Migrating SAP HANA Systems to X1 Instances on AWS

Scenarios and Patterns for AWS Migrations

May 2016

Notices

Contents

Abstract

Using the SAP HANA in-memory management system on Amazon Web Services
(AWS) enables many new business and technology capabilities that allow you to
innovate, be more agile, and benefit from the economies of scale and efficiencies
provided by AWS. The availability of the AWS X1 instance type, which offers
1,952 GiB of memory, accelerates and creates new opportunities for using SAP
HANA on AWS. This guide discusses migration options for SAP HANA on AWS
X1, for users who would like to know more about migrating from other AWS
instance types or from other SAP systems, both on premises and in the cloud.

Introduction

AWS has worked closely with SAP to certify the AWS cloud platform for SAP solutions, and offers certified SAP HANA systems, as shown in the certified SAP HANA directory on the SAP website. AWS also offers innovative deployment options that help you deploy SAP workloads quickly and easily on AWS.

The AWS SAP HANA Quick Start deploys and configures SAP HANA in your AWS account by using the AWS CloudFormation technology. Quick Start deployments have been configured to meet SAP HANA key performance indicators (KPIs) and are fully certified by SAP for production systems.

The introduction of the AWS X1 instance type (x1.32xlarge) creates new possibilities and options for deploying and implementing SAP HANA on AWS. The X1 instance type provides 1,952 GiB of memory on a single instance. This extra memory capacity allows for a greater number of consolidation options, workloads, and deployment combinations. X1 instances are powered by quad-socket 2.3-GHz Intel Xeon E7-8880 v3 (Haswell) processors offering up to 128 vCPUs, high-memory bandwidth, and larger L3 caches, and are optimized to boost the performance and reliability of in-memory applications. X1 instances also offer up to 10 Gbps of dedicated throughput to Amazon Elastic Block Store (Amazon EBS), and are EBS-optimized by default and at no additional cost. To learn more about X1, see Amazon EC2 X1 Instances and the Amazon EC2 User Guide.

This guide describes the most common scenarios, use cases, and migration options for X1, including migrating SAP systems from another platform (on-premises or cloud) to AWS, and migrating an existing SAP HANA instance on AWS to X1. For both scenarios, this guide covers SAP HANA scale-up and scale-out deployments. More specific scenarios are discussed in the next section.

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 Further 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.

AWS Migration Scenarios

This guide covers four migration scenarios, as summarized in the following table. The tools and methodologies listed in the table are discussed in the next section.

Migration scenario	Source database	Target database	Migration tool or methodology
Migration within AWS	SAP HANA scale-up*	SAP HANA scale-up*	[✓] Backup and restore [✓] SAP HANA classical migration (considered a homogeneous system copy in this scenario)** [✓] SAP HANA system replication (HSR) [✓] Amazon EC2 instance type resizing
Migration within AWS	SAP HANA scale-out*	SAP HANA scale-up*	[✓] Backup and restore [✓] SAP HANA classical migration (considered a homogeneous system copy in this scenario)**
Migration from other platforms to AWS*	*anyDB* (any non-HANA database such as DB2, Oracle, or SQL Server)	SAP HANA	[✓] SAP HANA classical migration [✓] SUM DMO
Migration from other platforms to AWS*	SAP HANA (scale-up and scale-out considerations apply here as well)	SAP HANA	[✓] Backup and restore [✓] SAP HANA classical migration (considered a homogeneous system copy in this scenario)** [✓]SAP HANA system replication (HSR)

* Scale-up and scale-out are two ways to increase the capacity of your system on AWS, by changing the instance type (scale-up) and by adding instances (scale-out). These options are discussed in detail in the Scale-up and Scale-out Sizing section.

** See SAP Note 1844468 – Homogeneous system copy on SAP HANA (SAP support portal login is required).

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

AWS and SAP Migration Tools and Methodologies

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

- **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. For details, see Classical Migration to SAP HANA on the SAP website.

- **SAP Software Update Manager (SUM) Database Migration Option (DMO)**

  SAP offers the SUM DMO approach as a one-step migration option from other database systems to SAP HANA. This option uses the SAP DMO process and tool to automate multiple required migration steps. For details, see Database Migration Option (DMO) of SUM - Introduction on the SAP website.

- **Backup and restore**

  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, including AWS Import/Export Snowball, AWS Storage Gateway, Amazon Simple Storage Service (Amazon S3), and Amazon Elastic Block Store (Amazon EBS). For more information, see the Storage Options in the AWS Cloud whitepaper.

- **SAP HANA System Replication (HSR)**

  SAP HANA 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.

