Migrating SAP HANA to AWS

Patterns for AWS Migrations

September 2018

Rahul Kabra, Harpreet Singh, and Somckit Khemmanivanh
Solutions Architects, Amazon Web Services (AWS)

Table of Contents

Overview

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 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 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.

> **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.

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.
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.

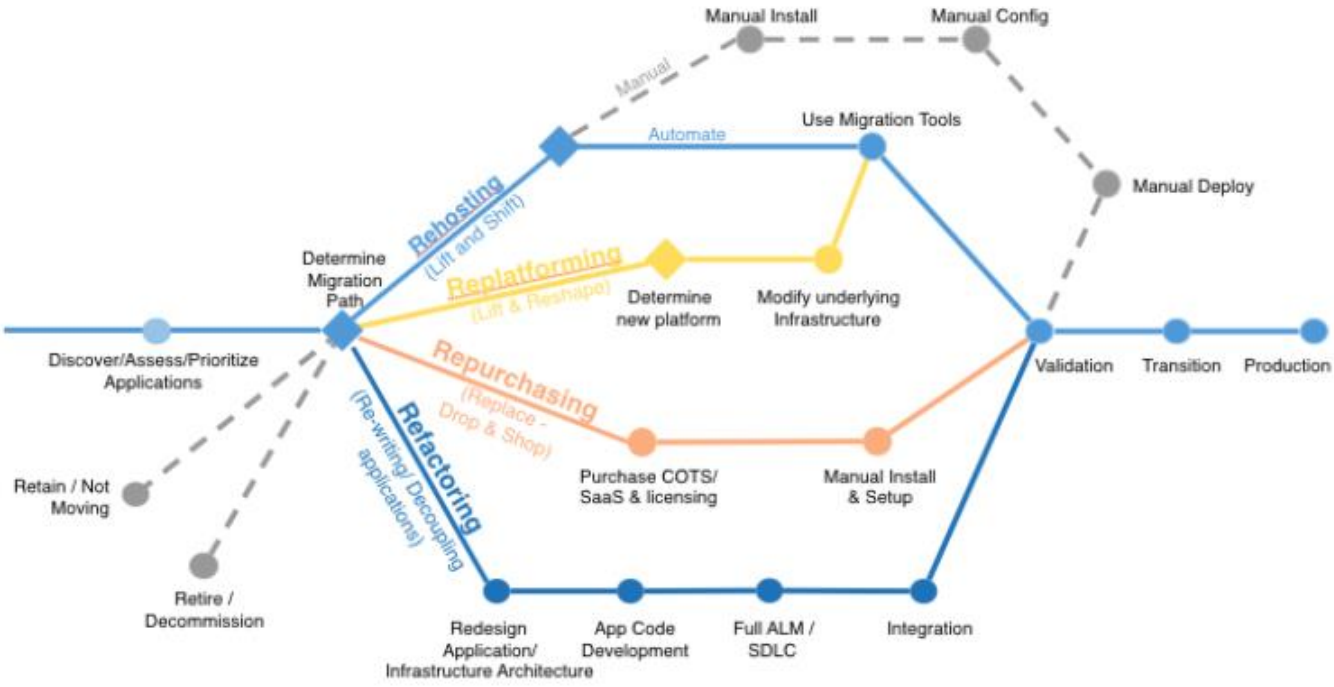

Figure 1: 6 Rs framework

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. 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 NetWeaver application servers. For details, see the Migration Tools and Methodologies 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.

