SAP HANA on the AWS Cloud

Quick Start Reference Deployment

Sabareesan Radhakrishnan and Karthik Krishnan
Solutions Architects, AWS Partner Program

July 2014
Last updated: June 2016 (see revisions)

This guide is also available in HTML format at
http://docs.aws.amazon.com/quickstart/latest/sap-hana/.

Supports SAP HANA Platform SPS 9–12

Contents

About This Guide

This Quick Start reference deployment guide provides detailed instructions for deploying SAP HANA on the Amazon Web Services (AWS) cloud by using AWS CloudFormation templates. The Quick Start builds and configures the AWS environment for SAP HANA by provisioning AWS resources such as Amazon Elastic Compute Cloud (Amazon EC2), Amazon Elastic Block Store (Amazon EBS), and Amazon Virtual Private Cloud (Amazon VPC).

This guide is for IT infrastructure architects, system administrators, SAP Basis architects, and SAP Basis administrators who are planning to implement or extend their SAP HANA workloads on the AWS cloud.

Quick Starts are automated reference deployments for key workloads on the AWS cloud. Each Quick Start launches, configures, and runs the AWS compute, network, storage, and other services required to deploy a specific workload on AWS, using AWS best practices for security and availability.

> **Note**   If you are implementing SAP Business One, version for SAP HANA, follow the instructions in the SAP Business One, version for SAP HANA, Quick Start reference deployment guide.

Overview

SAP HANA on AWS

SAP HANA is an in-memory relational database and application platform that provides high-performance analytics and real-time data processing.

The AWS cloud provides a suite of infrastructure services that enable you to deploy SAP HANA in a highly available, fault-tolerant, and affordable way. By deploying this solution on the AWS cloud, you can take advantage of the functionality of SAP HANA along with the flexibility and security of AWS.

This Quick Start helps you rapidly deploy fully functional SAP HANA systems on the AWS cloud, following best practices from AWS and SAP. The Quick Start ensures that Amazon Elastic Compute Cloud (Amazon EC2), Amazon Elastic Block Store (Amazon EBS), and the operating system (SUSE Linux Enterprise Server or Red Hat Enterprise Linux) are optimally configured to achieve the best performance for your SAP HANA system right out of the box.

This Quick Start supports R3 instances and the X1 instance type (x1.32xlarge), which offers 1,952 GiB of instance memory to boost the performance of large-scale, memory-intensive SAP HANA workloads on AWS. To learn more about X1, see [Amazon EC2 X1 Instances](#) and the [Amazon EC2 User Guide](#).

For a full list of new features and changes in this Quick Start release, see [Appendix B](#).

Note   In this Quick Start, we use the terms *SAP HANA platform*, *system*, *database*, *server*, *node*, and *instance*, depending on context, to refer to the SAP HANA deployment on AWS.

This Quick Start currently supports SAP HANA Platform SPS 9*x* through 12*x*.

Quick Links

The links in this section are for your convenience. Before you launch the Quick Start, please review the architecture, configuration, network security, and other considerations discussed in this guide.

- If you have an AWS account, and you're already familiar with AWS services and SAP HANA, you can launch the Quick Start to build the architecture shown in Figure 1 in a new or existing Amazon VPC. You can choose the SUSE Linux Enterprise Server (SLES) or Red Hat Enterprise Linux (RHEL) operating system (see limitations and subscription information). The deployment takes approximately 25 minutes for single-node deployment, or 35-60 minutes for multi-node deployment. If you're new to AWS or SAP HANA, please review the implementation details and follow the step-by-step instructions provided later in this guide.

Launch
(for new VPC)   **Launch**
                           (for existing VPC)

- If you want to take a look under the covers, you can view the AWS CloudFormation templates that automate the deployment.

View template
(for new VPC)   **View template**
                           (for existing VPC)

Cost and Licenses

You are responsible for the cost of the AWS services used while running this Quick Start reference deployment. There is no additional cost for using the Quick Start. The following table provides example prices for r3.8xlarge instances. For X1 instances, see the Amazon EC2 Pricing webpage.

> **Note**   The prices in the table cover two of many possible configuration options and do not include NAT gateway or data transfer costs. You can use the AWS Simple Monthly Calculator links provided in the table and customize the configuration in the calculator to get more accurate pricing. For more information about configuration options, see the Planning the Deployment section. Other factors that influence the cost are billing options (All Upfront, Partial Upfront, No Upfront, or On-Demand) and the region in which the AWS resources are launched.