- **Amazon EC2 instance type resizing**

    You can resize an Amazon Elastic Compute Cloud (Amazon EC2) instance by changing its instance type. Each instance type provides different compute, memory, and storage capabilities. By changing the instance type, you can increase or decrease the amount of resources to suit the computing needs of your SAP HANA system. For more information, see Resizing Your Instance in the *Amazon EC2 User Guide*.

Prerequisites for Migration

Specialized Knowledge

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

- 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 Answers
- Amazon Virtual Private Cloud (Amazon VPC)
- Amazon EC2
- Amazon EBS
- Amazon S3
- 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
- High Availability and Disaster Recovery Options for SAP HANA on AWS

SAP HANA Sizing

For general information about sizing SAP systems on AWS, see the Capacity Planning for SAP Systems on AWS whitepaper.

If you're migrating to an SAP HANA system, see the following SAP notes for HANA sizing considerations (SAP support portal login is required):

- 706478 – Preventing Basis tables from increasing considerably
- 1514966 – SAP HANA 1.0: Sizing SAP In-Memory Database
- 1637145 – SAP BW on HANA: Sizing SAP In-Memory Database
- 1736976 – Sizing Report for BW on HANA
- 1793345 – Sizing for SAP Suite on HANA
- 1855041 – Sizing Recommendation for Master Node in BW-on-HANA
- 1702409 – HANA DB: Optimal number of scale out nodes for BW on HANA
- 1872170 – Suite on HANA - S/4 sizing report

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 Amazon 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 Amazon EC2 instance type.

Note   Only production SAP HANA systems need to run on certified configurations that meet SAP HANA KPIs. SAP provides more flexibility when running SAP HANA non-production systems. For more information, see the SAP HANA TDI – FAQ on the SAP website.

Network Sizing

In addition to instance sizing (see the Amazon EC2 Instance Types page for instance capacity limits), you will need to consider network sizing to plan for data transfer into AWS. As an alternative, you can use the AWS Import/Export Snowball service for a physical (non-network) transport of data into AWS.
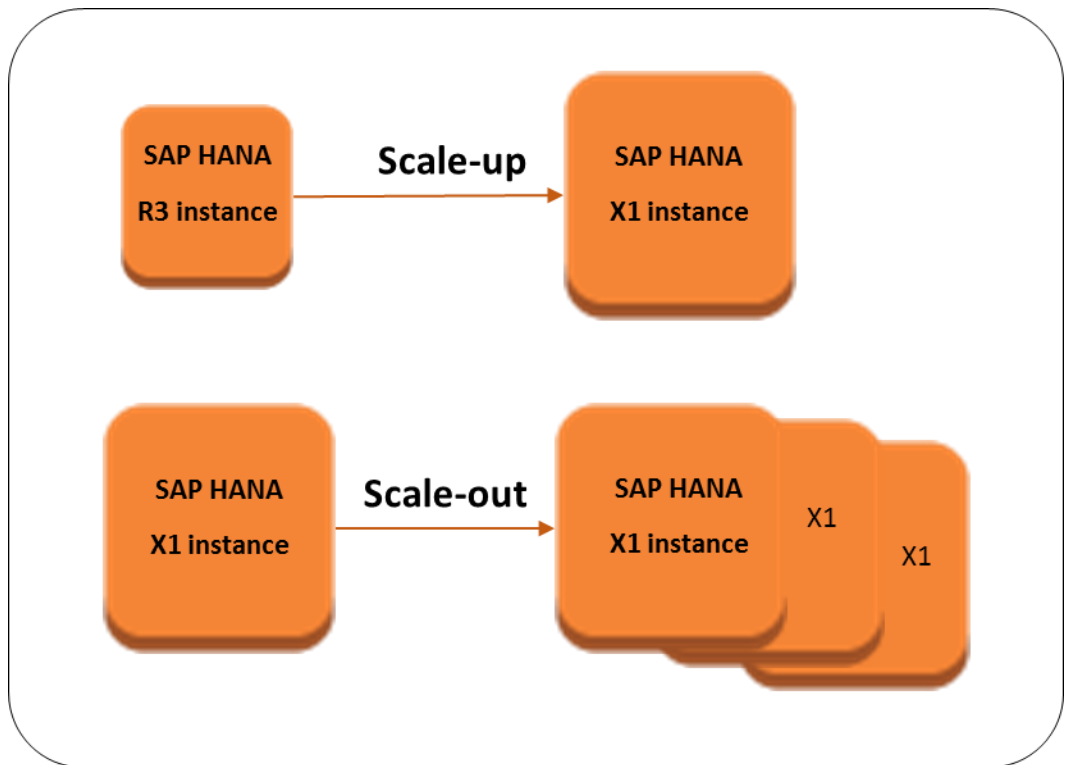
Scale-up and Scale-out Sizing

The release of the AWS X1 instance type increases the options available to you for SAP HANA scale-up and scale-out deployments. When you *scale up*, you utilize the computing capacity of a single SAP HANA node by switching it to an EC2 instance type that provides a larger capacity. For example, if you're using an R3 instance type, and it becomes too small for your workload, you can change it to an X1 instance. 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 are able to scale up to keep pace with change.