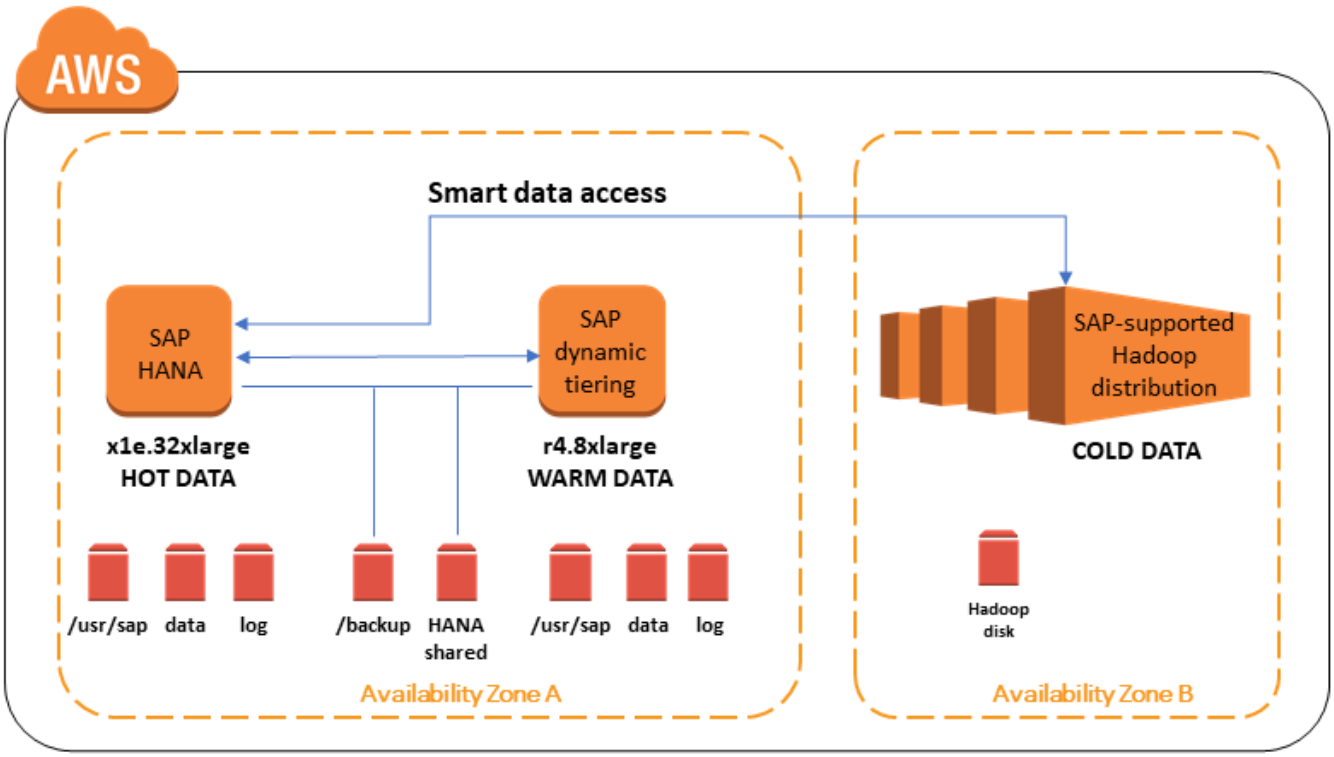

Figure 2: Data tiering

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 HANA on AWS Implementation and Operations Guide
- SAP HANA Quick Start Reference Deployment
- 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 RHEL to 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 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 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 50 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 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:

