) volumes and create new EBS volumes from it. Mount them in the new environment. (Make sure that the HANA SID stays the same for minimal post-processing.)
- Use new SAP HANA 2.0 functionality such as SAP HANA Cloning. The new system will become a clone of the original system.

At the end of this process, you will have a new SAP HANA system that is ready to test.

SAP Note 1984882 (*Using HANA System Replication for Hardware Exchange with Minimum/Zero Downtime*) has specific recommendations and guidelines for promoting your system to production.

Backup and Recovery

This section provides an overview of the AWS services used in the backup and recovery of SAP HANA systems and provides an example backup and recovery scenario. This guide does not include detailed instructions on how to execute database backups using native HANA backup and recovery features or third- party backup tools. Please refer to the standard OS, SAP, and SAP HANA documentation or the documentation provided by backup software vendors. In addition, backup schedules, frequency, and retention periods might vary with your system type and business requirements. See the following standard SAP documentation for guidance on these topics.

Note
For a discussion of both general and advanced backup and recovery concepts for SAP systems on AWS, see the SAP on AWS Backup and Recovery Guide.

SAP Note	Description
1642148	FAQ: SAP HANA Database Backup & Recovery
1821207	Determining required recovery files
1869119	Checking backups using hdbbackupcheck
1873247	Checking recoverability with hdbbackupdiag --check
1651055	Scheduling SAP HANA Database Backups in Linux
2484177	Scheduling backups for multi-tenant SAP HANA Cockpit 2.0

Creating an Image of an SAP HANA System

You can use the AWS Management Console or the command line to create your own AMI based on an existing instance. For more information, see the AWS documentation. You can use an AMI of your SAP HANA instance for the following purposes:

- **To create a full offline system backup** (of the OS /usr/sap, HANA shared, backup, data, and log files) – AMIs are automatically saved in multiple Availability Zones within the same AWS Region.
- **To move a HANA system from one AWS Region to another** – You can create an image of an existing EC2 instance and move it to another AWS Region by following the instructions in the AWS documentation. When the AMI has been copied to the target AWS Region, you can launch the new instance there.
- **To clone an SAP HANA system** – You can create an AMI of an existing SAP HANA system to create an exact clone of the system. See the next section for additional information.

Note
See Restoring SAP HANA Backups and Snapshots (p. 56) later in this whitepaper to view the recommended restoration steps for production environments.

Tip
The SAP HANA system should be in a consistent state before you create an AMI. To do this, stop the SAP HANA instance before creating the AMI or by following the instructions in SAP Note 1703435.

AWS Services and Components for Backup Solutions

AWS provides a number of services and options for storage and backup, including Amazon Simple Storage Service (Amazon S3), AWS Identity and Access Management (IAM), and Glacier.

Amazon S3

Amazon S3 is the center of any SAP backup and recovery solution on AWS. It provides a highly durable storage infrastructure designed for mission-critical and primary data storage. It is designed to provide 99.999999999% durability and 99.99% availability over a given year. See the Amazon S3 documentation for detailed instructions on how to create and configure an S3 bucket to store your SAP HANA backup files.

IAM

With IAM, you can securely control access to AWS services and resources for your users. You can create and manage AWS users and groups and use permissions to grant user access to AWS resources. You can create roles in IAM and manage permissions to control which operations can be performed by the entity, or AWS service, that assumes the role. You can also define which entity is allowed to assume the role.

During the deployment process, AWS CloudFormation creates an IAM role that allows access to get objects from and/or put objects into Amazon S3. That role is subsequently assigned to each EC2 instance that is hosting SAP HANA master and worker nodes at launch time as they are deployed.

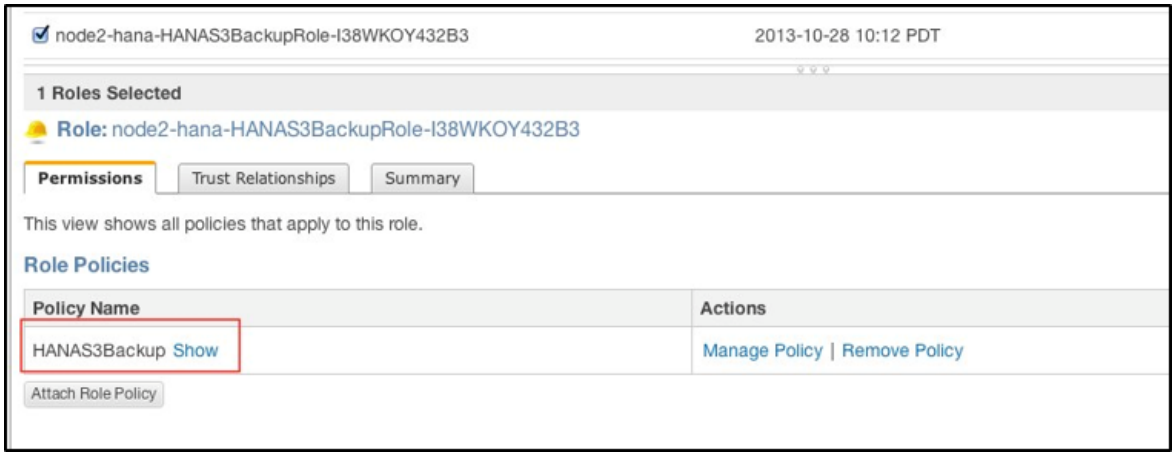

Figure 2: IAM role example

To ensure security that applies the principle of least privilege, permissions for this role are limited only to actions that are required for backup and recovery.

```
{"Statement":[
   {"Resource":"arn:aws:s3::: <your-s3-bucket-name>/*",
      "Action":["s3:GetObject","s3:PutObject","s3:DeleteObject",
"s3:ListBucket","s3:Get*","s3:List*"], "Effect":"Allow"},
```

```
{"Resource":"*","Action":["s3:List*","ec2:Describe*","ec2:Attach NetworkInterface",

"ec2:AttachVolume","ec2:CreateTags","ec2:CreateVolume","ec2:RunI nstances",
   "ec2:StartInstances"],"Effect":"Allow"}]}
```

To add functions later, you can use the AWS Management Console to modify the IAM role.

Glacier

Glacier is an extremely low-cost service that provides secure and durable storage for data archiving and backup. Glacier is optimized for data that is infrequently accessed and provides multiple options such as expedited, standard, and bulk methods for data retrieval. With standard and bulk retrievals, data is available in 3-5 hours or 5-12 hours, respectively.

However, with expedited retrieval, Glacier provides you with an option to retrieve data in 3-5 minutes, which can be ideal for occasional urgent requests. With Glacier, you can reliably store large or small amounts of data for as little as $0.01 per gigabyte per month, a significant savings compared to on-premises solutions. You can use lifecycle policies, as explained in the *Amazon S3 Developer Guide*, to push SAP HANA backups to Glacier for long-term archiving.

Backup Destination

The primary difference between backing up SAP systems on AWS compared with traditional on-premises infrastructure is the backup destination. Tape is the typical backup destination used with on-premises infrastructure. On AWS, backups are stored in Amazon S3. Amazon S3 has many benefits over tape, including the ability to automatically store backups offsite from the source system, since data in Amazon S3 is replicated across multiple facilities within the AWS Region.

SAP HANA systems provisioned by using the SAP HANA Quick Start reference deployment are configured with a set of EBS volumes to be used as an initial local backup destination. HANA backups are first stored on these local EBS volumes and then copied to Amazon S3 for long-term storage.

You can use SAP HANA Studio, SQL commands, or the DBA Cockpit to start or schedule SAP HANA data backups. Log backups are written automatically unless disabled. The /backup file system is configured as part of the deployment process.

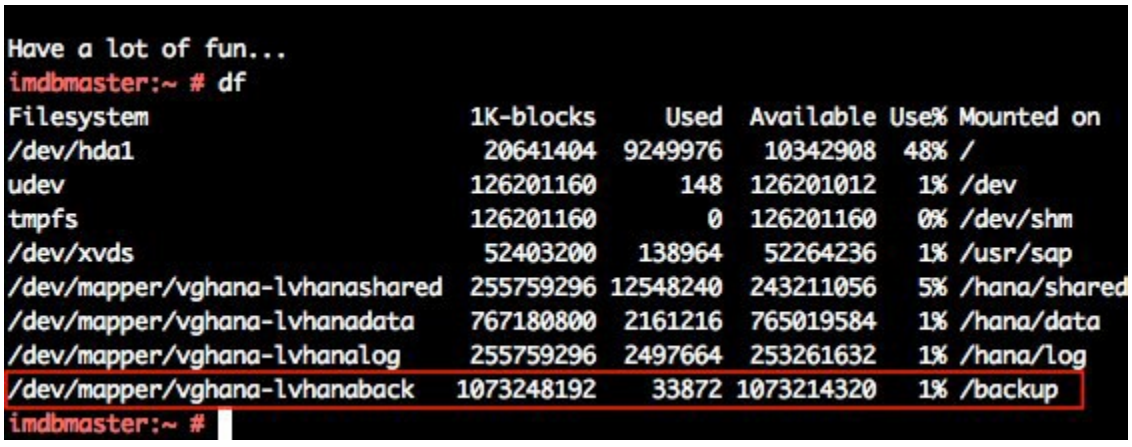

Figure 3: SAP HANA file system layout

The SAP HANA global.ini configuration file has been customized by the SAP HANA Quick Start reference deployment process as follows: database backups go directly to `/backup/data/<SID>`, while automatic log archival files go to `/backup/log/<SID>`.

```
[persistence]
```

```
basepath_shared = no
savepoint_intervals = 300
basepath_datavolumes = /hana/data/<SID>
basepath_logvolumes = /hana/log/<SID>
basepath_databackup = /backup/data/<SID>
basepath_logbackup = /backup/log/<SID>
```

Some third-party backup tools like Commvault, NetBackup, and IBM Tivoli Storage Manager (IBM TSM) are integrated with Amazon S3 capabilities and can be used to trigger and save SAP HANA backups directly into Amazon S3 without needing to store the backups on EBS volumes first.

AWS CLI

The AWS Command Line Interface (AWS CLI), which is a unified tool to manage AWS services, is installed as part of the base image. Using various commands, you can control multiple AWS services from the command line directly and automate them through scripts. Access to your S3 bucket is available through the IAM role assigned to the instance (as discussed earlier (p. 49)). Using the AWS CLI commands for Amazon S3, you can list the contents of the previously created bucket, back up files, and restore files, as explained in the AWS CLI documentation.

```
imdbmaster:/backup # aws s3 ls --region=us-east-1 s3://node2- hana-s3bucket-gcynh5v2nqs3

Bucket: node2-hana-s3bucket-gcynh5v2nqs3
Prefix:
      LastWriteTime      Length      Name
      -------------      ------      ----
```

Backup Example

Here are the steps you might take for a typical backup task:

1. In the SAP HANA Backup Editor, choose **Open Backup Wizard**. You can also open the Backup Wizard by right-clicking the system that you want to back up and choosing **Back Up**.
    1. Select the destination type **File**. This will back up the database to files in the specified file system.
    2. Specify the backup destination (`/backup/data/<SID>`) and the backup prefix.

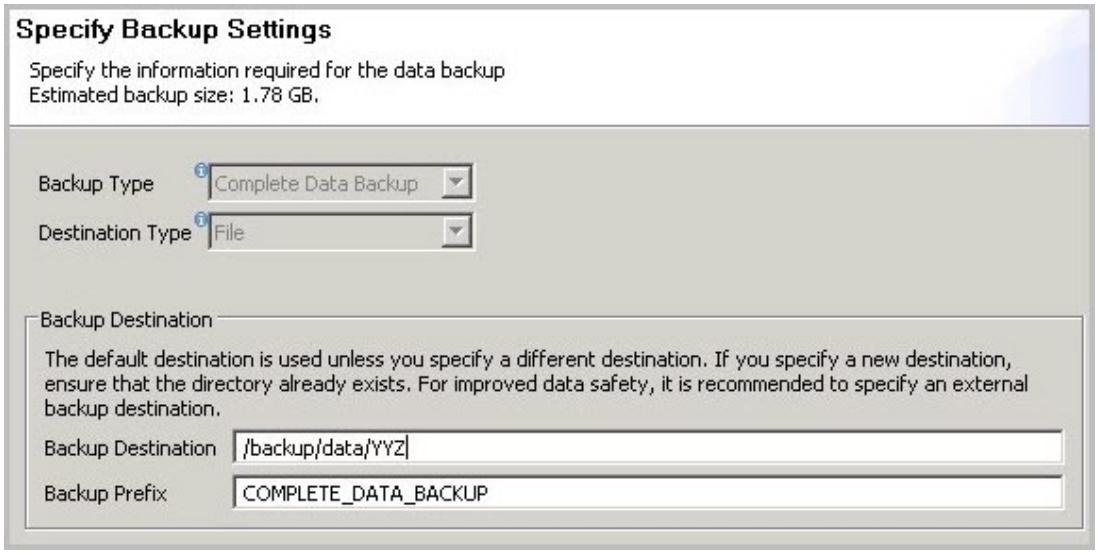

Figure 4: SAP HANA backup example

3. Choose **Next** and then **Finish**. A confirmation message will appear when the backup is complete.
4. Verify that the backup files are available at the OS level. The next step is to push or synchronize the backup files from the /backup file system to Amazon S3 by using the aws s3 sync command.

```
imdbmaster:/ # aws s3 sync backup s3://node2-hana-s3bucket- gcynh5v2nqs3 --region=us-
east-1
```

2. Use the AWS Management Console to verify that the files have been pushed to Amazon S3. You can also use the aws s3 ls command shown previously in the AWS Command Line Interface section (p. 51).

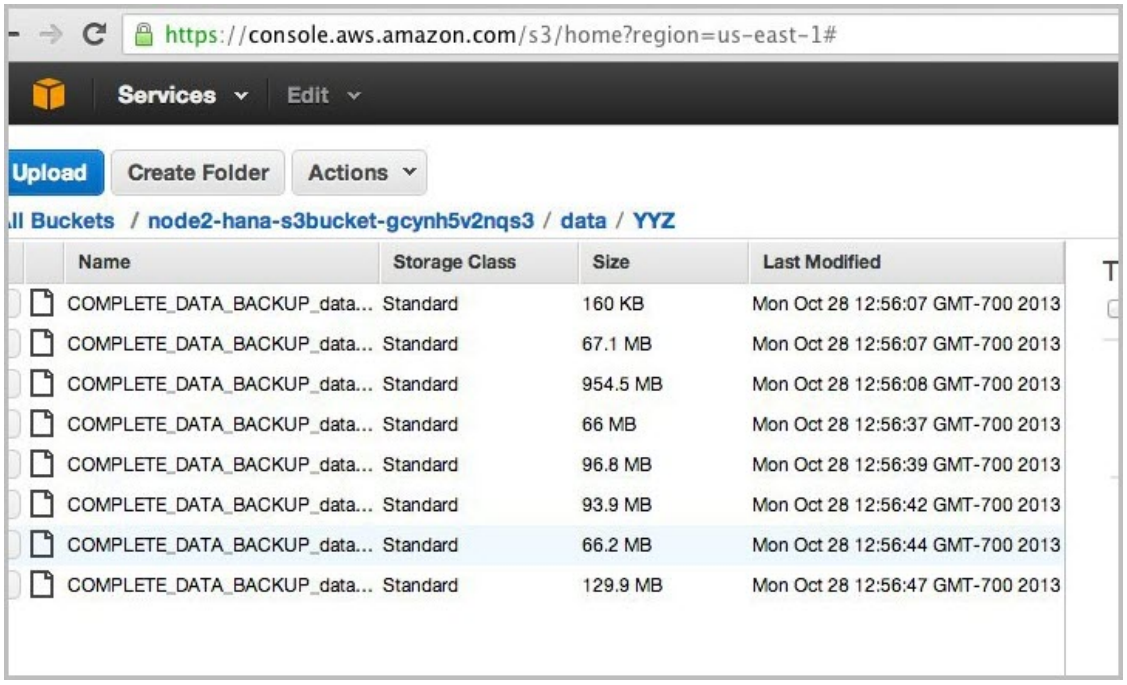

Figure 5: Amazon S3 bucket contents after backup

Tip
The `aws s3 sync` command will only upload new files that don't exist in Amazon S3. Use a periodically scheduled `cron` job to sync, and then delete files that have been uploaded. See SAP Note 1651055 for scheduling periodic backup jobs in Linux, and extend the supplied scripts with `aws s3 sync` commands.