In addition to SAP HANA sizing, you may also need to size your SAP application tier. For the SAP Application Performance Standard (SAPS) ratings for the X1 instance, see SAP Standard Application Benchmarks in Cloud Environments and the SAP on AWS support note (SAP support portal login required).

When you *scale out*, you add SAP HANA capacity by adding nodes (new instances) to the SAP HANA cluster. For example, once you have reached the maximum memory capacity of a single EC2 instance (1,952 GiB of RAM on X1), you may decide to scale out your SAP HANA cluster and add more instances. The more likely scenario is that your sizing has revealed the need to plan for a scale-out configuration from the start. AWS gives you the option of deploying SAP HANA scale-out configurations when you use the SAP HANA Quick Start.

Figure 1 illustrates scale-up and scale-out sizing.

Figure 1: Scale-up and Scale-out Sizing

After your SAP sizing and SAP HANA deployment models have been finalized, you can plan your migration strategy.

Migrating to X1 on AWS

Migrating from an existing EC2 instance to an X1 instance type is straightforward. In the following table, migration scenarios are listed by source system configuration and scale-up or scale-out option. The scenarios are discussed in subsequent sections. Please refer to the section for the scenario that best fits your current configuration.

Scenario		Source	Target	Example
1	Scale-up	One existing instance (SAP HANA scale-up)	One X1 instance (SAP HANA scale-up)	Migrating one R3 instance to one X1 instance
2	Scale-out	Multiple existing instances (SAP HANA scale-out)	One X1 instance (SAP HANA scale-up)	Migrating multiple R3 instances to one X1 instance

Scenario		Source	Target	Example
3	Scale-out	Multiple existing instances (SAP HANA scale-out)	Multiple X1 instances (SAP HANA scale-out)	Migrating multiple R3 instances to multiple X1 instances

Scale-up to X1 Instance

You can migrate your existing EC2 instance (for example, r3.8xlarge) to an X1 instance with these general steps. (Additional steps may be required depending on your use case.)

1. Back up your source SAP HANA database. Detach your Amazon EBS volume that holds the HANA backup files.

2. Provision your SAP HANA system and landscape on AWS. ([AWS Quick Start templates](#) can help expedite and automate this process for you.) We recommend that you keep the same virtual name on the X1 instance as on your existing system.

3. Attach your Amazon EBS volume with backup files to your X1 instance, and recover your database.

4. Update the DNS record and hosts file of your client applications with the new IP address of the X1 instance.

5. Replicate any SAP HANA configurations to target the X1 instance.

6. Apply a new license if necessary.

7. Start your SAP application servers.

Scale-out to X1 Instance

For scale-out configurations, we will cover two possible scenarios:

- Migrate your existing EC2 scale-out instances to one X1 scale-up instance.

- Migrate your existing EC2 scale-out instances to multiple X1 scale-out instances.

Migrating Existing EC2 Instances to One X1 Instance

To migrate an existing EC2 instance (for example, r3.8xlarge) in a scale-out configuration to one X1 instance, follow these steps:

1. Calculate the total memory required. For example, if your current scale-out configuration is three r3.8xlarge instances, with each instance providing 244 GiB of RAM, your current total memory is 732 GiB. In this scenario, you would choose one X1 instance, which offers up to 1,952 GiB of memory.

2. Back up your source SAP HANA database. Detach your backup EBS volume.

3. Provision your SAP HANA system and landscape on AWS. (AWS Quick Start templates can help expedite and automate this process for you.) We recommend that you keep the same virtual name on the X1 instance as on your existing system.

4. Attach your backup EBS volume to your X1 instance and recover your database. For details and prerequisites (such as adding required SAP HANA services), see SAP Note 2093572 – SAP HANA Migration from Multi-Node to Single-Node.