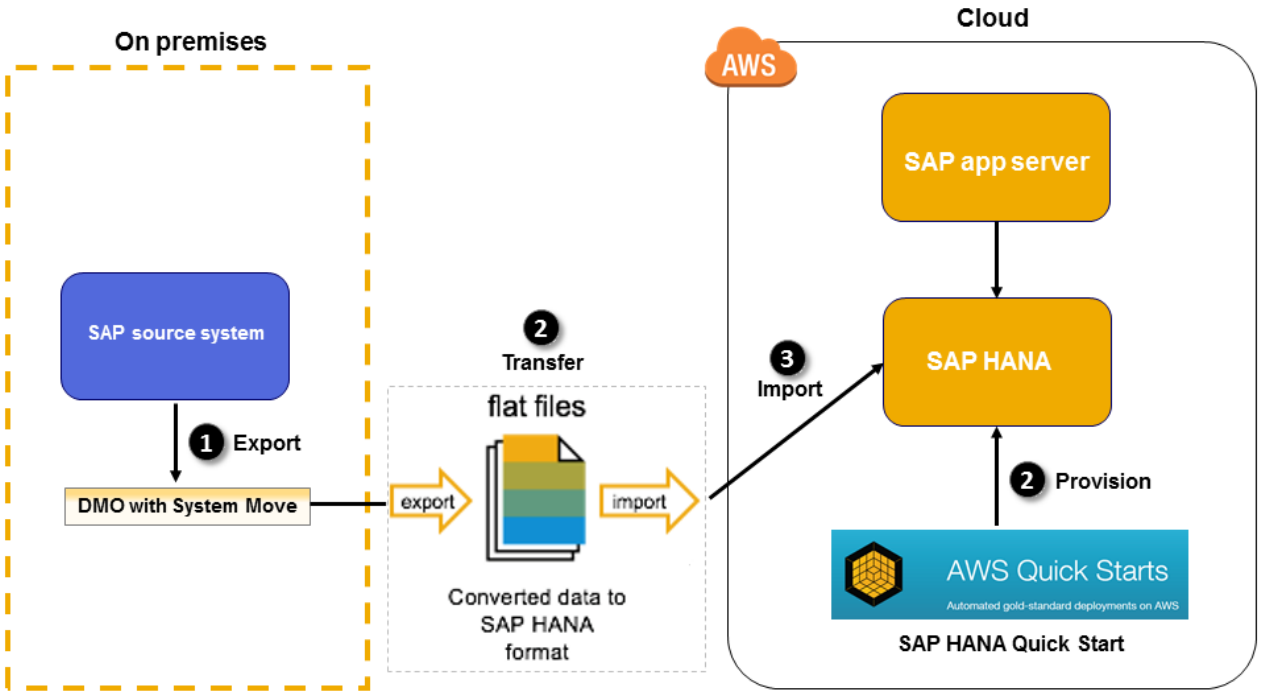

Figure 3: DMO with System Move

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 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.