Scheduling and Executing Backups Remotely

You can use the AWS Systems Manager Run Command, along with Amazon CloudWatch Events, to schedule backups of your SAP HANA system remotely without the need to log in to the EC2 instances. You can also use `cron` or any other instance-level scheduling mechanism.

The Systems Manager Run Command lets you remotely and securely manage the configuration of your managed instances. A managed instance is any EC2 instance or on-premises machine in your hybrid environment that has been configured for Systems Manager. The Run Command enables you to automate common administrative tasks and perform ad hoc configuration changes at

scale. You can use the Run Command from the Amazon EC2 console, the AWS CLI, Windows PowerShell, or the AWS SDKs.

Systems Manager Prerequisites

Systems Manager has the following prerequisites.

Supported operating system (Linux)	Instances must run a supported version of Linux. 64-bit and 32-bit systems: • Amazon Linux 2014.09, 2014.03 or later • Ubuntu Server 16.04 LTS, 14.04 LTS, or 12.04 LTS • Red Hat Enterprise Linux (RHEL) 6.5 or later • CentOS 6.3 or later 64-bit systems only: • Amazon Linux 2015.09, 2015.03 or later • Red Hat Enterprise Linux (RHEL) 7.x or later • CentOS 7.1 or later • SUSE Linux Enterprise Server (SLES) 12 or higher For the latest information about supported operating systems, see the AWS Systems Manager documentation.
Roles for Systems Manager	Systems Manager requires an IAM role for instances that will process commands and a separate role for users who are executing commands. Both roles require permission policies that enable them to communicate with the Systems Manager API. You can choose to use Systems Manager managed policies or you can create your own roles and specify permissions. For more information, see Configuring Security Roles for Systems Manager in the AWS documentation. If you are configuring on-premises servers or virtual machines (VMs) that you want to configure using Systems Manager, you must also configure an IAM service role. For more information, see Create an IAM Service Role in the AWS documentation.
SSM Agent (EC2 Linux instances)	AWS Systems Manager Agent (SSM Agent) processes Systems Manager requests and configures your machine as specified in the request. You must download and install SSM Agent to your EC2 Linux instances. For more information, see Installing SSM Agent on Linux in the AWS documentation.

To schedule remote backups, follow these high-level steps:

1. Install and configure SSM Agent on the EC2 instance. For detailed installation steps, see the AWS Systems Manager documentation.
2. Provide SSM access to the EC2 instance role that is assigned to the SAP HANA instance. For detailed information on how to assign SSM access to a role, see the AWS Systems Manager documentation.
3. Create an SAP HANA backup script. You can use the following sample script as a starting point and modify it to meet your requirements.

```
#!/bin/sh
set -x
S3Bucket_Name=<Name of the S3 bucket where backup files will be copied>
TIMESTAMP=$(date +\%F\_%H\%M)
exec 1>/backup/data/${SAPSYSTEMNAME}/${TIMESTAMP}_backup_log.out 2>&1
echo "Starting to take backup of Hana Database and Upload the backup files to S3"
echo "Backup Timestamp for $SAPSYSTEMNAME is $TIMESTAMP" BACKUP_PREFIX=
${SAPSYSTEMNAME}_${TIMESTAMP}
echo $BACKUP_PREFIX
# source HANA environment
source $DIR_INSTANCE/hdbenv.sh
# execute command with user key
hdbsql -U BACKUP "backup data using file ('$BACKUP_PREFIX')" echo "HANA Backup is
 completed"
echo "Continue with copying the backup files in to S3" echo $BACKUP_PREFIX
sudo -u root /usr/local/bin/aws s3 cp --recursive
/backup/data/${SAPSYSTEMNAME}/ s3://${S3Bucket_Name}/bkps/${SAPSYSTEMNAME}/data/ --
exclude "*" --include "${BACKUP_PREFIX}*"
echo "Copying HANA Database log files in to S3"
sudo -u root /usr/local/bin/aws s3 sync
/backup/log/${SAPSYSTEMNAME}/ s3://${S3Bucket_Name}/bkps/${SAPSYSTEMNAME}/log/ --
exclude "*" --include "log_backup*"
sudo -u root /usr/local/bin/aws s3 cp
/backup/data/${SAPSYSTEMNAME}/${TIMESTAMP}_backup_log.out
s3://${S3Bucket_Name}/bkps/${SAPSYSTEMNAME}
```

   Note
   This script takes into consideration that `hdbuserstore` has a key named `Backup`.

4. Test a one-time backup by executing an `ssm` command directly.

   Note
   For this command to execute successfully, you will have to enable `<sid>adm login` using `sudo`.

```
aws ssm send-command --instance-ids <HANA master instance ID> --document-name AWS-
RunShellScript
--parameters commands="sudo - u <HANA_SID>adm TIMESTAMP=$(date +\%F\_%H\%M)
 SAPSYSTEMNAME=<HANA_SID>
DIR_INSTANCE=/hana/shared/${SAPSYSTEMNAME}/HDB00 -i /usr/sap/HDB/HDB00/hana_backup.sh"
```

5. Using CloudWatch Events, you can schedule backups remotely at any desired frequency. Navigate to the CloudWatch Events page and create a rule.

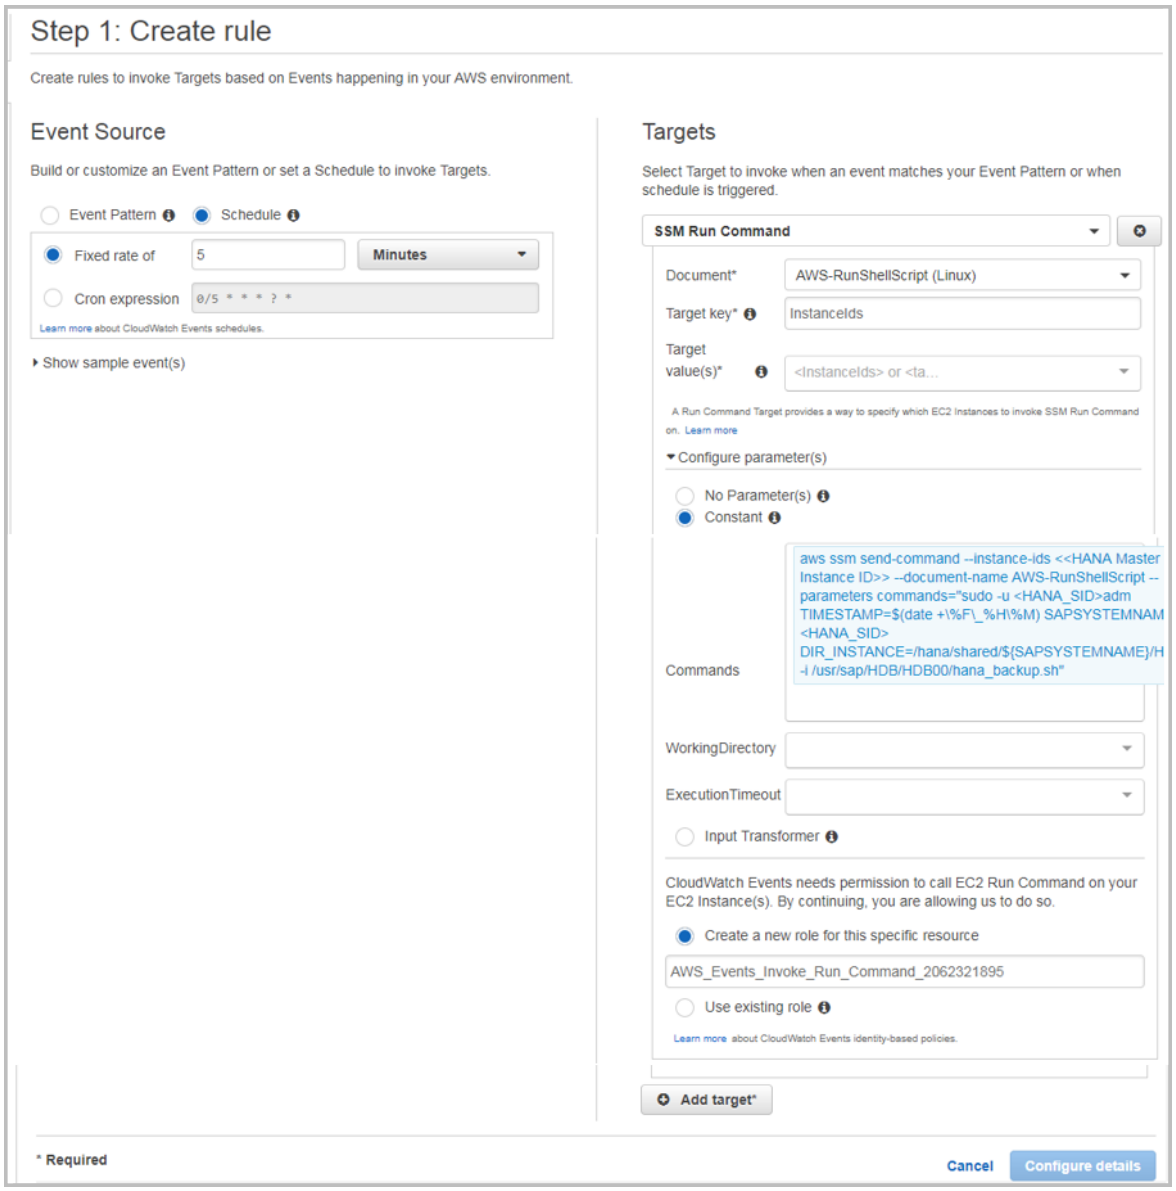

Figure 6: Creating Amazon CloudWatch Events rules

When configuring the rule:

1. Choose **Schedule**.
2. Select **SSM Run Command** as the target.
3. Select **AWS-RunShellScript (Linux)** as the document type.
4. Choose **InstanceIds** or **Tags** as the target key.
5. Choose **Constant** under **Configure Parameters**, and type the `run` command.

Restoring SAP HANA Backups and Snapshots

Restoring SAP Backups

To restore your SAP HANA database from a backup, perform the following steps:

1. If the backup files are not already available in the /backup file system but are in Amazon S3, restore the files from Amazon S3 by using the aws s3 cp command. This command has the following syntax:

```
aws --region <region> cp <s3-bucket/path> --recursive <backup- prefix>*.
```

   For example:

```
imdbmaster:/backup/data/YYZ # aws --region us-east-1 s3 cp s3://node2-hana-s3bucket-
gcynh5v2nqs3/data/YYZ . --recursive -- include COMPLETE*
```

2. Recover the SAP HANA database by using the Recovery Wizard as outlined in the SAP HANA Administration Guide. Specify **File** as the destination type and enter the correct backup prefix.

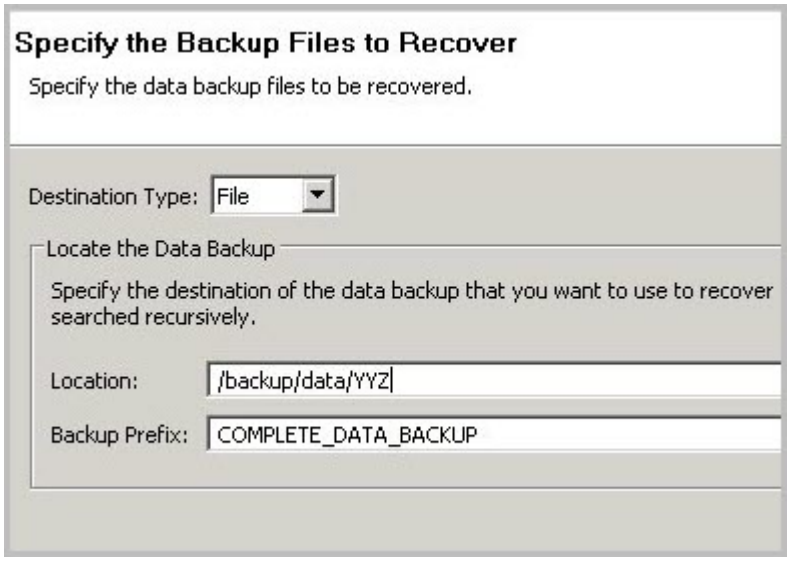

Figure 7: Restore example

3. When the recovery is complete, you can resume normal operations and clean up backup files from the `/backup/<SID>/*` directories.

Restoring EBS Snapshots

To restore EBS snapshots, perform the following steps:

1. Create a new volume from the snapshot:

```
aws ec2 create-volume --region us-west-2 --availability-zone us- west-2a --snapshot-id
 snap-1234abc123a12345a --volume-type gp2
```

2. Attach the newly created volume to your EC2 host:

```
aws ec2 attach-volume --region=us-west-2 --volume-id vol- 4567c123e45678dd9 --instance-
id i-03add123456789012 --device /dev/sdf
```

3. Mount the logical volume associated with SAP HANA data on the host:

```
mount /dev/sdf /hana/data
```

4. Start your SAP HANA instance.

Note
For large mission-critical systems, we highly recommend that you execute the volume initialization command on the database data and log volumes after restoring the AMI but before starting the database. Executing the volume initialization command will help you avoid extensive wait times before the database is available. Here is the sample `fio` command that you can use:

```
sudo fio -filename=/dev/xvdf -rw=read -bs=128K -iodepth=32 -
 ioengine=libaiodirect=1 -name=volume-initialize
```

For more information about initializing Amazon EBS volumes, see the AWS documentation.

Restoring AMI Snapshots

You can restore your SAP HANA AMI snapshots through the AWS Management Console. Open the Amazon EC2 console, and choose **AMIs** in the navigation pane.

Choose the AMI that you want to restore, expand **Actions**, and then choose **Launch**.