SAP HANA instance	Scenario	Includes	See
r3.8xlarge	Single-node (scale-up)	Bastion host: t2.small RDP instance: c4.xlarge SAP HANA: Single instance, SLES, r3.8xlarge Storage: General Purpose (SSD)	Calculator
r3.8xlarge	Multi-node (scale-out)	Bastion host: t2.small RDP instance: c4.xlarge SAP HANA: Three instances, SLES, r3.8xlarge Storage: General Purpose (SSD)	Calculator

This deployment uses a Bring Your Own License (BYOL) model for SAP HANA. You must already own a license for SAP HANA, and you must have access to the SAP Software Download Center to download the software.

This Quick Start launches the Amazon Machine Image (AMI) for the operating system you choose—SLES or RHEL—and the license cost for the operating system is included in the Amazon EC2 hourly price.

AWS Services

The core AWS components used by this Quick Start include the following AWS services. (If you are new to AWS, see the Getting Started section of the AWS documentation.)

- Amazon VPC – The Amazon Virtual Private Cloud (Amazon VPC) service lets you provision a private, isolated section of the AWS cloud where you can launch AWS services and other resources in a virtual network that you define. You have complete control over your virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables and network gateways.

- Amazon EC2 – The Amazon Elastic Compute Cloud (Amazon EC2) service enables you to launch virtual machine instances with a variety of operating systems. You can choose from existing AMIs or import your own virtual machine images.

- Automatic recovery – Automatic recovery is a feature of Amazon EC2 that is designed to increase instance availability. You can enable automatic recovery for an instance by creating an Amazon CloudWatch alarm that monitors an Amazon EC2 instance and automatically recovers the instance if it becomes impaired due to an underlying hardware failure or a problem that requires AWS involvement to repair. A recovered instance is identical to the original instance, including the instance ID, private IP

addresses, Elastic IP addresses, and all instance metadata. This Quick Start optionally enables automatic recovery on SAP HANA nodes for you.

- Amazon EBS – Amazon Elastic Block Store (Amazon EBS) provides persistent block-level storage volumes for use with Amazon EC2 instances in the AWS cloud. Each Amazon EBS volume is automatically replicated within its Availability Zone to protect you from component failure, offering high availability and durability. Amazon EBS volumes provide the consistent and low-latency performance needed to run your workloads.

- AWS CloudFormation – AWS CloudFormation gives you an easy way to create and manage a collection of related AWS resources, and provision and update them in an orderly and predictable way. You use a template to describe all the AWS resources (e.g., Amazon EC2 instances) that you want. You don't have to individually create and configure the resources or figure out dependencies—AWS CloudFormation handles all of that.

- Amazon CloudWatch – Amazon CloudWatch monitors your AWS resources and the applications you run on AWS in real time. You can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, set alarms, and automatically react to changes in your AWS resources.

- NAT Gateway – NAT Gateway is an AWS managed service that controls NAT gateway resources. A NAT gateway is a type of network address translation (NAT) device that enables instances in a private subnet to connect to the Internet or to other AWS services, but prevents the Internet from connecting to those instances.

- IAM – AWS Identity and Access Management (IAM) enables you to securely control access to AWS services and resources for your users. With IAM, you can manage users, security credentials such as access keys, and permissions that control which AWS resources users can access, from a central location.

Architecture

This Quick Start uses AWS CloudFormation, the AWS Command Line Interface (AWS CLI) for Linux, and custom scripts to deploy SAP HANA on AWS. AWS CloudFormation provides an easy way to create and manage a collection of related AWS resources, provisioning and updating them in an orderly and predictable fashion. AWS CLI for Linux enables you to configure AWS resources from the command line. This Quick Start includes options for single-node and multi-node SAP HANA configurations.