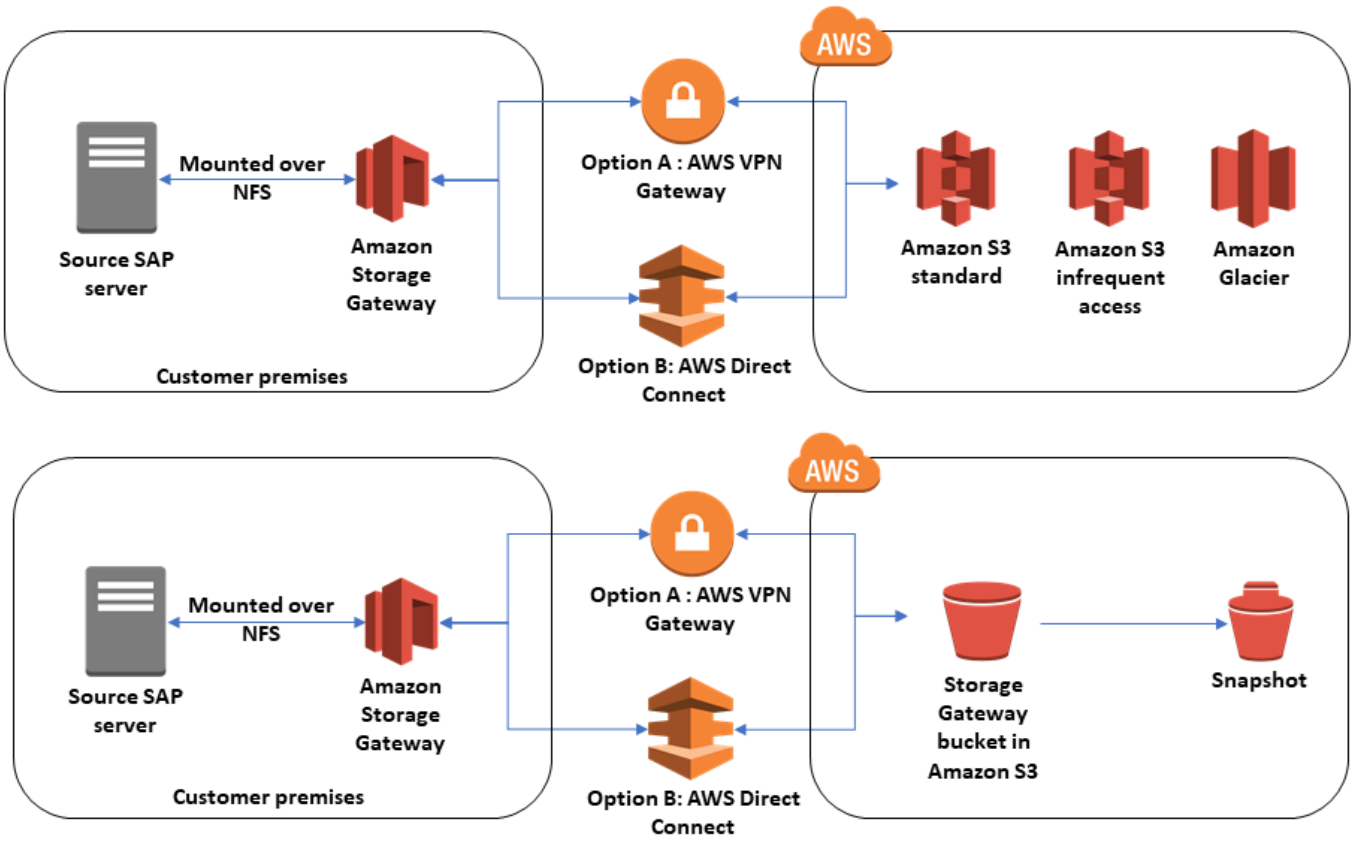

Figure 4: SAP file replication with AWS Storage Gateway

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 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.

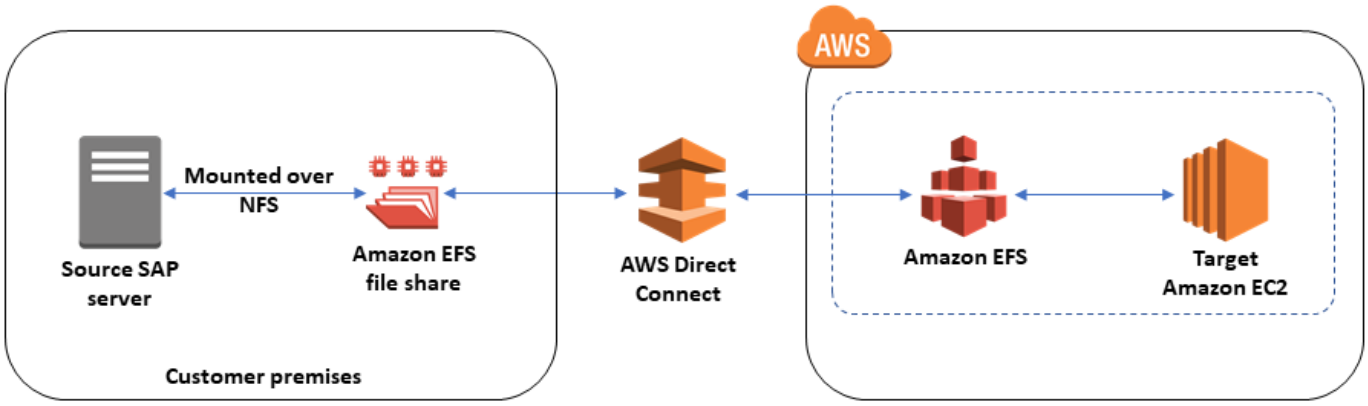

Figure 5: Transferring SAP files with Amazon EFS

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.

Migration Steps

This section covers the migration scenarios summarized in the following table. 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-HANA database such as IBM DB2, Oracle Database, or SQL Server)	HANA	[✓] SAP HANA classical migration [✓] SAP DMO with System Move
Migration of SAP HANA from other platforms to AWS*	HANA (scale-up and scale-out considerations apply here as well)	HANA	[✓] Backup and restore [✓] SAP HANA classical migration (considered a homogeneous system copy in this scenario)** [✓]SAP HANA HSR [✓]SAP HANA HSR with initialization via backup and restore

* 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.) AWS customers can use the SAP SUM DMO tool 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 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 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.