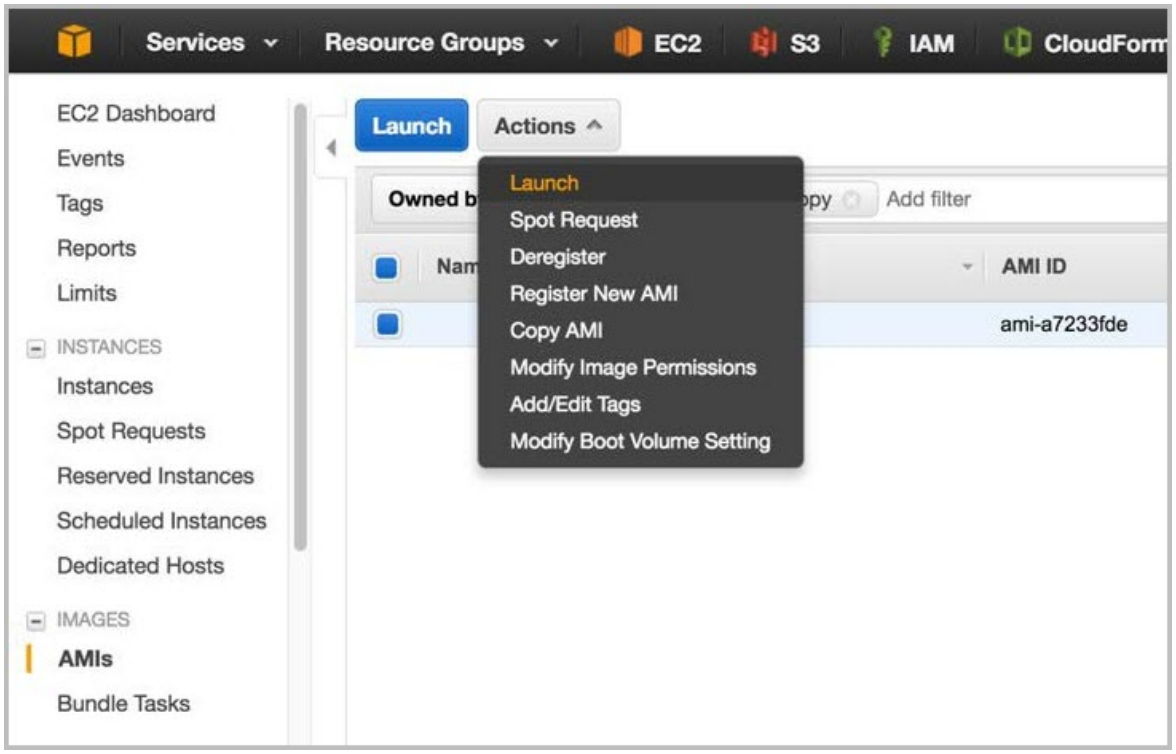

Figure 8: Restoring an AMI snapshot

Networking

SAP HANA components communicate over the following logical network zones:

- Client zone – to communicate with different clients such as SQL clients, SAP Application Server, SAP HANA Extended Application Services (XS), and SAP HANA Studio
- Internal zone – to communicate with hosts in a distributed SAP HANA system as well as for SAP HSR
- Storage zone – to persist SAP HANA data in the storage infrastructure for resumption after start or recovery after failure

Separating network zones for SAP HANA is considered both an AWS and an SAP best practice, because it enables you to isolate the traffic required for each communication channel.

In a traditional, bare-metal setup, these different network zones are set up by having multiple physical network cards or virtual LANs (VLANs). Conversely, on the AWS Cloud, you can use elastic network interfaces combined with security groups to achieve this network isolation. Amazon EBS-optimized instances can also be used for further isolation for storage I/O.

EBS-Optimized Instances

Many newer Amazon EC2 instance types such as the X1 use an optimized configuration stack and provide additional, dedicated capacity for Amazon EBS I/O. These are called EBS-optimized instances. This optimization provides the best performance for your EBS volumes by minimizing contention between Amazon EBS I/O and other traffic from your instance.

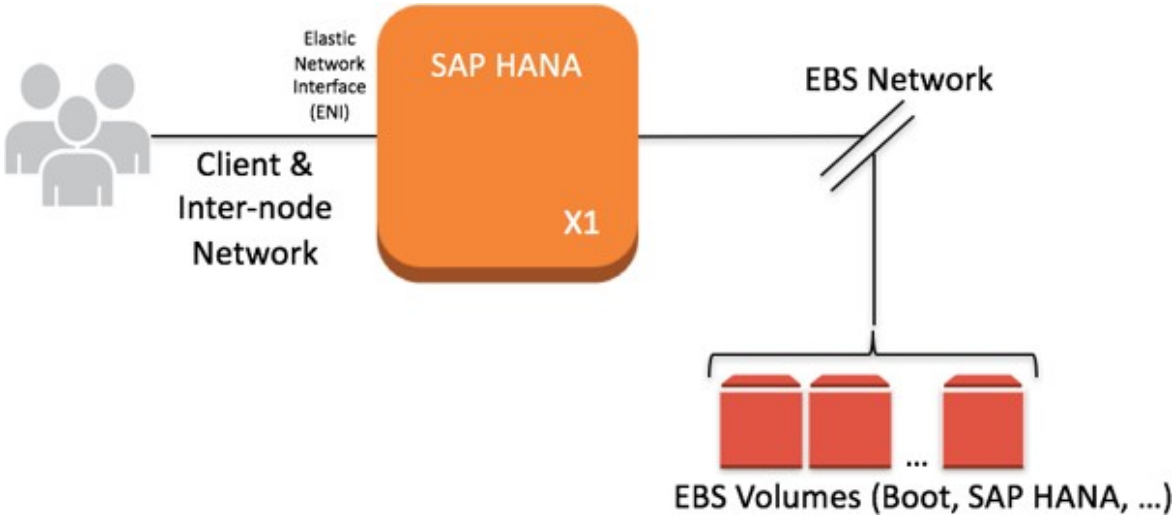

Figure 9: EBS-optimized instances

Elastic Network Interfaces

An elastic network interface is a virtual network interface that you can attach to an EC2 instance in an Amazon Virtual Private Cloud (Amazon VPC). With an elastic network interface (referred to as *network interface* in the remainder of this guide), you can create different logical networks by specifying multiple private IP addresses for your instances.

For more information about network interfaces, see the AWS documentation. In the following example, two network interfaces are attached to each SAP HANA node as well as in a separate communication channel for storage.

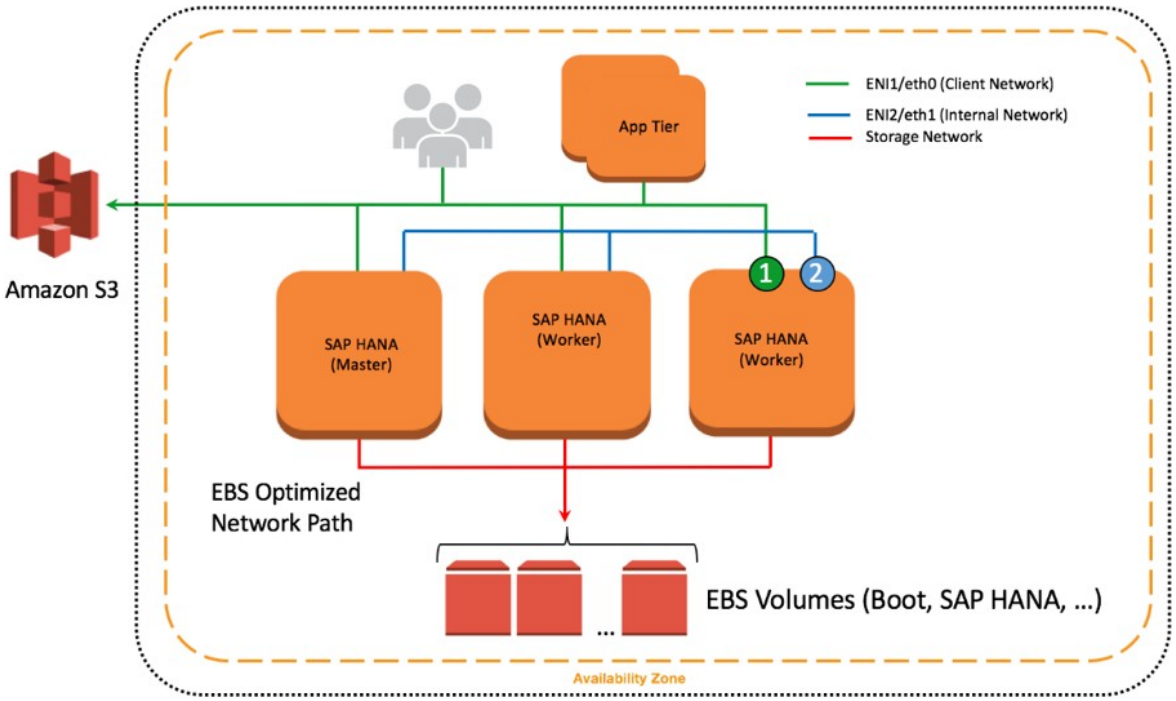

Figure 10: Network interfaces attached to SAP HANA nodes

Security Groups

A security group acts as a virtual firewall that controls the traffic for one or more instances. When you launch an instance, you associate one or more security groups with the instance. You add rules to each security group that allow traffic to or from its associated instances. You can modify the rules for a security group at any time. The new rules are automatically applied to all instances that are associated with the security group. To learn more about security groups, see the AWS documentation. In the following example, ENI-1 of each instance shown is a member of the same security group that controls inbound and outbound network traffic for the client network.

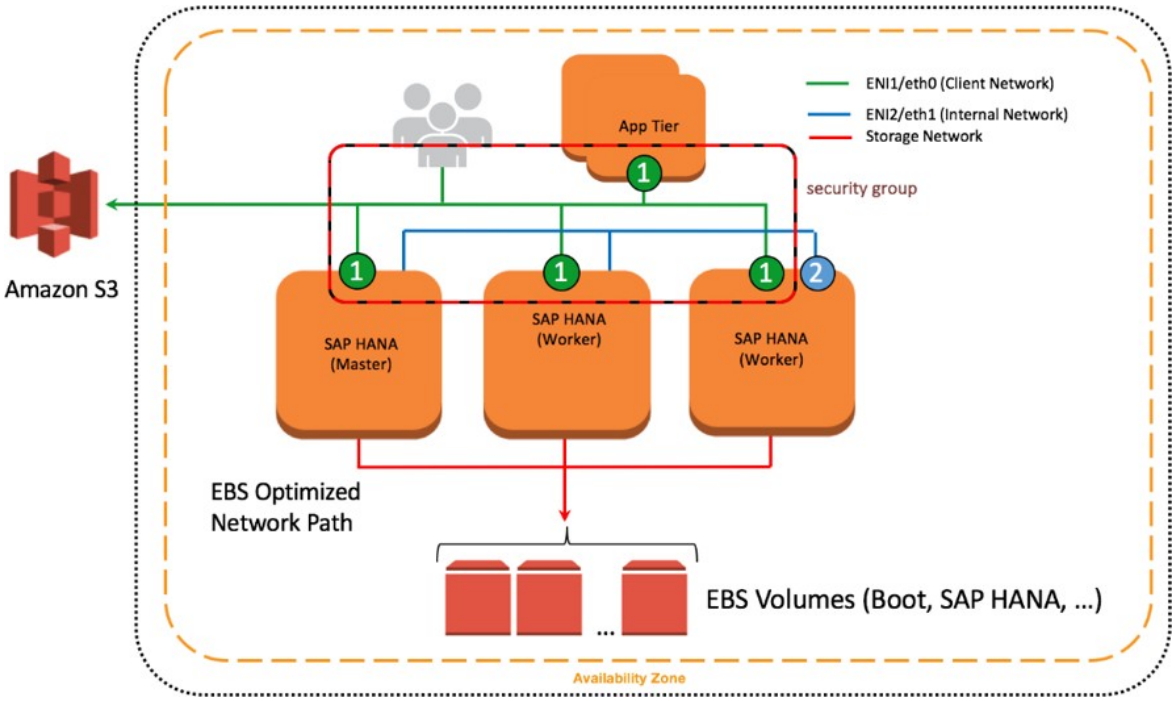

Figure 11: Network interfaces and security groups

Network Configuration for SAP HANA System Replication (HSR)

You can configure additional network interfaces and security groups to further isolate inter-node communication as well as SAP HSR network traffic. In Figure 10, ENI-2 is has its own security group (not shown) to secure client traffic from inter-node communication. ENI-3 is configured to secure SAP HSR traffic to another Availability Zone within the same Region. In this example, the target SAP HANA cluster would be configured with additional network interfaces similar to the source environment, and ENI-3 would share a common security group.

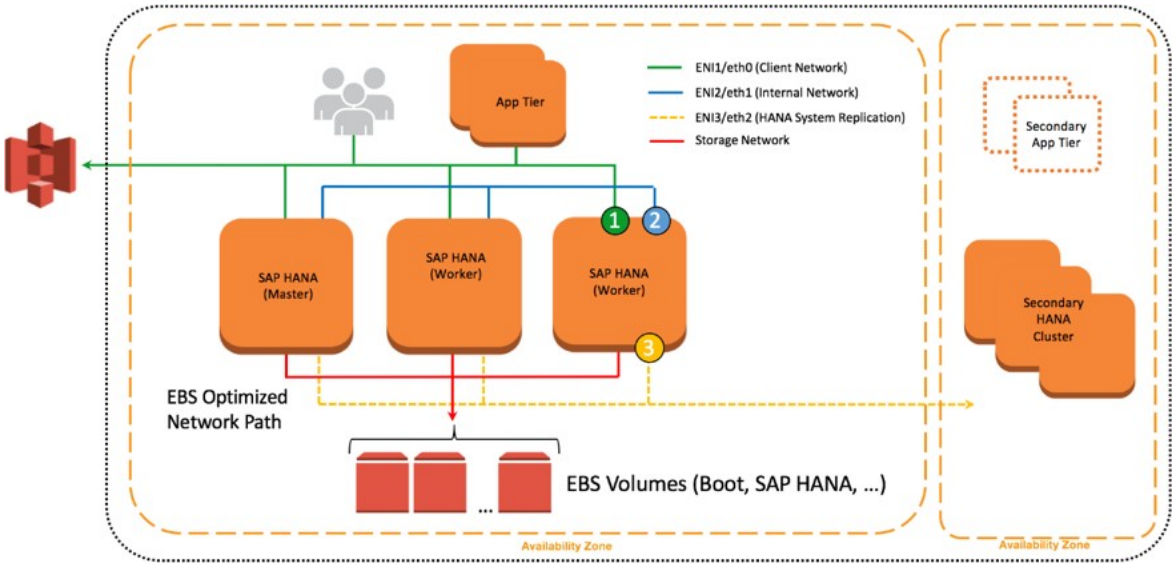

Figure 12: Further isolation with additional ENIs and security groups

Configuration Steps for Logical Network Separation

To configure your logical network for SAP HANA, follow these steps:

1. Create new security groups to allow for isolation of client, internal communication, and, if applicable, SAP HSR network traffic. See Ports and Connections in the SAP HANA documentation to learn about the list of ports used for different network zones. For more information about how to create and configure security groups, see the AWS documentation.
2. Use Secure Shell (SSH) to connect to your EC2 instance at the OS level. Follow the steps described in the appendix (p. 64) to configure the OS to properly recognize and name the Ethernet devices associated with the new network interfaces you will be creating.
3. Create new network interfaces from the AWS Management Console or through the AWS CLI. Make sure that the new network interfaces are created in the subnet where your SAP HANA instance is deployed. As you create each new network interface, associate it with the appropriate security group you created in step 1. For more information about how to create a new network interface, see the AWS documentation.
4. Attach the network interfaces you created to your EC2 instance where SAP HANA is installed. For more information about how to attach a network interface to an EC2 instance, see the AWS documentation.
5. Create virtual host names and map them to the IP addresses associated with client, internal, and replication network interfaces. Ensure that host name-to-IP-address resolution is working by creating entries in all applicable host files or in the Domain Name System (DNS). When complete, test that the virtual host names can be resolved from all SAP HANA nodes and clients.
6. For scale-out deployments, configure SAP HANA inter-service communication to let SAP HANA communicate over the internal network. To learn more about this step, see Configuring SAP HANA Inter-Service Communication in the SAP HANA documentation.
7. Configure SAP HANA hostname resolution to let SAP HANA communicate over the replication network for SAP HSR. To learn more about this step, see Configuring Hostname Resolution for SAP HANA System Replication in the SAP HANA documentation.