5. Update the DNS record and hosts file of your client applications with the new IP address of the X1 instance.

6. Replicate any SAP HANA configurations to target the X1 instance.

7. Apply a new license if necessary.

8. Start your SAP application servers.

Migrating Existing EC2 Instances to Multiple X1 Instances

To migrate an existing EC2 instance (for example, r3.8xlarge) in a scale-out configuration to a scale-out X1 configuration, follow these steps:

1. Calculate the total memory required. For example, if your current scale-out configuration is ten r3.8xlarge instances, with each instance providing 244 GiB of RAM, your current total memory is 2,440 GiB. In this scenario, you would choose an X1 scale-out configuration with at least two x1.32xlarge

nodes. Depending on your SAP workload, three SAP HANA nodes may work best; please refer to the SAP HANA Sizing section previously in this document for details.

2. Back up your source SAP HANA database. Detach your backup EBS volume.

3. Provision your AWS SAP HANA system and landscape. (AWS Quick Start templates can help expedite and automate this process for you.) We recommend that you keep the same virtual name on the X1 instance as on your existing system.

4. Attach your backup EBS volume to your X1 instance and recover your database. For details and prerequisites (such as adding required SAP HANA services), see SAP Note 2093572 – SAP HANA Migration from Multi-Node to Single-Node.