HANA Backup and Restore

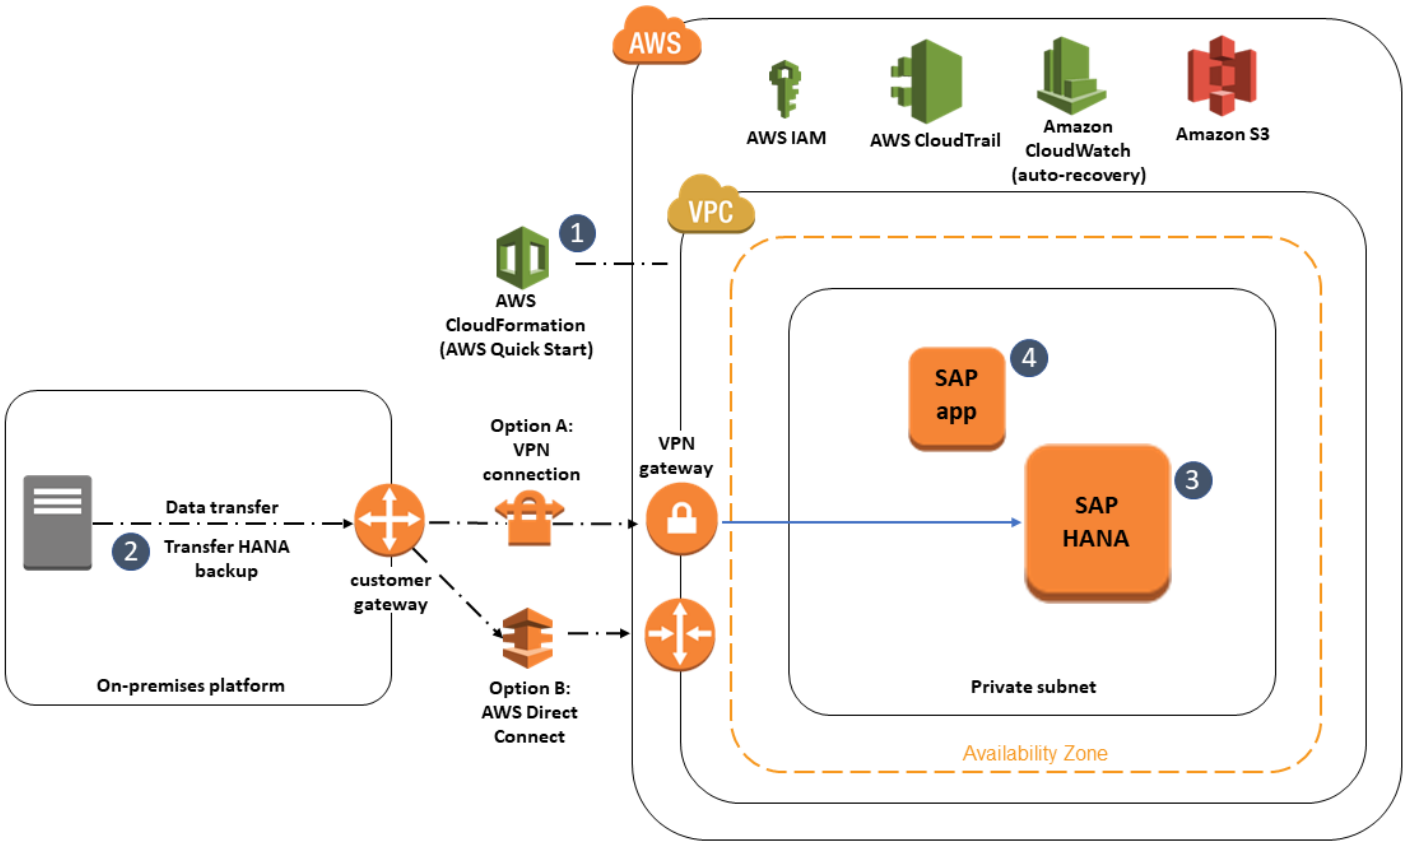

Figure 6: Backup and restore

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 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.