SAP Support Access

In some situations it may be necessary to allow an SAP support engineer to access your SAP HANA systems on AWS. The following information serves only as a supplement to the information contained in the “Getting Support” section of the SAP HANA Administration Guide.

A few steps are required to configure proper connectivity to SAP. These steps differ depending on whether you want to use an existing remote network connection to SAP, or you are setting up a new connection directly with SAP from systems on AWS.

Support Channel Setup with SAProuter on AWS

When setting up a direct support connection to SAP from AWS, consider the following steps:

1. For the SAProuter instance, create and configure a specific SAProuter security group, which only allows the required inbound and outbound access to the SAP support network. This should be limited to a specific IP address that SAP gives you to connect to, along with TCP port 3299. See the Amazon EC2 security group documentation for additional details about creating and configuring security groups.

2. Launch the instance that the SAProuter software will be installed on into a public subnet of the VPC and assign it an Elastic IP address.
3. Install the SAProuter software and create a saprouttab file that allows access from SAP to your SAP HANA system on AWS.
4. Set up the connection with SAP. For your internet connection, use **Secure Network Communication (SNC)**. For more information, see the SAP Remote Support – Help page.
5. Modify the existing SAP HANA security groups to trust the new SAProuter security group you have created.

   Tip
   For added security, shut down the EC2 instance that hosts the SAProuter service when it is not needed for support purposes

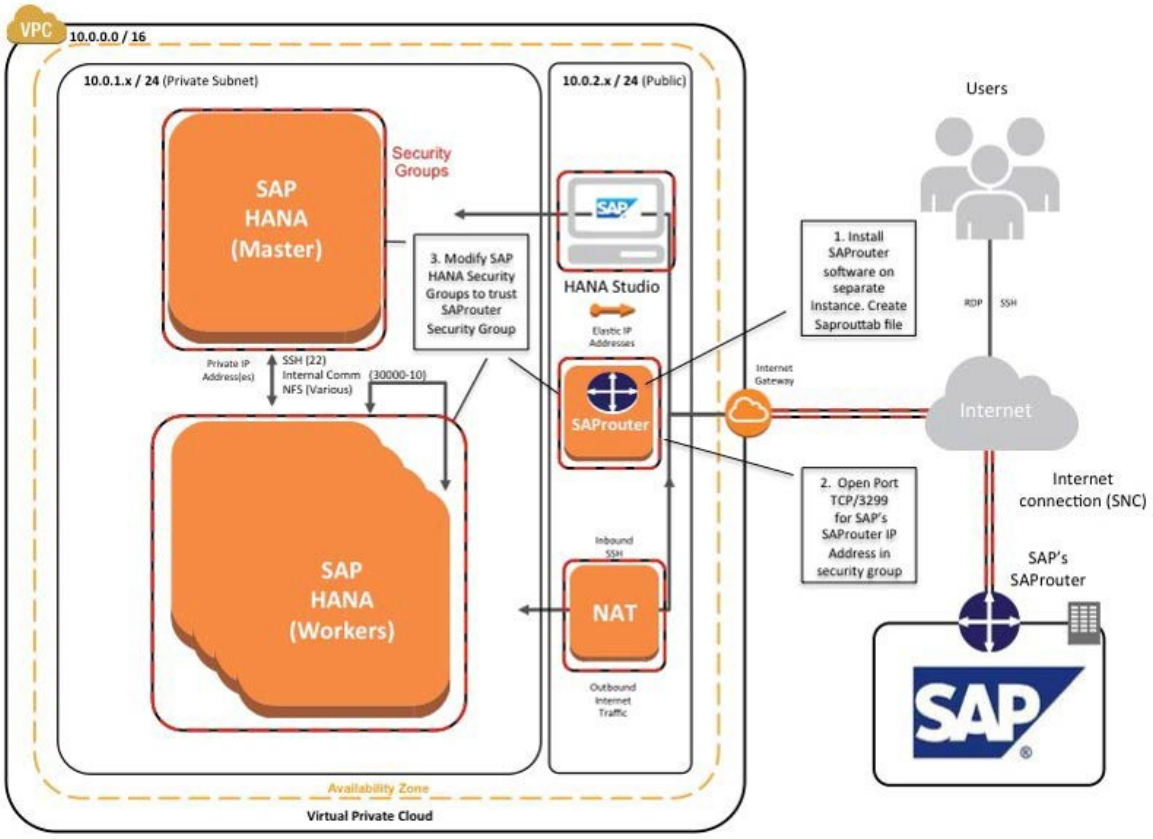

Figure 13: Support connectivity with SAProuter on AWS

Support Channel Setup with SAProuter on Premises

In many cases, you may already have a support connection configured between your data center and SAP. This can easily be extended to support SAP systems on AWS. This scenario assumes that connectivity between your data center and AWS has already been established, either by way of a secure VPN tunnel over the internet or by using AWS Direct Connect.

You can extend this connectivity as follows:

1. Ensure that the proper saprouttab entries exist to allow access from SAP to resources in the VPC.

2. Modify the SAP HANA security groups to allow access from the on- premises SAProuter IP address.
3. Ensure that the proper firewall ports are open on your gateway to allow traffic to pass over TCP port 3299.

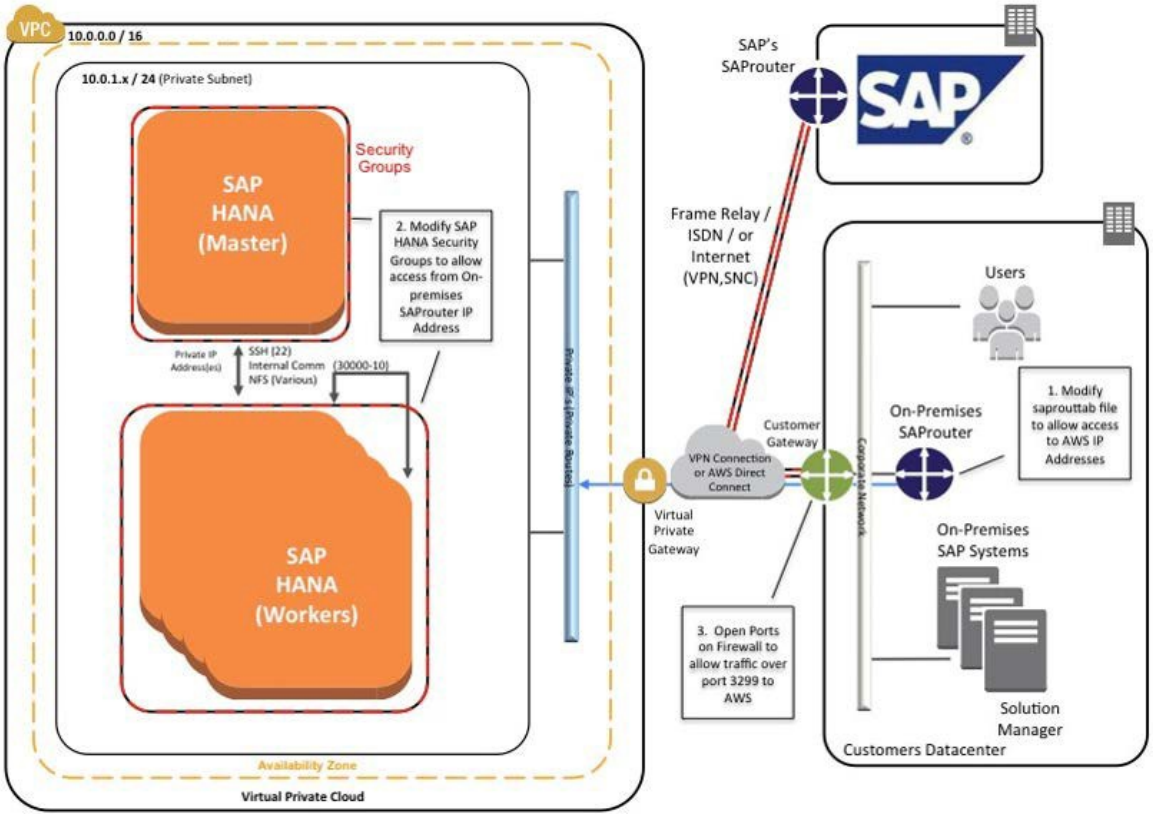

Figure 14: Support connectivity with SAProuter on premises

Security

This section discusses additional security topics you may want to consider that are not covered in the SAP HANA Quick Start reference deployment guide.

Here are additional AWS security resources to help you achieve the level of security you require for your SAP HANA environment on AWS:

- AWS Cloud Security Center
- CIS AWS Foundation whitepaper
- AWS Cloud Security whitepaper
- AWS Cloud Security Best Practices whitepaper

OS Hardening

You may want to lock down the OS configuration further, for example, to avoid providing a DB administrator with root credentials when logging into an instance.

You can also refer to the following SAP notes:

- 1730999: *Configuration changes in HANA appliance*
- 1731000: *Unrecommended configuration changes*

Disabling HANA Services

HANA services such as HANA XS are optional and should be deactivated if they are not needed. For instructions, see SAP Note 1697613: *Remove XS Engine out of SAP HANA database*. In case of service deactivation, you should also remove the TCP ports from the SAP HANA AWS security groups for complete security.

API Call Logging

AWS CloudTrail is a web service that records AWS API calls for your account and delivers log files to you. The recorded information includes the identity of the API caller, the time of the API call, the source IP address of the API caller, the request parameters, and the response elements returned by the AWS service.

With CloudTrail, you can get a history of AWS API calls for your account, including API calls made via the AWS Management Console, AWS SDKs, command line tools, and higher-level AWS services (such as AWS CloudFormation). The AWS API call history produced by CloudTrail enables security analysis, resource change tracking, and compliance auditing.

Notifications on Access

You can use Amazon Simple Notification Service (Amazon SNS) or third-party applications to set up notifications on SSH login to your email address or mobile phone.

High Availability and Disaster Recovery

For details and best practices for high availability and disaster recovery of SAP HANA systems running on AWS, see High Availability and Disaster Recovery Options for SAP HANA on AWS.

Appendix: Configuring Linux to Recognize Ethernet Devices for Multiple Network Interfaces

Follow these steps to configure the Linux operating system to recognize and name the Ethernet devices associated with the new elastic network interfaces created for logical network separation, which were discussed earlier in this paper (p. 61).

1. Use SSH to connect to your SAP HANA host as `ec2-user`, and `sudo` to root.
2. Remove the existing `udev` rule; for example:

```
hanamaster:# rm -f /etc/udev/rules.d/70-persistent-net.rules
```

3. Create a new `udev` rule that writes rules based on MAC address rather than other device attributes. This will ensure that on reboot, `eth0` is still `eth0`, `eth1` is `eth1`, and so on. For example:

```
hanamaster:# cat <<EOF >/etc/udev/rules.d/75-persistent-net- generator.rules
```

```
# Copyright (C) 2012 Amazon.com, Inc. or its affiliates. # All Rights Reserved.
#
# Licensed under the Apache License, Version 2.0 (the "License").
# You may not use this file except in compliance with the License.
# A copy of the License is located at #
#     http://aws.amazon.com/apache2.0/ #
# or in the "license" file accompanying this file. This file is # distributed on an "AS
 IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS
# OF ANY KIND, either express or implied. See the License for the
# specific language governing permissions and limitations under the
# License.
# these rules generate rules for persistent network device naming
SUBSYSTEM!="net", GOTO="persistent_net_generator_end" KERNEL!="eth*",
 GOTO="persistent_net_generator_end" ACTION!="add", GOTO="persistent_net_generator_end"
 NAME=="?*", GOTO="persistent_net_generator_end"

# do not create rule for eth0
ENV{INTERFACE}=="eth0", GOTO="persistent_net_generator_end" # read MAC address
ENV{MATCHADDR}="\$attr{address}" # do not use empty address
ENV{MATCHADDR}=="00:00:00:00:00:00",
GOTO="persistent_net_generator_end"
# discard any interface name not generated by our rules ENV{INTERFACE_NAME}=="?*",
 ENV{INTERFACE_NAME}=""
# default comment
ENV{COMMENT}="elastic network interface" # write rule IMPORT{program}="write_net_rules"
# rename interface if needed ENV{INTERFACE_NEW}=="?*", NAME="\$env{INTERFACE_NEW}"
LABEL="persistent_net_generator_end" EOF
```

4. Ensure proper interface properties. For example:

```
hanamaster:# cd /etc/sysconfig/network/

hanamaster:# cat <<EOF >/etc/sysconfig/network/ifcfg-ethN
BOOTPROTO='dhcp4'
MTU="9000"
REMOTE_IPADDR=''
STARTMODE='onboot'
LINK_REQUIRED=no
LINK_READY_WAIT=5
EOF
```

5. Ensure that you can accommodate up to seven more Ethernet devices or network interaces, and restart `wicked`. For example:

```
hanamaster:# for dev in eth{1..7} ; do
ln -s -f ifcfg-ethN /etc/sysconfig/network/ifcfg-${dev} done

hanamaster:# systemctl restart wicked
```

6. Create and attach a new network interface to the instance.
7. Reboot.
8. Modify `/etc/iproute2/rt_tables`.

   Important
   Repeat the following for each ENI that you attach to your instance.

   For example:

```
hanamaster:# cd /etc/iproute2
hanamaster:/etc/iproute2 # echo "2 eth1_rt" >> rt_tables
hanamaster:/etc/iproute2 # ip route add default via 172.16.1.122 dev eth1 table eth1_rt
```

```
hanamaster:/etc/iproute2 # ip rule
0: from all lookup local
32766: from all lookup main
32767: from all lookup default

hanamaster:/etc/iproute2 # ip rule add from <ENI IP Address>
lookup eth1_rt prio 1000

hanamaster:/etc/iproute2 # ip rule 0: from all lookup local
1000: from <ENI IP address> lookup eth1_rt
32766: from all lookup main
32767: from all lookup default
```

Document Revisions

Date	Change	Location
December 2017	Initial publication	–

SAP HANA Data Tiering on AWS Overview

SAP specialists, Amazon Web Services (AWS)

Last updated (p. 81): July 2019

This guide is part of a content series that provides detailed information about hosting, configuring, and using SAP technologies in the Amazon Web Services (AWS) Cloud. For the other guides in the series, ranging from overviews to advanced topics, see SAP on AWS Technical Documentation home page.