5. Start your SAP application servers.

Migrating to X1 on AWS from a Different Platform

Migrating from another SAP platform (for example, an on-premises infrastructure or cloud platform outside of AWS) involves two major migration scenarios:

- Migrate from *anyDB* to an SAP HANA database on AWS. *AnyDB* refers to non-HANA databases such as DB2, Oracle, or SQL Server.

- Migrate from an SAP HANA database on another platform to an SAP HANA database on AWS.

These two scenarios are discussed in the following sections.

Migrating *AnyDB* to SAP HANA on AWS X1

Migrating from *anyDB* to HANA typically involves operating system and database platform 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. (However, note that the SAP HANA solution on AWS automatically incorporates the best practices from SAP Note 2205917 – SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12.)

- 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 AWS and SAP Migration Tools and Methodologies.) AWS customers can utilize the SAP SUM DMO tool to migrate their database to SAP HANA on AWS. Some considerations for the SAP SUM DMO method are the network bandwidth, amount of data to be transferred, and the amount of time available for the data to be transferred. Figure 2 shows the generally recommended migration option for this scenario.
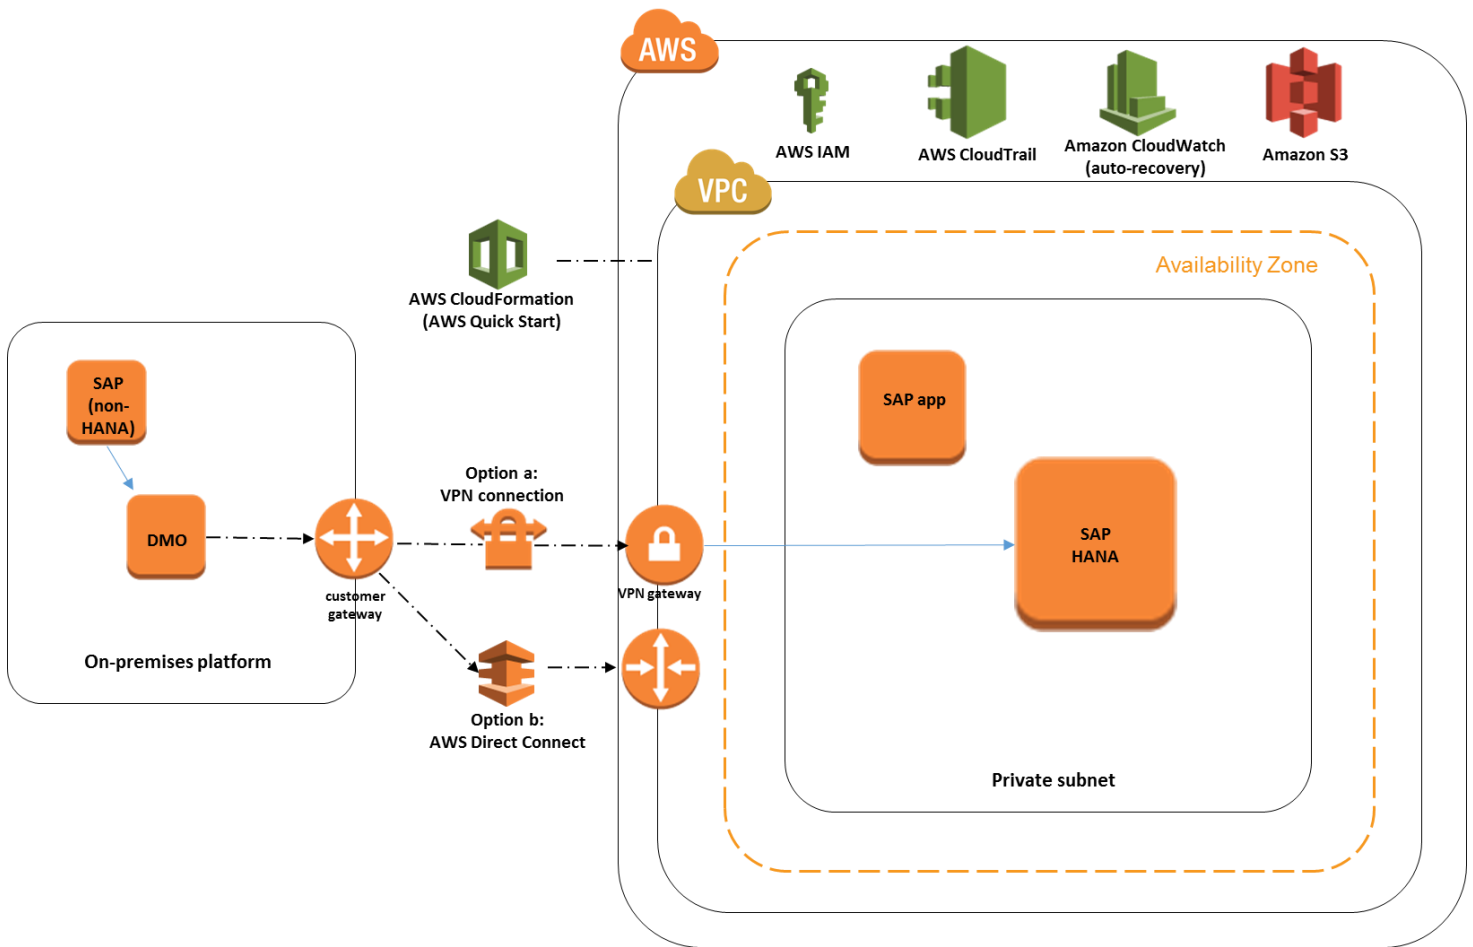
Figure 2: Migrating from Other Databases to SAP HANA on AWS