SAP HANA Classical Migration

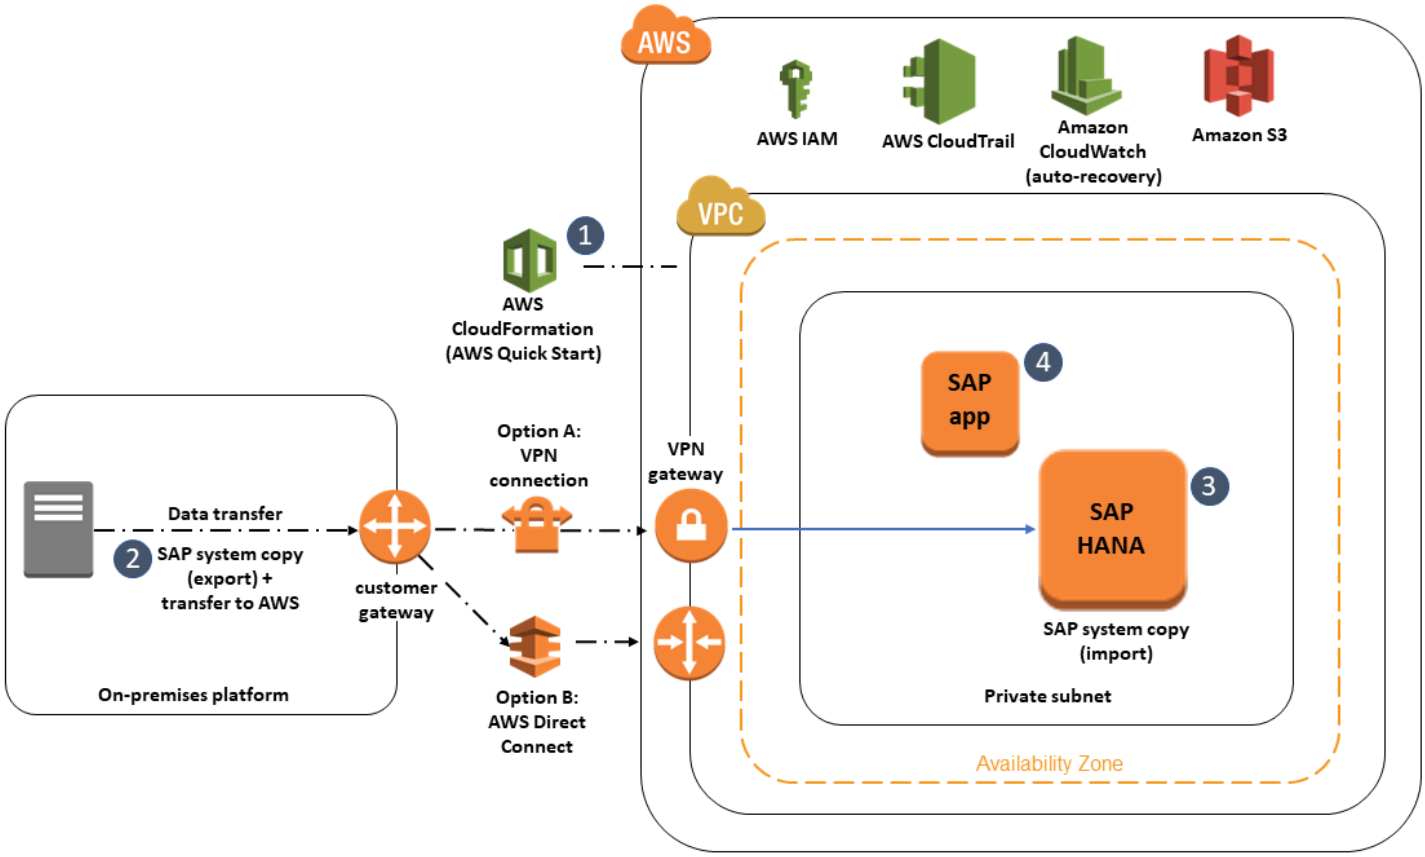

Figure 7: SAP HANA classical migration

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.