Overview

This guide provides an overview of data tiering for SAP customers and partners who are considering implementing or migrating SAP environments or systems to the Amazon Web Services (AWS) Cloud.

This guide is for users who architect, design, deploy, and support SAP systems directly and IT professionals that support these same functions for their SAP systems.

Prerequisites

Specialized Knowledge

You should have previous experience installing, migrating, and operating SAP environments and systems.

Technical Requirements

To access the SAP notes referenced in this guide, you must have an SAP One Support Launchpad user account.

SAP Data Tiering

SAP data tiering is a data management strategy that's used to separate your data into different categories (hot, warm, and cold tiers) by various characteristics of the data. The most common characteristics used to assign the data to the correct categories are:

- frequency of access to the data
- requirement to update the data
- performance requirement and timely access to the data
- criticality of the data for operational business processes

Assigning your data to the correct category is a process that is specific to your business and IT requirements. Here are some ways to align these categories with your specific requirements.

Hot Tier: The hot tier is for storing data that is used (read, accessed or updated) in real time and that must be available in a performant and timely manner. This hot data is critical and valuable to the business for its operational and analytical processes.

Warm Tier: The warm tier is for data that is read less often than hot data, has less stringent performance requirements, but must still be updatable. The warm tier is integrated with the hot tier in the SAP HANA database. The benefit of this integration is a more transparent view of the data in the hot and warm data tiers. Applications accessing the data are unware that the data physically resides on different data tiers.

Cold Tier: The cold tier is for storing data that is infrequently accessed, does not require updates, can be accessed in a longer timeframe, and is not critical for daily operational or analytical processes.

The following table summarizes the data tiers and their characteristics.

Data tier characteristics

	Data access frequency	Performance requirement	Data criticality	Data updatability
Hot	High	High	High	Required
Warm	Medium	Medium	Medium	Required
Cold	Low	Low	Low	N/A

After you have assigned the data to your preferred tiers, you can map your SAP product to the data tiering solution that is supported by SAP on AWS. For more information, see SAP HANA on AWS: Supported AWS EC2 products and SAP HANA on AWS: Dynamic Tiering.

For the hot tier, this guide does not cover SAP HANA on AWS specifically. See SAP HANA on AWS documentation for more information about running SAP HANA on AWS. For the warm and cold tier, you have the following technology options shown in the following table, depending on your SAP product:

Warm and cold tier options

	Native SAP HANA	SAP BW on HANA or SAP BW/4 HANA	SAP Business Suite on HANA or SAP S/4 HANA
Hot	Certified SAP HANA EC2 instances	Amazon EC2 instances certified for SAP HANA	Amazon EC2 instances certified for SAP HANA
Warm	SAP HANA dynamic tiering SAP HANA extension node Native Storage Extension	SAP HANA extension node	Data aging
Cold	Data Lifecycle Manager (DLM) with SAP Data Hub and Amazon S3 DLM with SAP HANA Spark Controller	SAP BW NLS with SAP IQ SAP BW NLS with Hadoop and Amazon S3 SAP BW/4 HANA Data Tiering Optimization (DTO) with SAP Data Hub and Amazon S3	ILM Store with SAP IQ Data archiving and Amazon S3

Warm Data Tiering Options

The following sections discuss the warm data tiering options you have on AWS.

SAP HANA Dynamic Tiering

SAP HANA dynamic tiering is an optional add-on to the SAP HANA database to manage historical data that can be used for your native SAP HANA use case. SAP HANA dynamic tiering's purpose is to extend SAP HANA memory with a disk-centric columnar store (as opposed to SAP HANA's in-memory store) for managing less frequently accessed warm data. In this disk-centric solution, dynamic tiering service (esserver) runs on a separate dedicated server. Note that the SAP HANA dynamic tiering solution does not support all use cases. As noted in the solution table, SAP HANA dynamic tiering:

- can only be used for native SAP HANA use cases.
- provides online data storage in extended store, available for both queries and updates.
- is fully validated and supported on the AWS Cloud beginning with SAP HANA 2 SPS 2.
- is an integrated component of the SAP HANA database and cannot be operated separately from the SAP HANA database.
- allows you to store up to 5 times more data in the warm tier than in your hot tier.

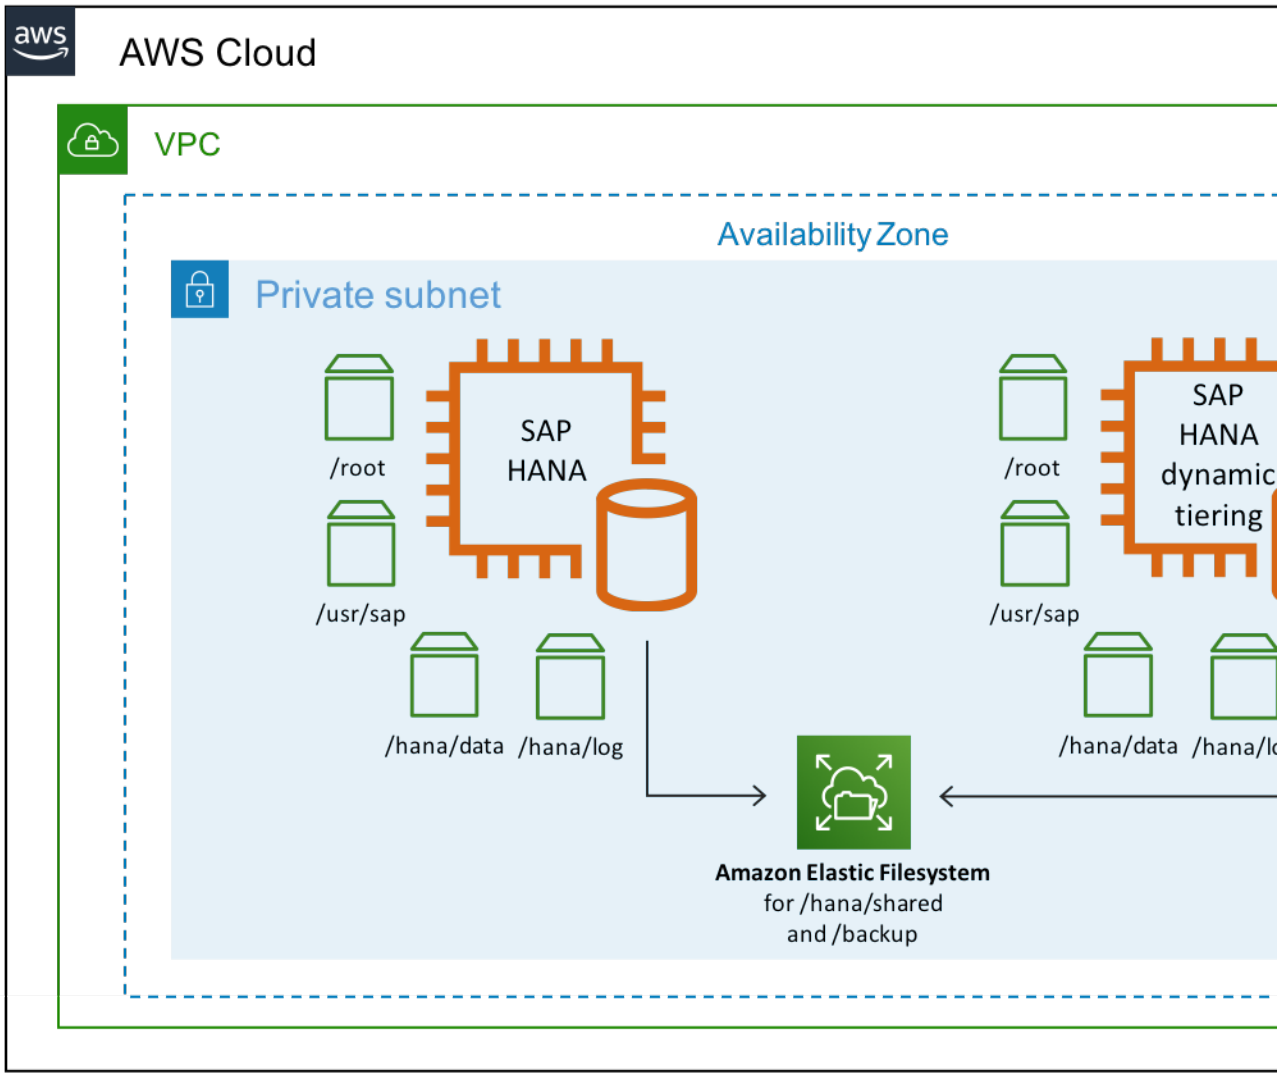

Figure 1: SAP HANA dynamic tiering on AWS (single-AZ)

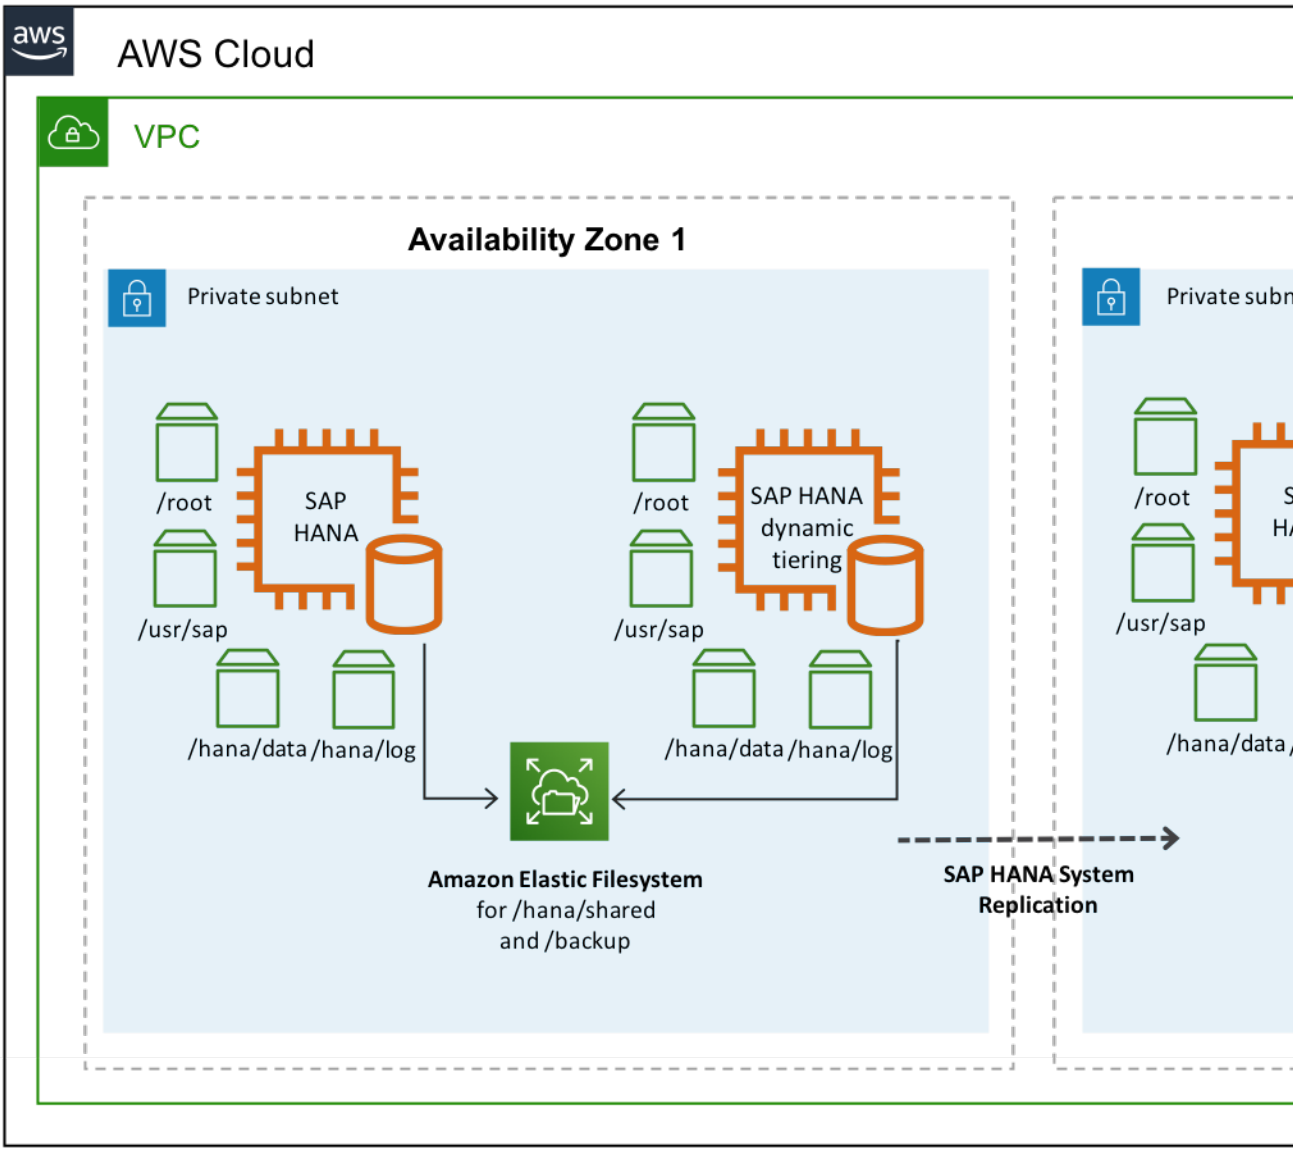

Figure 2: SAP HANA dynamic tiering on AWS (multi-AZ)

SAP HANA Extension Node

SAP HANA extension node is a special purpose SAP HANA worker node that is specifically set up and reserved for storing warm data. An important difference between SAP HANA dynamic tiering and SAP HANA extension node is that the extension node is a separate SAP HANA instance. It is not a separate process (esserver) like dynamic tiering. Because of this, the SAP HANA extension node offers the full feature set of the SAP HANA database. SAP HANA extension node allows you to store warm data for your SAP Business Warehouse (BW) or native SAP HANA use cases.

The total amount of data that can be stored on the SAP HANA extension node ranges from 1 to 2x of the total amount of memory of your extension node. For example, if your extension node had 2 TB of memory, you could potentially store up to 4 TB of warm data on your extension node.