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. 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. Also, the SAP HANA system on AWS comes
equipped with capabilities like SAP HANA system replication, which supports
jumbo frames and Multi-AZ (Availability Zone) deployments.

Migrating SAP HANA to AWS X1

This scenario is straightforward because you're already using SAP HANA. For this migration, you need to map your existing SAP HANA systems and sizing on a different platform to SAP HANA solutions on AWS.

With X1's memory capabilities, you also have the option to consolidate multiple SAP HANA databases on one (scale-up) or multiple (scale-out) X1 instances. SAP calls these options HANA and ABAP One Server, Multiple Components One Database (MCOD), Multiple Components One System (MCOS), and Multitenant Database Containers (MDC). It is beyond the scope of this paper to recommend specific consolidation combinations; for more information, see the [SAP documentation](#).

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. Depending on your specific scenario, you might use standard backup and restore, SAP HANA classical migration, or SAP HANA HSR for the data transfer; see the following sections for details on each option.
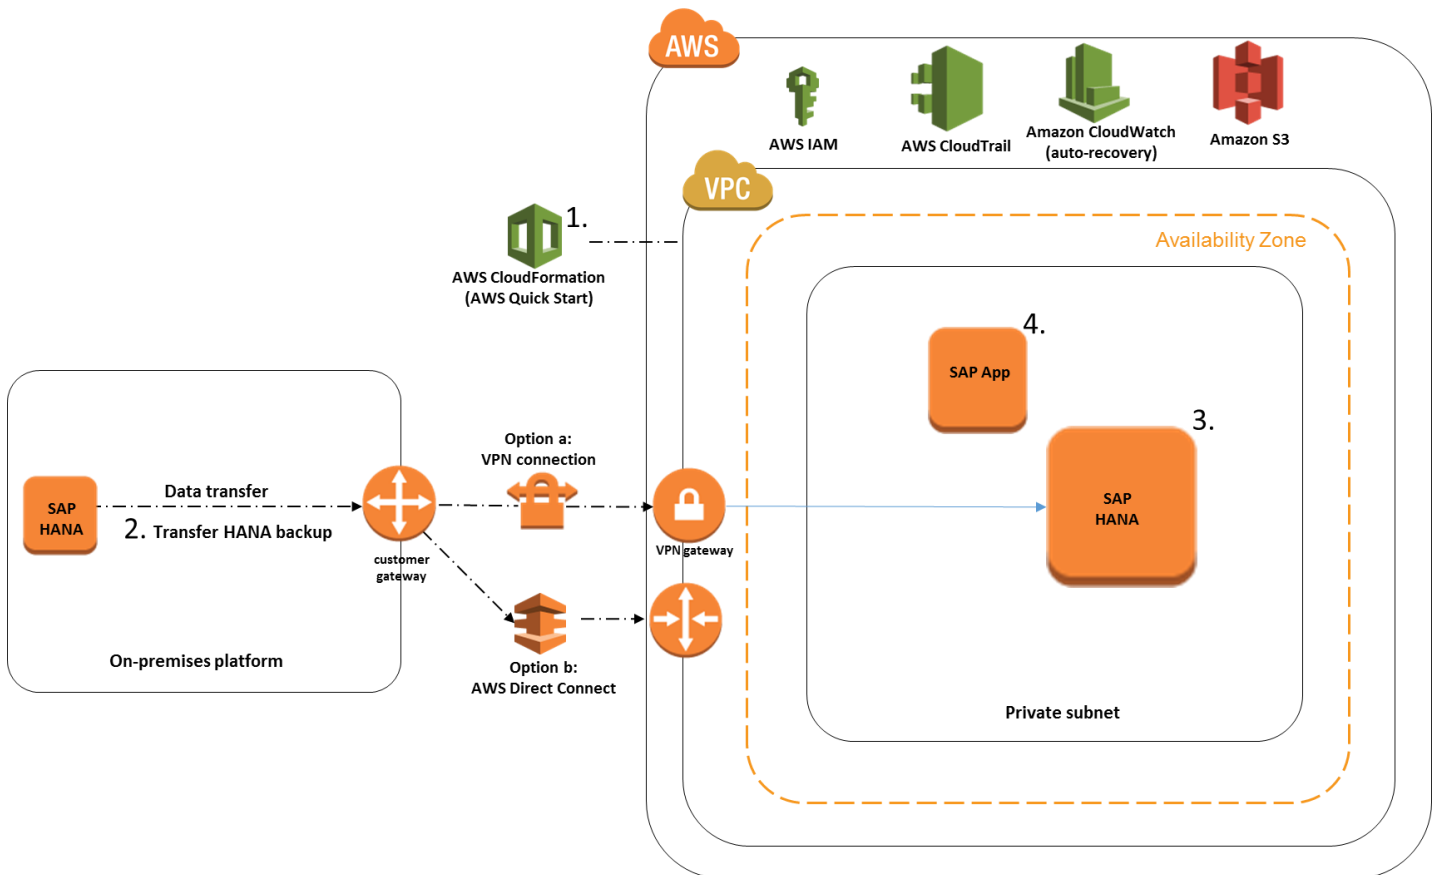
Backup and Restore

Figure 3: Backup and Restore

1. Provision your SAP HANA system and landscape on AWS. (AWS Quick Start templates can help expedite and automate this process for you.)

2. Transfer (**sftp**) a full SAP HANA backup, making sure to transfer any necessary SAP HANA logs to recover to your point in time, from your source system to your AWS X1 system (which has a /hana/backup file system). 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.