The Quick Start deploys and configures the following components:

- An Amazon VPC configured with public and private subnets according to AWS best practices, to provide you with your own virtual network on AWS

- An Internet gateway to allow access to the Internet

- Amazon EC2 instance(s) to host the SAP HANA database, with your choice of Linux (SLES or RHEL) and the appropriate Amazon EBS volumes configured to meet or exceed SAP HANA storage key performance indicators (KPIs)

- A bastion host in the public subnet with an Elastic IP address to allow inbound SSH (Secure Shell) access to the Amazon EC2 instance(s) that host the SAP HANA database

- A managed NAT gateway to allow outbound Internet access for resources in the private subnet

- An IAM instance role with fine-grained permissions for access to the AWS services necessary for the deployment process

- Security groups for each instance or function to restrict access to only necessary protocols and ports

- An optional automated installation of SAP HANA software

- An optional Amazon EC2 instance with Windows Server in the public subnet to host SAP HANA Studio. You can install SAP HANA Studio manually to administer your SAP HANA database.

Single-Node Architecture

The single-node deployment option provisions a single Amazon EC2 instance with your choice of Amazon EBS storage and operating system to host the SAP HANA platform. The Quick Start follows security best practices by establishing an Amazon VPC with public and private subnets. For secure access, the SAP HANA server is placed in the private subnet, which is not directly accessible from the Internet. You can also install SAP HANA Studio manually in the optional Windows Server instance that is provisioned in the public subnet. For SSH access to the SAP HANA server, you can use the bastion host or an SSH client on the optional Windows Server instance.
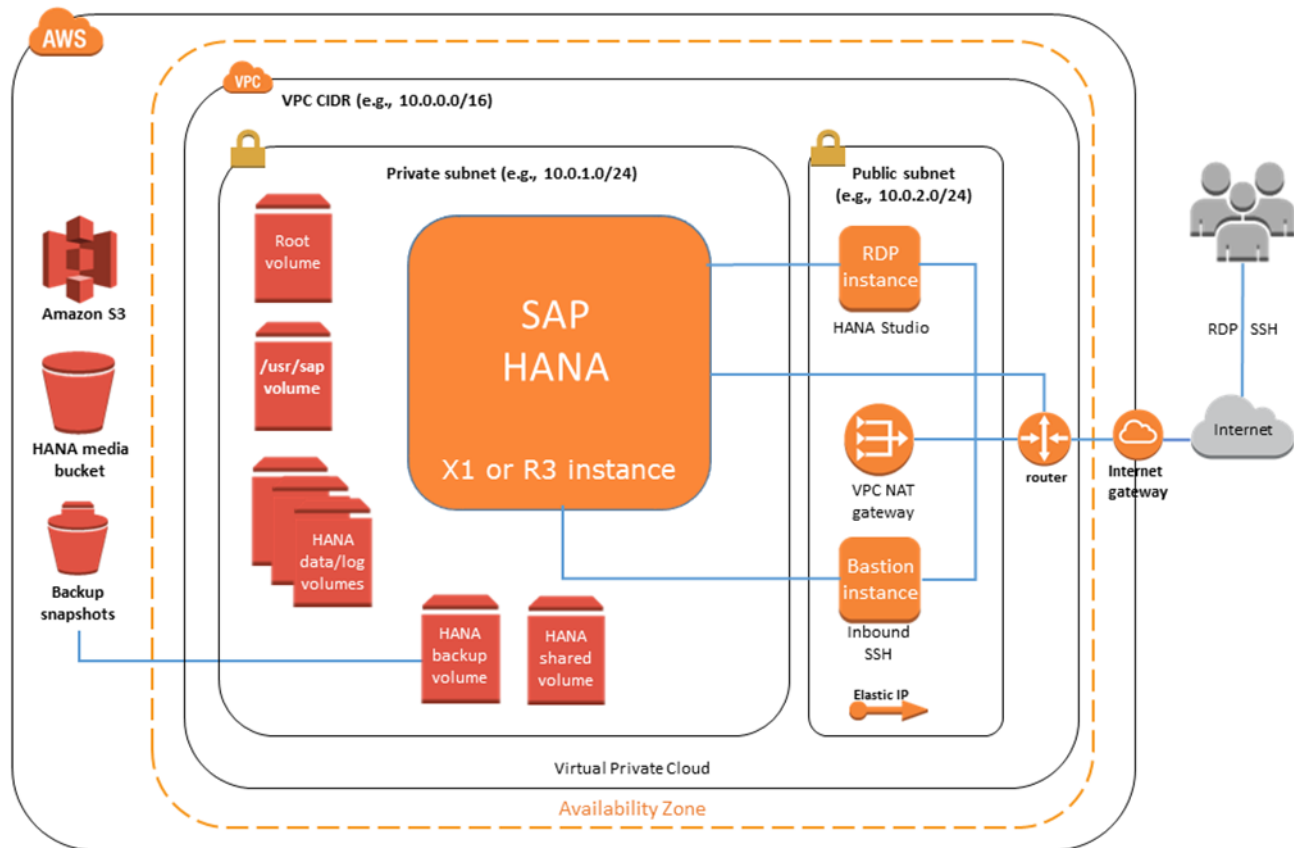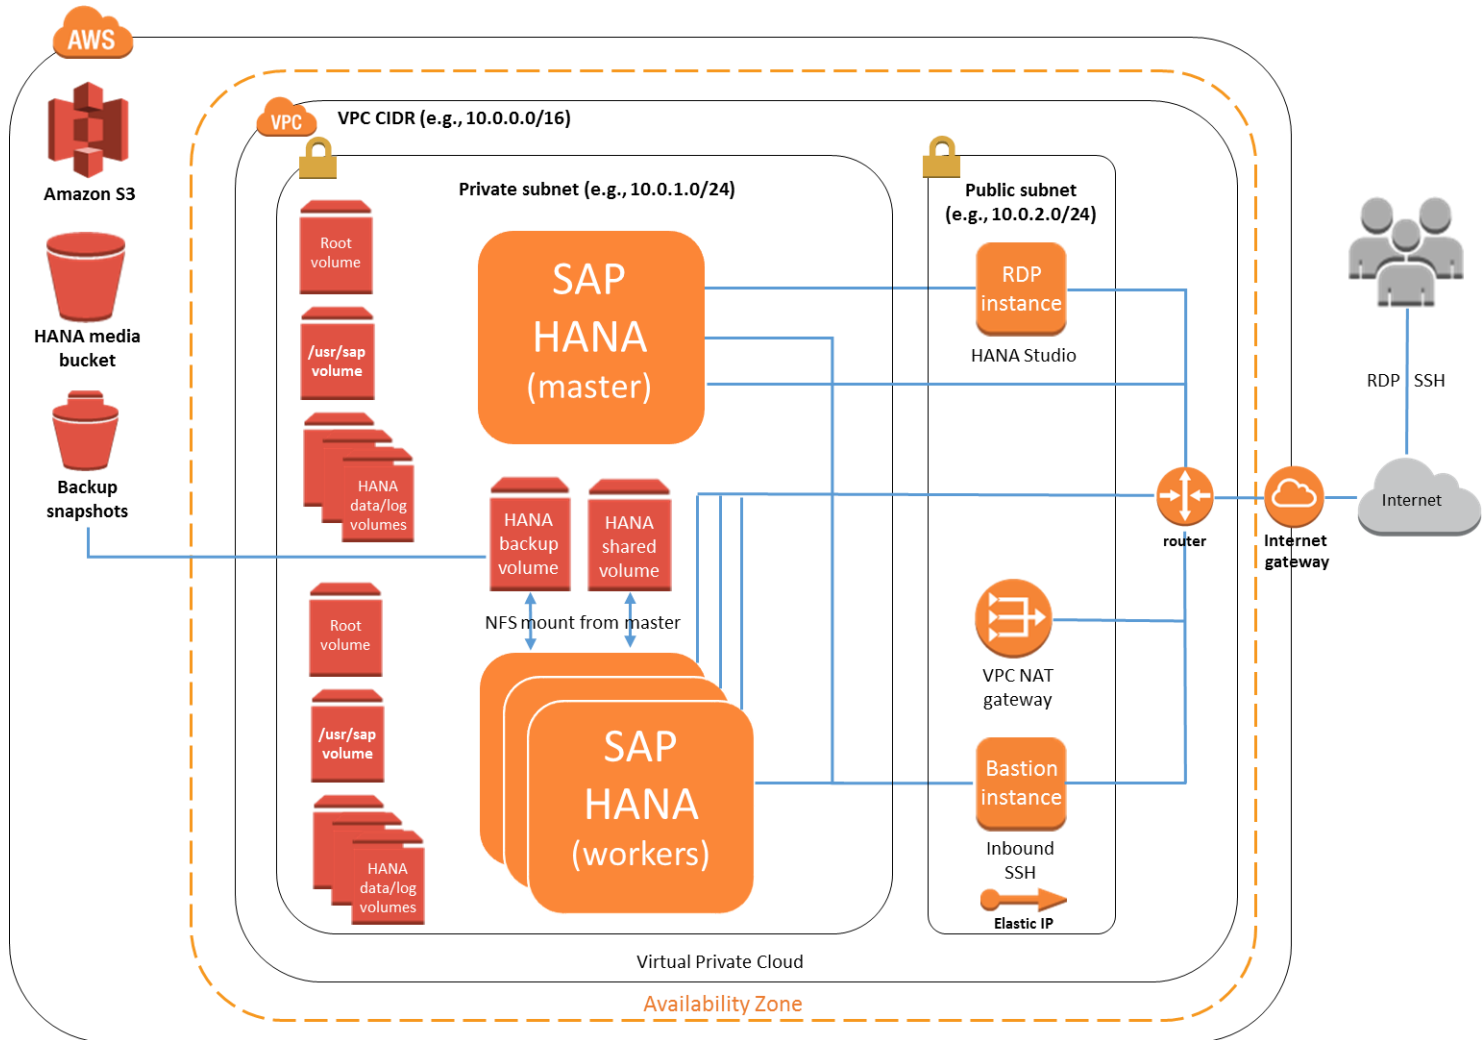
Figure 1: SAP HANA Single-Node Architecture on AWS