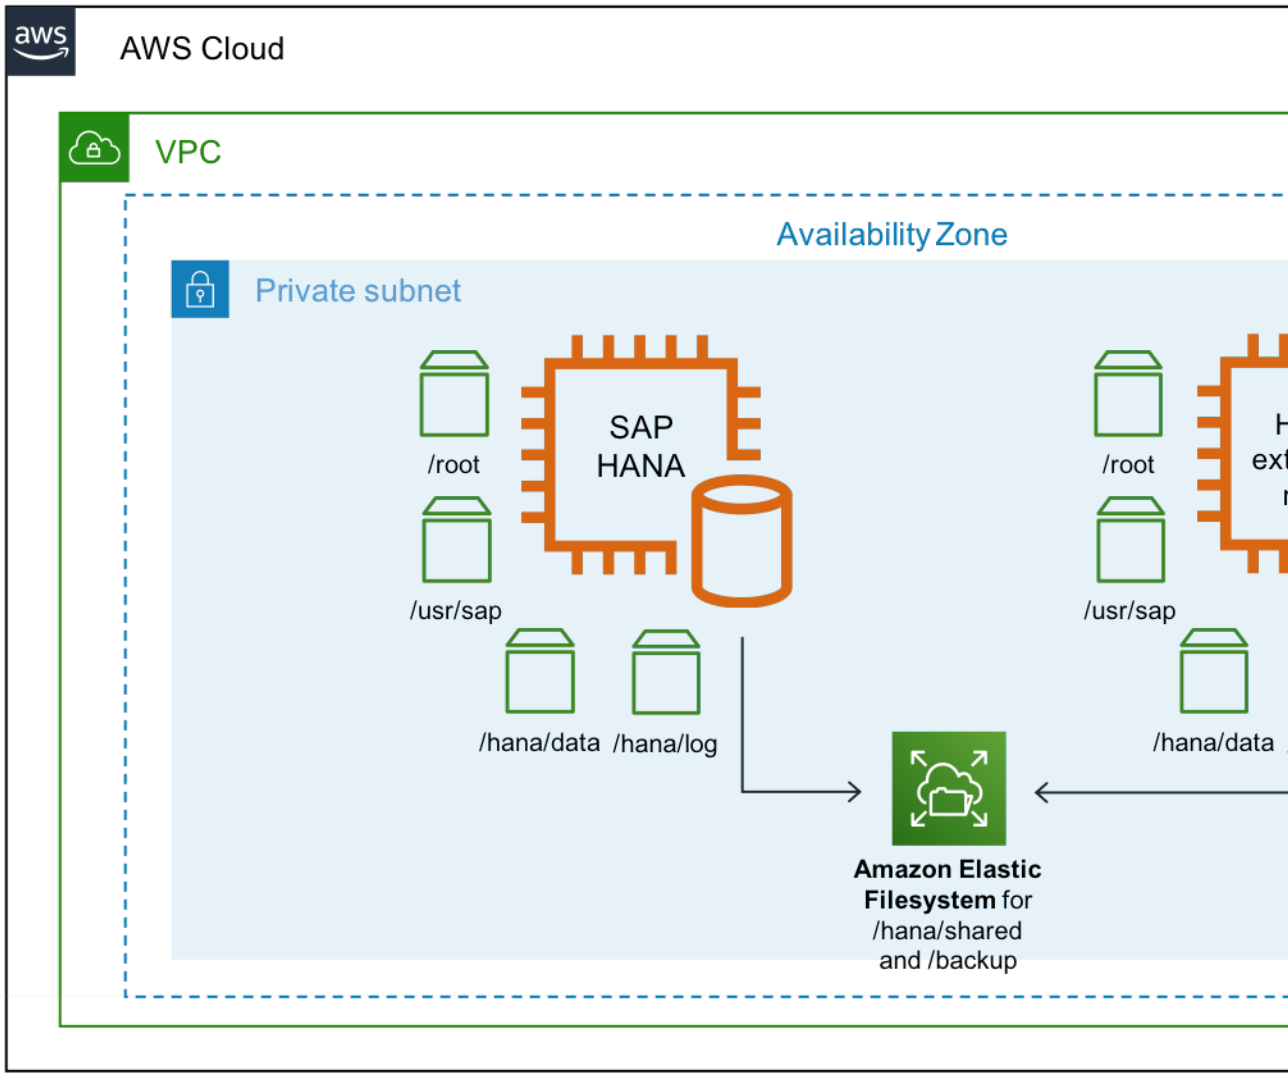

Figure 3: SAP HANA extension node on AWS

Data Aging

Data aging can be used for SAP products like SAP Business Suite on HANA (SoH) or SAP S/4HANA to move data from SAP HANA memory to the disk area. The disk area is additional disk space that is a part of the SAP HANA database. This helps free up more SAP HANA memory by storing older, less frequently accessed data in the disk area. When the data is read or updated, data aging uses the paged attribute property to selectively load the pages of a table into memory instead of loading the entire table into memory. This helps you conserve your memory space by only loading the required data (instead of the entire table) into memory. In addition, paged attributes are marked for a higher unload priority by SAP HANA and are paged out to disk first when SAP HANA needs to free up memory. To size your SAP HANA memory requirements for data aging, SAP recommends that you run the sizing report provided in the SAP Note 1872170 - ABAP on HANA sizing report (S/4HANA, Suite on HANA).

Cold Data Tiering Options

The following sections discuss cold data tiering options on AWS.

The Data Lifecycle Manager (DLM) tool, which is part of SAP HANA Data Warehousing Foundation, can be used to move data from SAP HANA memory to a cold storage location. For your native SAP HANA use case, you have two options.

DLM with SAP Data Hub

With this option, you can use the SAP Data Hub product to move data in and out of SAP HANA into your cold store location. On AWS, you are able to use native AWS services such as Amazon Simple Storage Service (Amazon S3) to store your cold data. Once your data is in Amazon S3, you can use Amazon S3 features such as S3 Intelligent-Tiering and Amazon S3 Lifecycle to optimize your costs. Once you have determined that you no longer need to access your cold data from SAP HANA, you can archive your data in Amazon S3 Glacier for long-term retention.

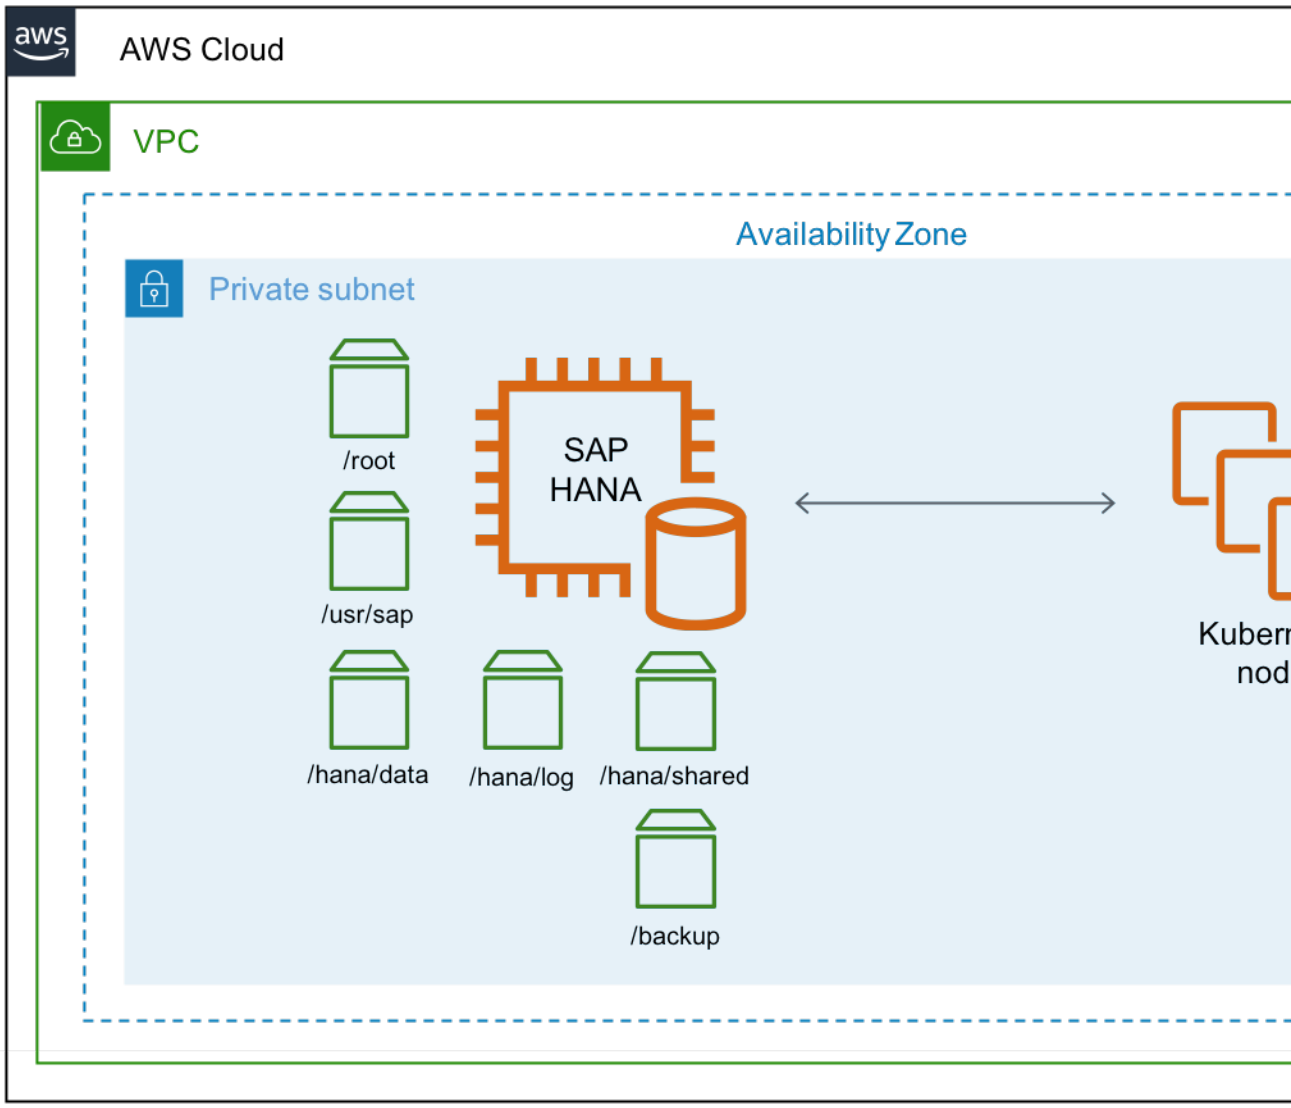

Figure 4: SAP Data Hub on Amazon EKS for cold tier

DLM with SAP HANA Spark Controller

With this option, you can use the SAP HANA Spark Controller to allow SAP HANA to access cold data through the Spark SQL SDA adapter. On AWS, you can use an AWS native service like Amazon EMR for the Hadoop cold tier storage location. To use Amazon EMR with SAP HANA, see DLM on Amazon Elastic Map Reduce documentation from SAP.

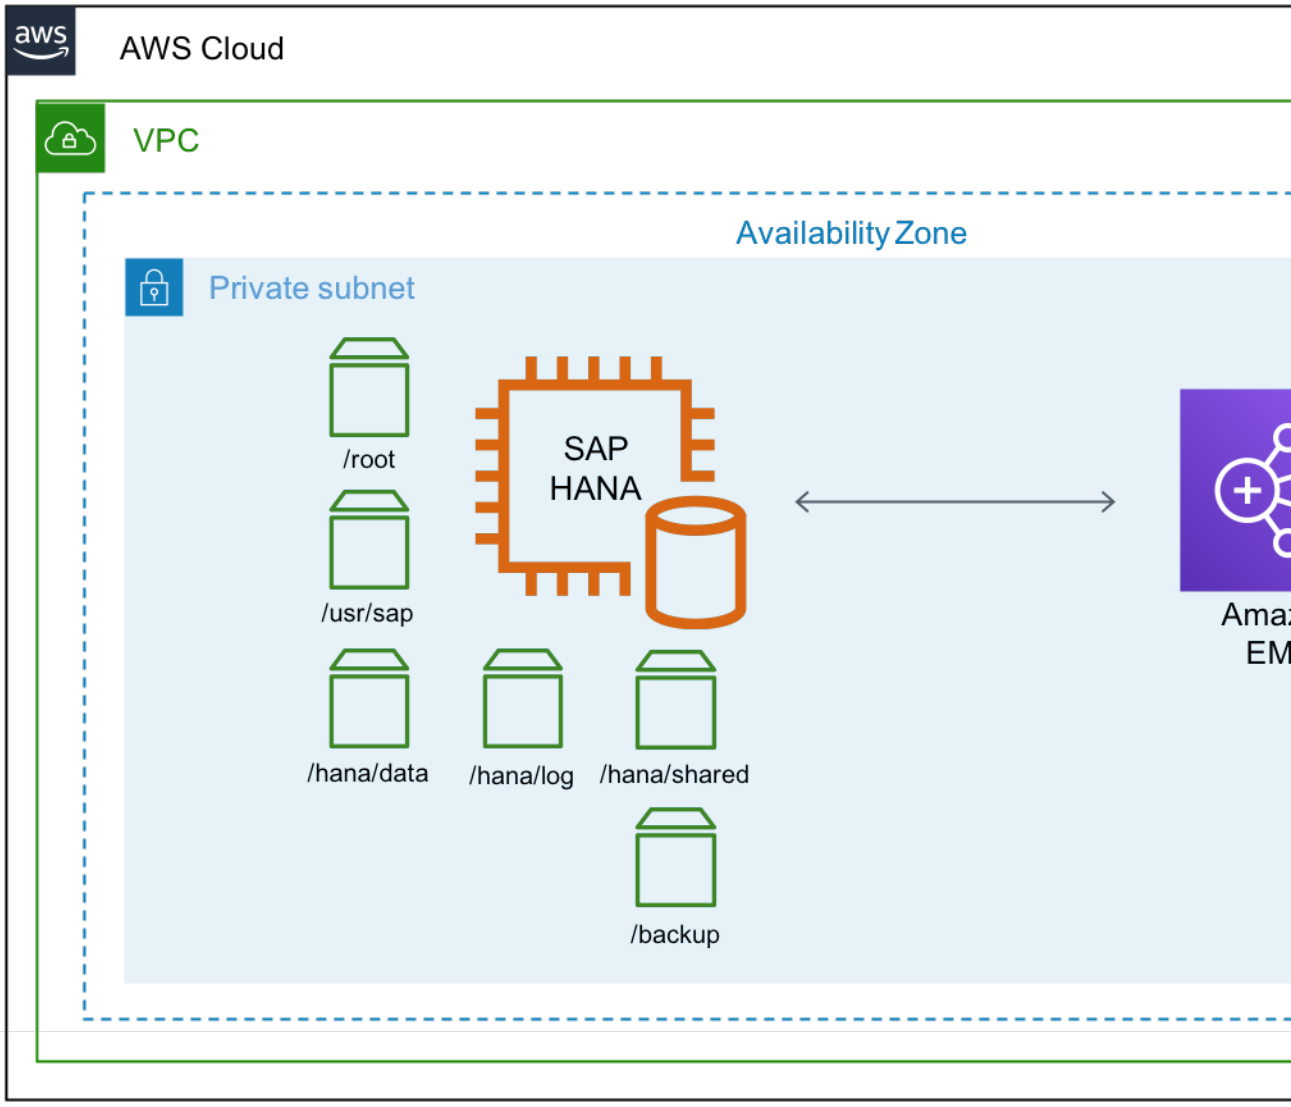

Figure 5: SAP HANA with Amazon EMR for cold tier

Cold Tier Options for SAP BW

For the SAP Business Warehouse (BW) on HANA or SAP BW/4 HANA use cases, you have additional options for cold tier storage.

SAP BW Near Line Storage (NLS) with SAP IQ

With this option, you can use SAP BW Near Line Storage (NLS) with SAP IQ or you can use Data Tiering Optimization (DTO) with SAP IQ to store your cold data. On AWS, you can run your SAP IQ server on Amazon Elastic Compute Cloud (Amazon EC2) instances for the cold tier storage.