SAP HANA HSR

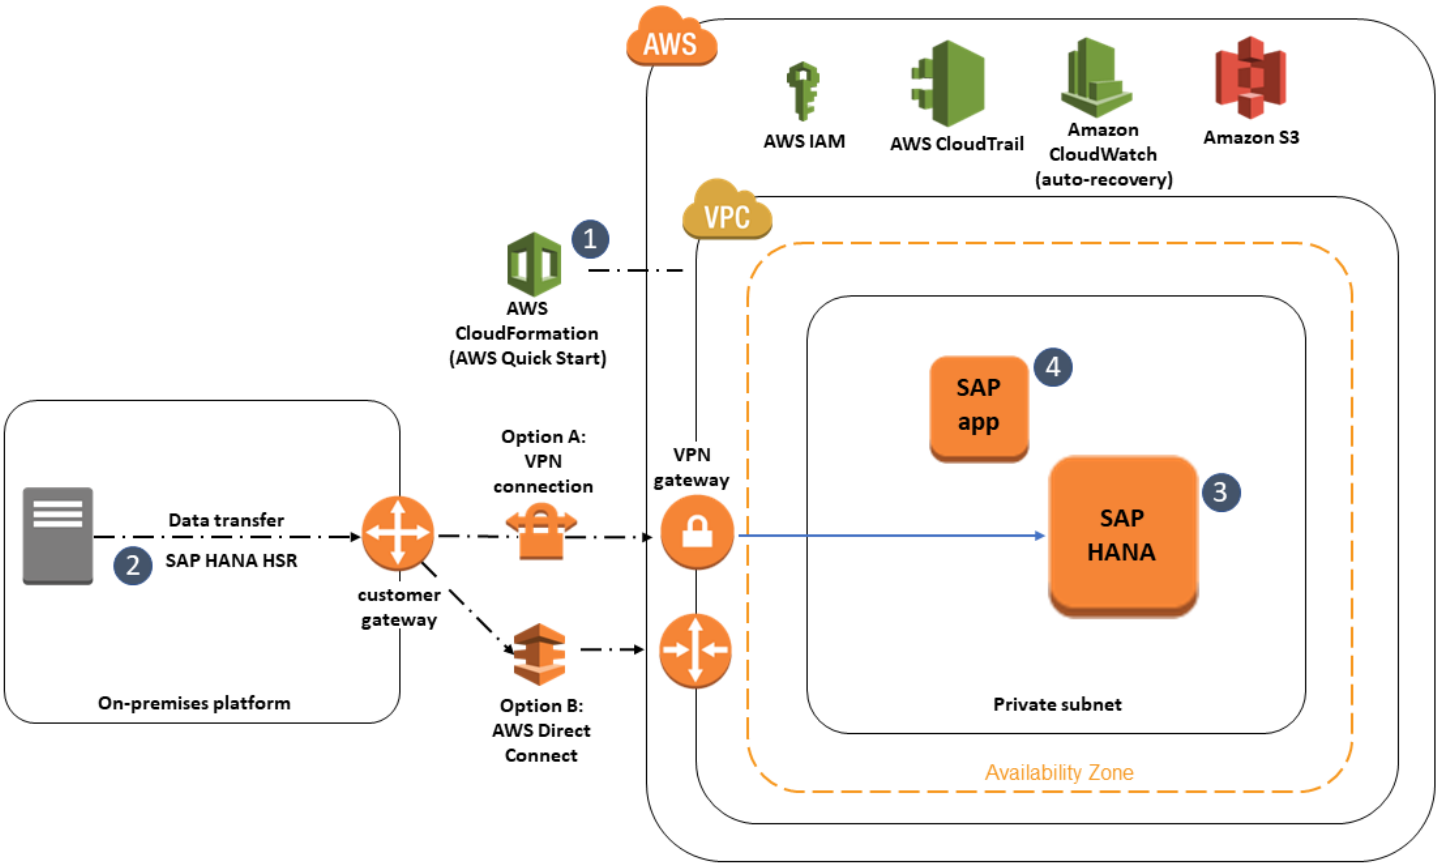

Figure 8: SAP HANA system replication

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.