3. Recover your SAP HANA database.

4. Install your SAP application servers.

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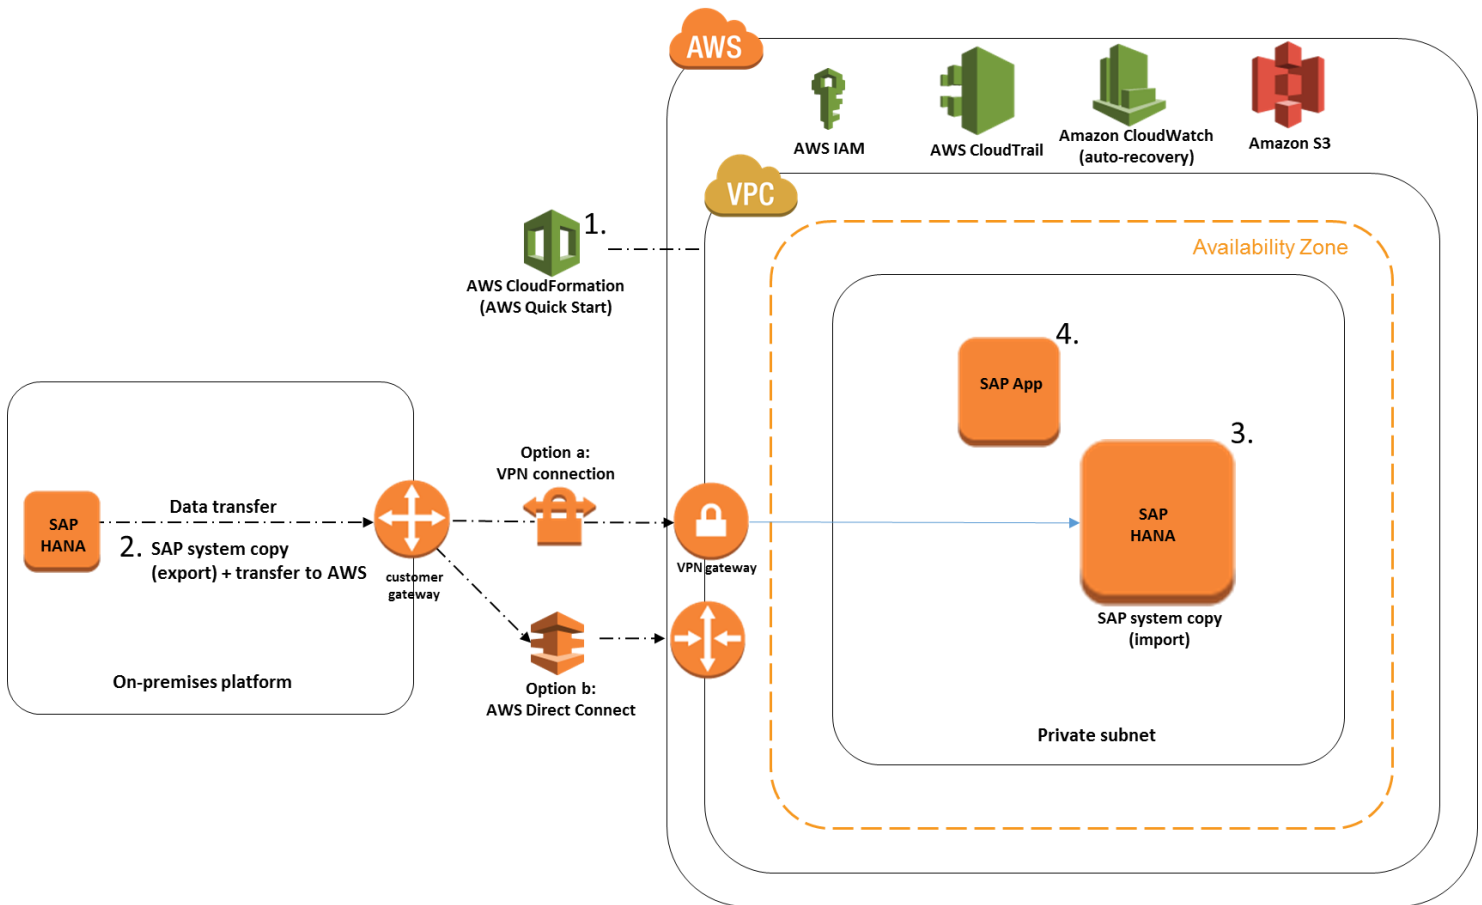
SAP HANA Classical Migration

Figure 4: SAP HANA Classical Migration

1. Provision your SAP HANA system and landscape on AWS. (AWS Quick Start templates 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.) 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.