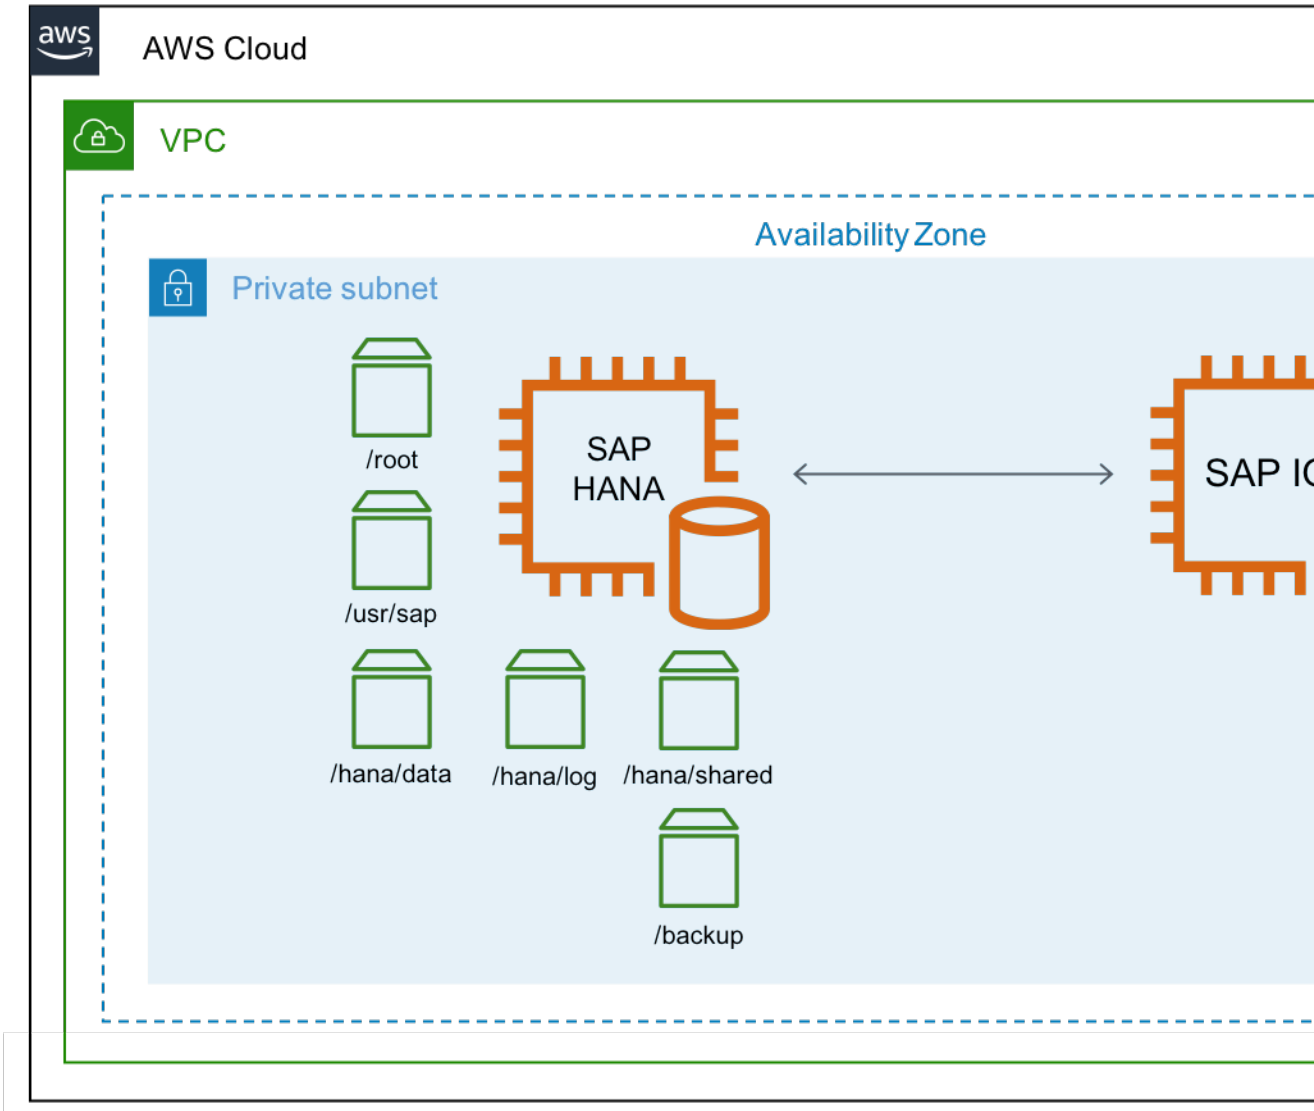

Figure 6: SAP BW NLS with SAP IQ for cold tier

SAP BW NLS with Hadoop

With this option, you can use SAP BW NLS with Apache Hadoop instead of SAP IQ, with this option you can persist your Hadoop data in Amazon S3 using a Hadoop third-party connector for Amazon S3. See Hadoop as a Near-Line Storage Solution documentation from SAP, SAP Note 2363218 – Hadoop NLS: Information, Recommendations and Limitations, and Cloud Data Access documentation from Hortonworks for details.

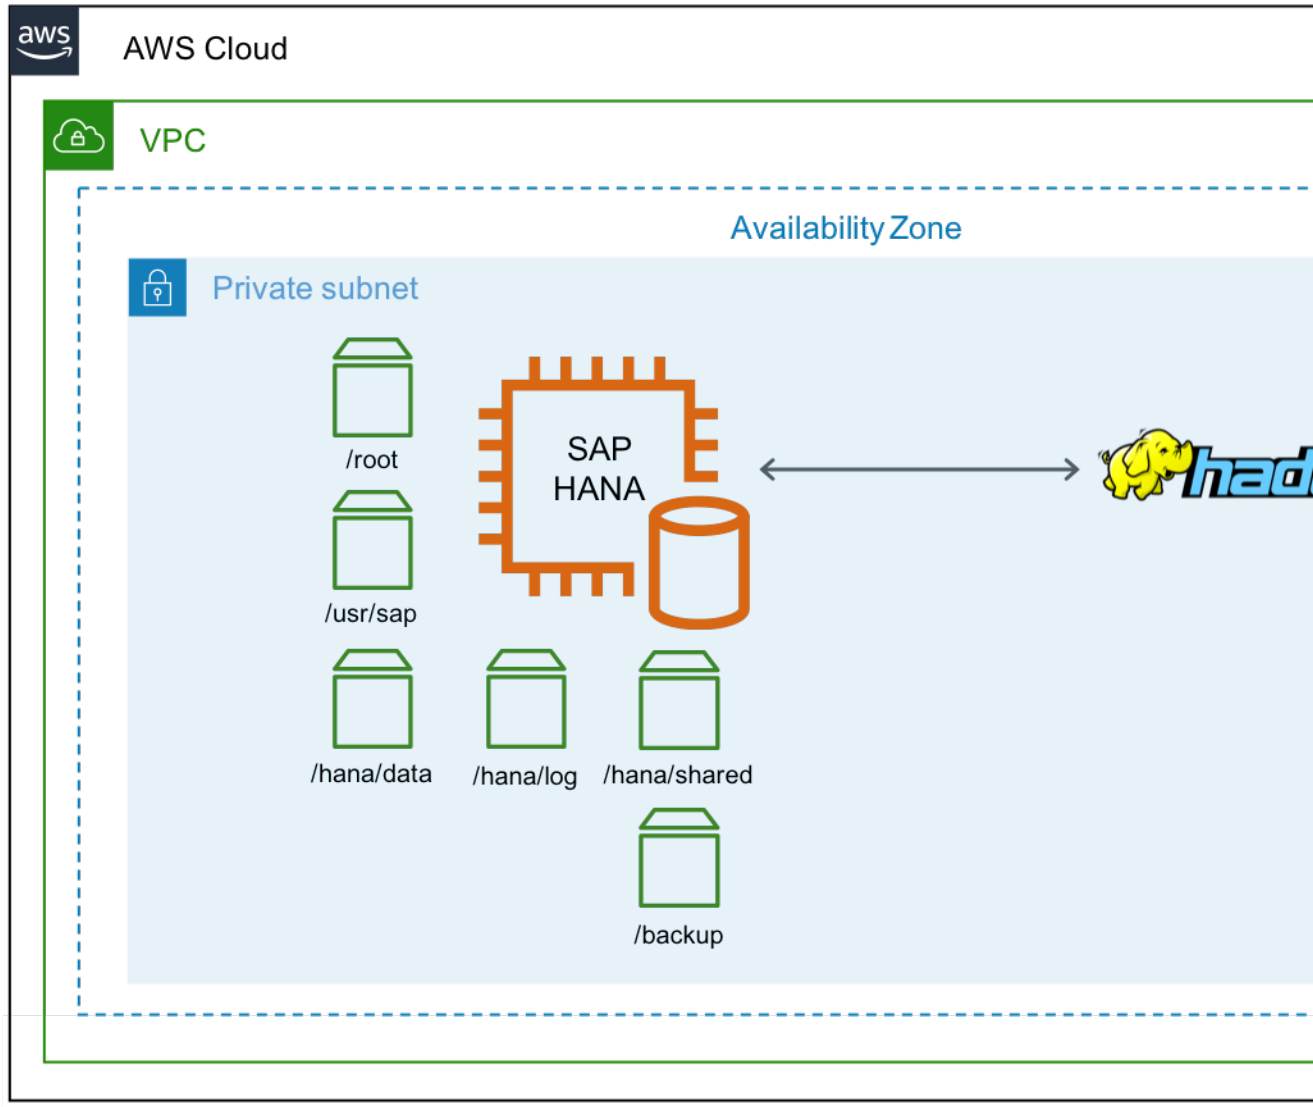

Figure 7: SAP BW NLS with Hadoop for cold tier

SAP BW/4HANA DTO with Data Hub

With this option, you can use DTO with SAP Data Hub to store your cold data in Amazon S3. This option only applies if you use SAP BW/4HANA.

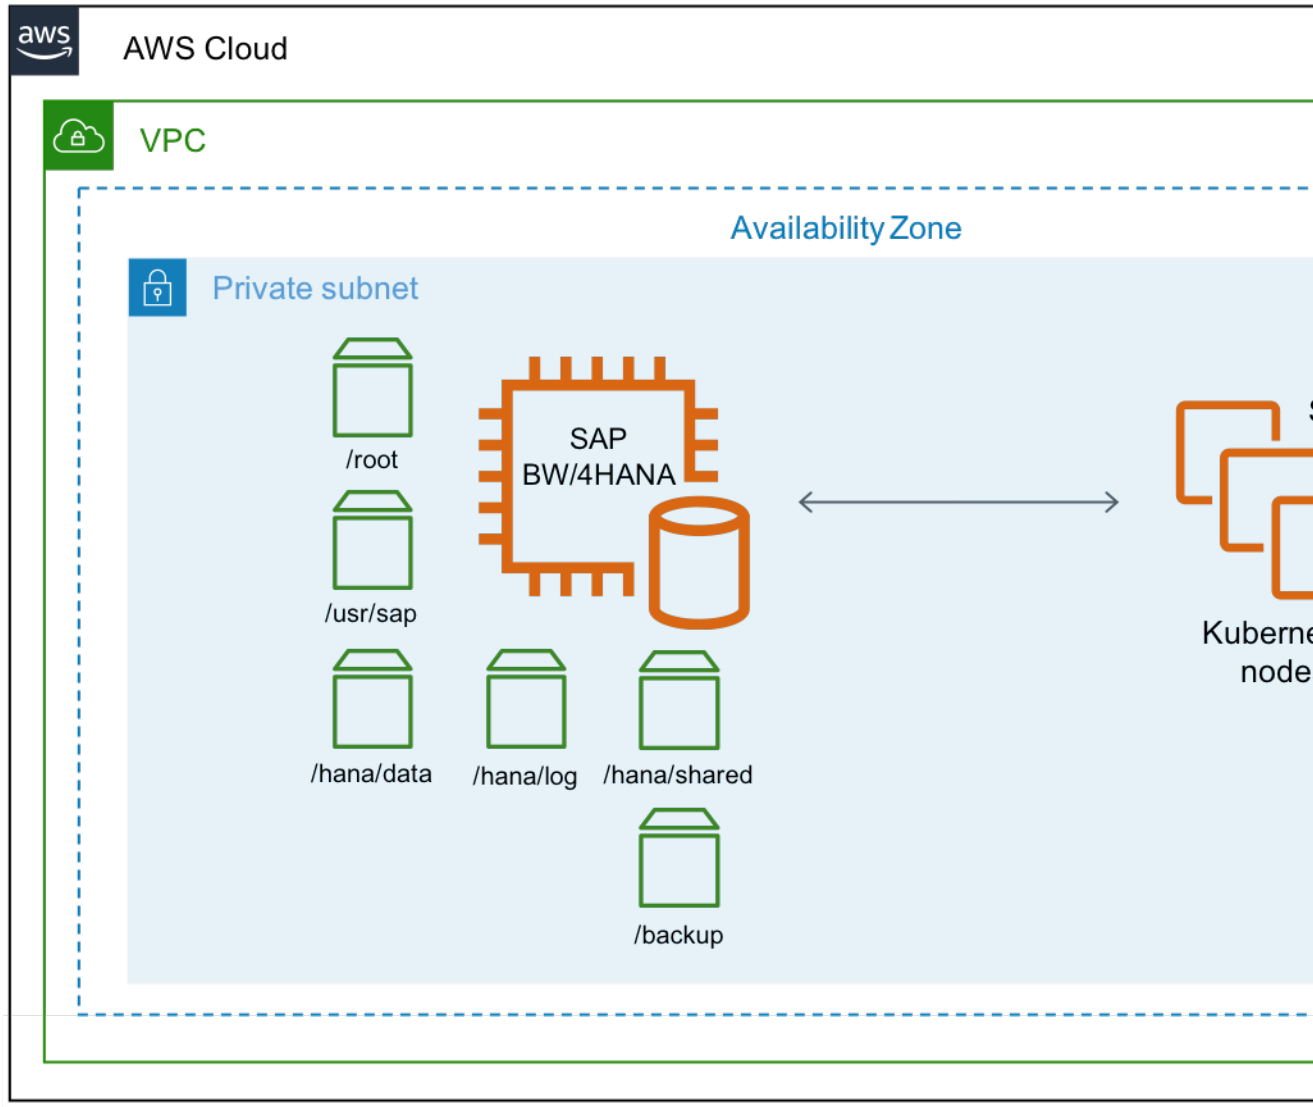

Figure 8: SAP Data Hub on Amazon EKS with BW4/HANA

Cold Tier Options for SAP S/4HANA or Suite on HANA

For S/4HANA or SOH, you can use SAP Information Life Cycle Management (ILM) for the cold data tiering. You have few options with ILM for cold tier. See ILM Store documentation from SAP for details.

SAP ILM with SAP IQ

With this option, you can use ILM with SAP IQ. Similar to the SAP BW NLS with SAP IQ scenario, you can run your SAP IQ server on AWS Amazon EC2 instances to store cold data.