SAP HANA HSR (with initialization via backup and restore)

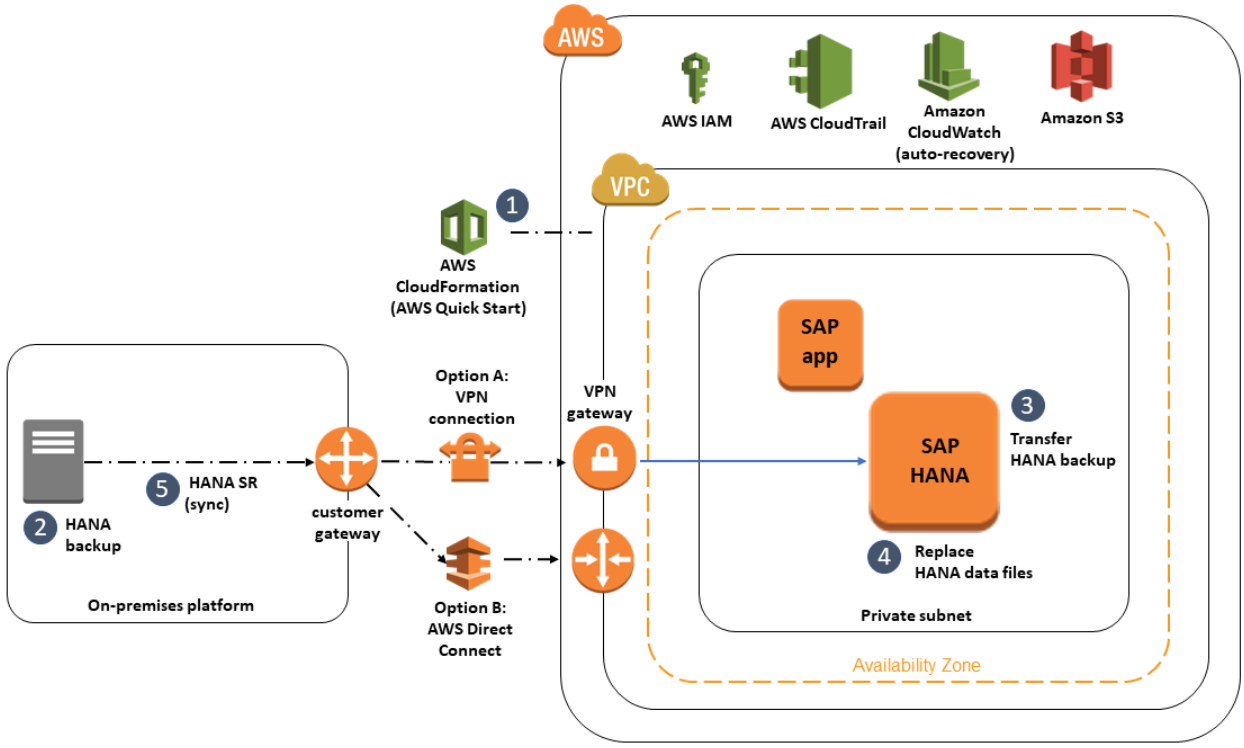

Figure 9: SAP HANA system replication (with initialization via backup and restore)

1. Provision your SAP NetWeaver and HANA systems 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.

Third-Party Migration Tools

If you are interested in using the rehosting option (see the 6 Rs Framework section) for your on-premises SAP HANA environment, you can also leverage third-party continuous data protection (CDP) tools such as Delphix, ATADATA, CloudEndure, 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

Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. 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.

The software included with this paper is 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 code 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.