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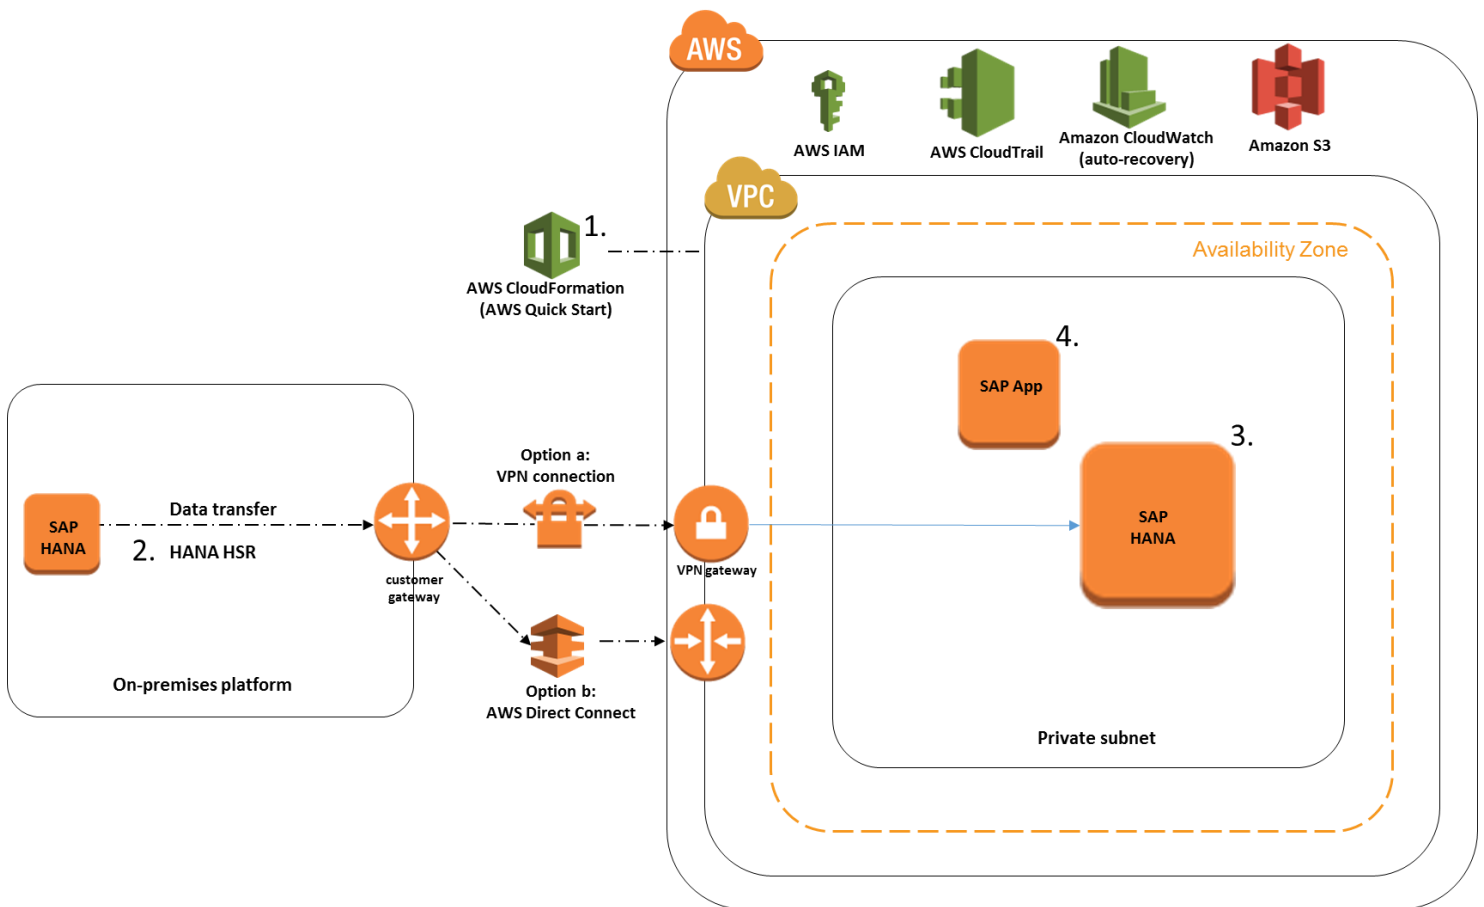
SAP HANA HSR

Figure 5: SAP HANA HSR

1. Provision your SAP HANA system and landscape on AWS. (The AWS Quick Start templates 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 SAP HANA HSR asynchronous 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.

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

Conclusion

AWS X1 instances are optimized for enterprise and database workloads, and are ideal for running in-memory databases such as SAP HANA. X1 offers up to 1,952 GiB of memory and creates new opportunities for deploying and using SAP HANA on AWS. This guide discusses migration options for SAP HANA on AWS. It covers migration scenarios from on-premises or other cloud environments to AWS, and also explains how to upgrade your SAP HANA system to X1 from other instance types on AWS. X1 instances, in combination with AWS and SAP migration tools and methodologies, will boost the performance of your SAP HANA system, and will enable you to architect and run your large production workloads on AWS with flexibility and ease.

Contributors

The following individuals and organizations contributed to this document:

- Somckit Khemmanivanh, AWS solutions architect

- Sabareesan Radhakrishnan, AWS solutions architect

- Bill Timm, AWS solutions architect

Further Reading

For additional help, please consult the following sources:

- [SAP HANA on the AWS Cloud: Quick Start Reference Deployment](#)

- [X1 Overview](#)

- [SAP and Amazon Web Services website](#)

- [SAP on AWS whitepapers](#)

- [AWS documentation](#)