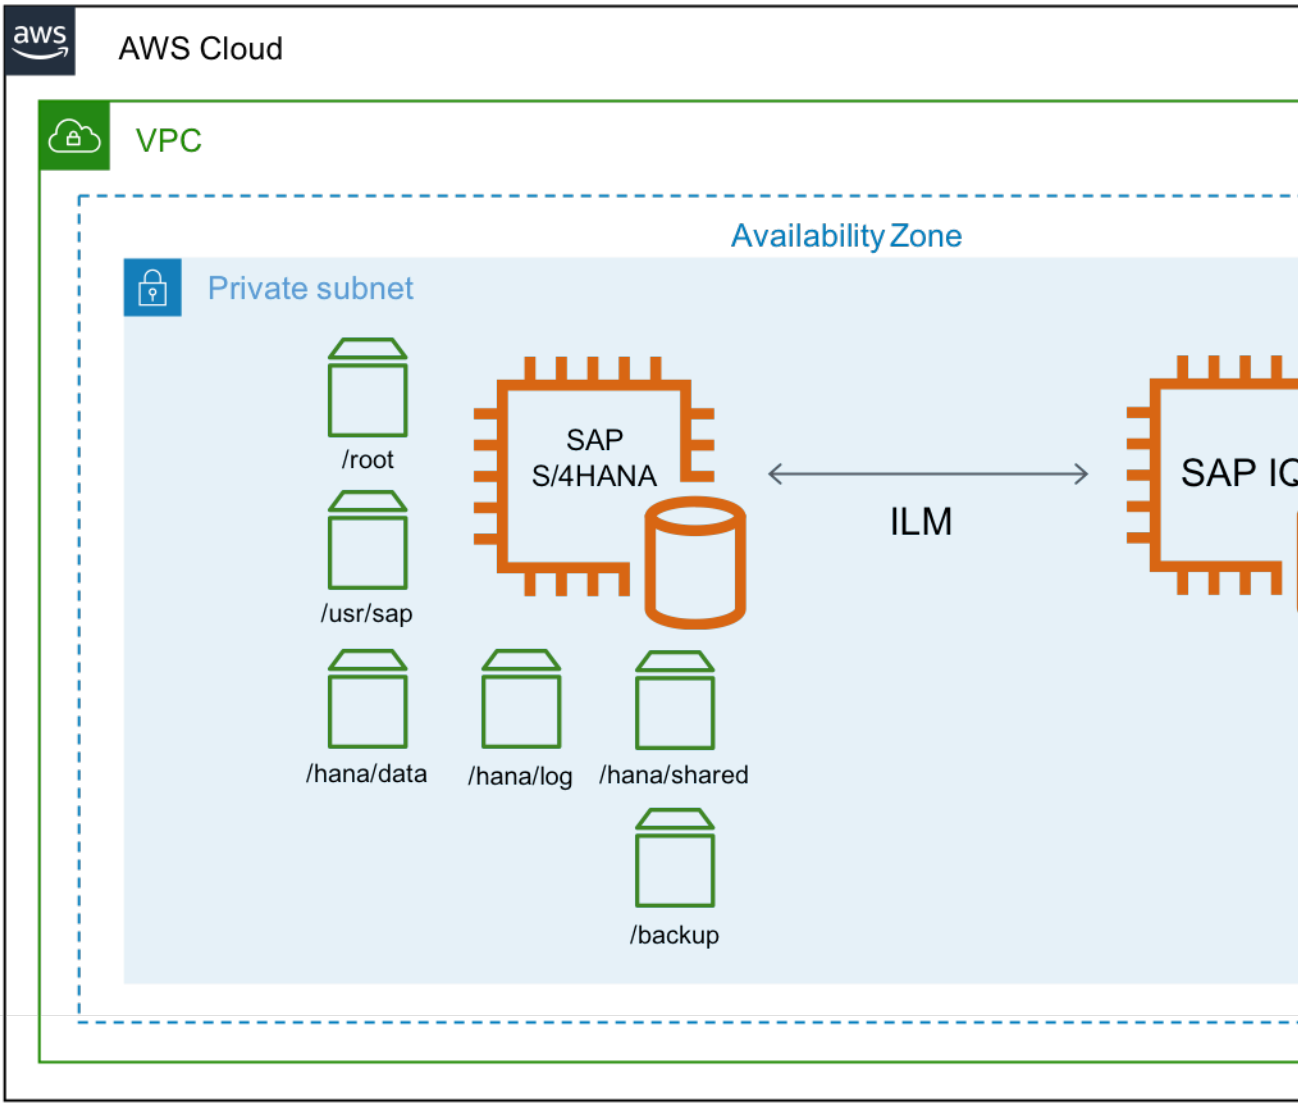

Figure 9: SAP ILM with SAP IQ for cold tier

SAP Archiving

With this option, you can use ILM or your standard data archiving process. You can use Amazon Elastic File System (Amazon EFS) to store your archive file in a highly available, scalable and durable manner. Similarly, for Windows based systems, you can use Amazon FSx to store your archive files. Amazon EFS and Amazon FSx can be mounted as your archive file system and you can archive your data from SAP to this file system through SAP transaction code SARA.

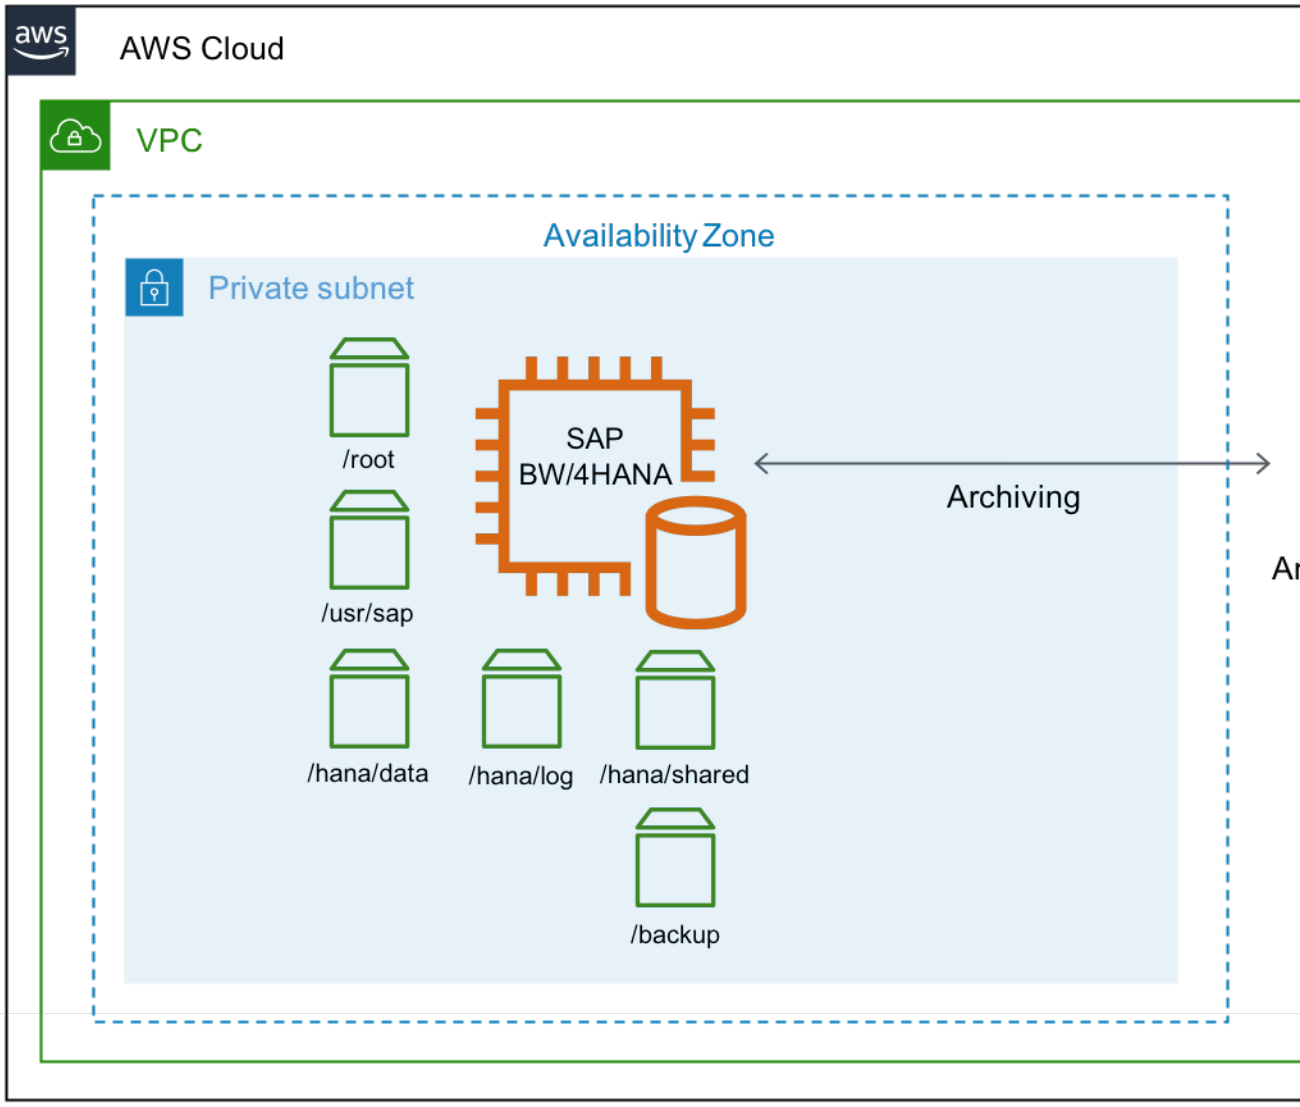

Figure 10: SAP archiving with Amazon EFS for cold tier

For archiving, another option is to use the Amazon Elastic Block Store (Amazon EBS) sc1 volume type as the underlying storage type for your archive file system. Amazon EBS sc1 volumes are inexpensive block storage and are designed for less frequently accessed workloads like data archiving. To increase durability and availability of your archived data, we recommend that you copy the data to Amazon S3 for backup and Amazon S3 Glacier for long term retention.

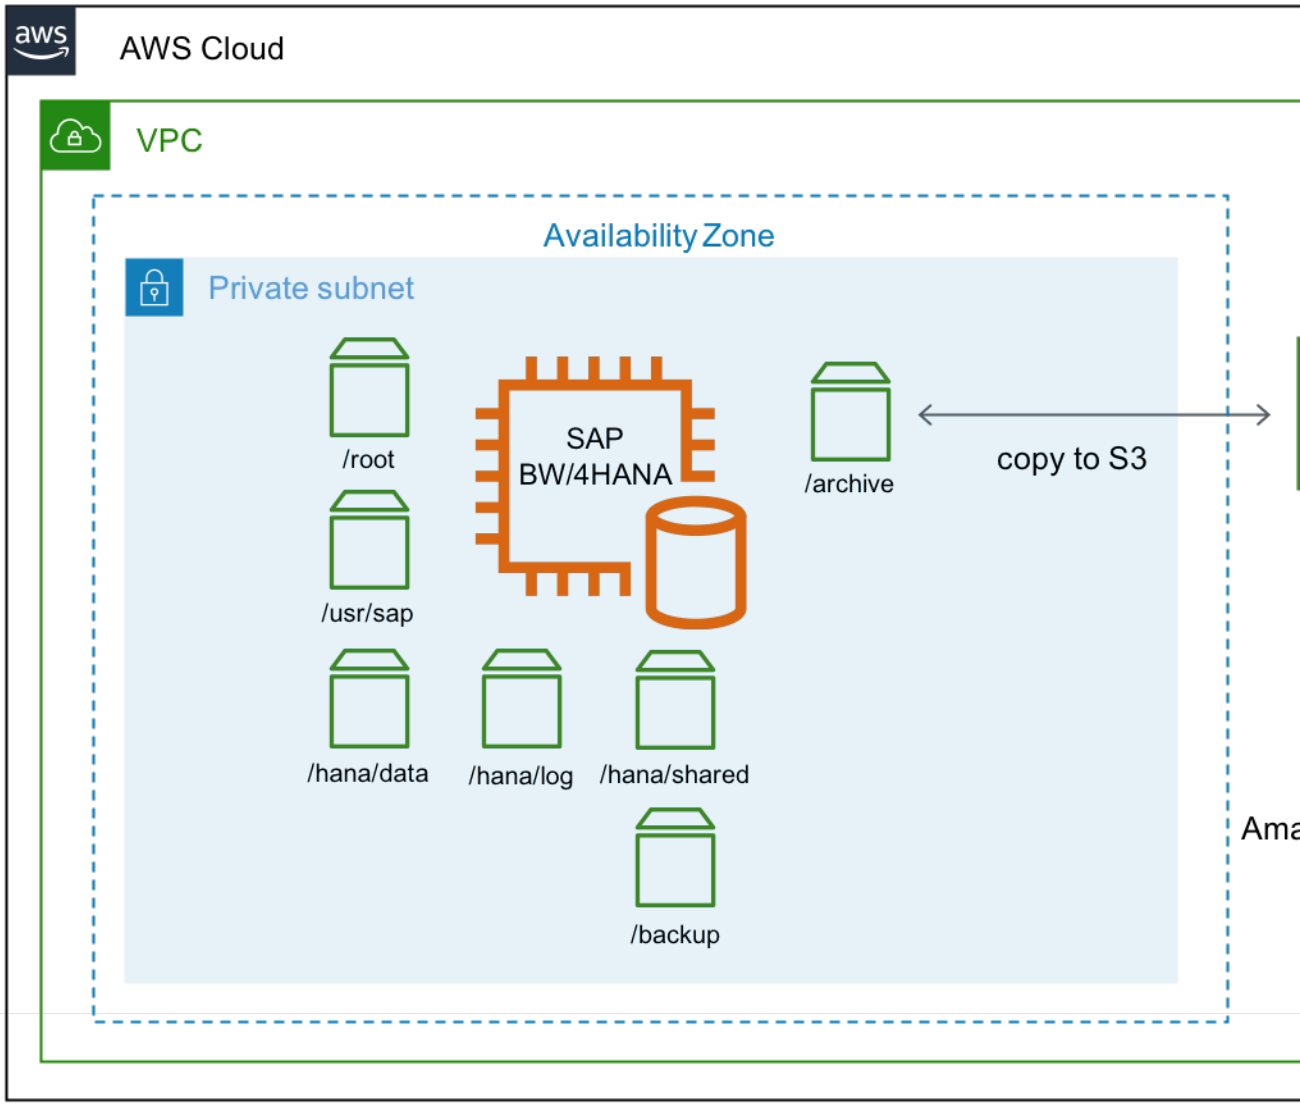

Figure 11: SAP archiving with Amazon EBS for cold tier

Additional Reading

SAP on AWS technical documentation

- SAP HANA on AWS Documentation
- SAP on AWS Technical Documentation

SAP documentation

- SAP Note 1872170 - ABAP on HANA sizing report (S/4HANA, Suite on HANA)
- SAP HANA Extension Nodes as a Warm Store
- SAP HANA Dynamic Tiering Architecture
- Extended Store Table Function Restrictions

- DLM on Amazon Elastic Map Reduce

Document Revisions

Date	Change
July 2019	Initial publication

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.