Multi-Node Architecture

The multi-node deployment option provisions up to five EC2 instances***** with your choice of Amazon EBS storage and operating system to host the SAP HANA platform on a cluster of servers. All the SAP HANA servers in the cluster are deployed into the same subnet regardless of function, in accordance with security best practices.

*Important   SAP HANA scale-out clusters with the X1 instance type are not currently certified or supported by SAP for production use. SAP HANA scale-out clusters that are larger than five nodes with r3.8xlarge instances are currently in controlled availability. SAP will need to verify your SAP Business Warehouse sizing report results before you implement an SAP HANA scale-out cluster that is larger than five nodes on AWS. Please contact SAP at HWC@sap.com and AWS at sap-on-aws@amazon.com before you implement SAP HANA scale-out clusters of this size. Refer to SAP OSS Note 1964437 – SAP HANA on AWS: Supported AWS EC2 Products (requires SAP support portal login) for details.

As in single-node deployment, access to all SAP HANA servers is provided either through the bastion host or by using the optional Windows Server instance with an SSH client of your choice. The security group or firewall rules are extended to allow for internode communication over the appropriate ports and protocols. For more information, see [Appendix C: Security Groups](#).

Figure 2: SAP HANA Multi-Node Architecture on AWS

Planning the Deployment

Before you deploy SAP HANA on AWS, please review the following sections for guidelines regarding instance types, storage, memory sizing, and operating system choices for the deployment.

AWS Instance Type for SAP HANA

This Quick Start supports four instance types for the SAP HANA host:

- x1.32xlarge

- r3.2xlarge (default)

- r3.4xlarge

- r3.8xlarge

The default instance type is r3.2xlarge, but you can switch to one of the other types during deployment. The r3.8xlarge and x1.32xlarge instance types are officially supported by SAP for production use. For more information about different instance types and their use cases, see the [Amazon EC2 Instances](#) webpage.

Storage Configuration for SAP HANA

SAP HANA stores and processes all or most of its data in memory, and provides protection against data loss by saving the data in persistent storage locations. To achieve optimal performance, the storage solution used for SAP HANA data and log volumes should meet SAP's storage KPI. AWS has worked with SAP to certify both Amazon EBS General Purpose SSD (`gp2`) and Provisioned IOPS SSD (`io1`) storage solutions for SAP HANA workloads with the configuration shown in the tables provided in this section.

`gp2` volumes balance price and performance for a wide variety of workloads, while `io1` volumes provide the highest performance consistently for mission-critical applications. With these two options, you can choose the right storage solution that meets your performance and cost requirements. We recommend using the `io1` configuration with the X1 instance type (x1.32xlarge) for your mission-critical SAP HANA workloads for production use.

Note that only x1.32xlarge and r3.8xlarge instances are certified for production use; you can use the other two R3 instance types for development and testing.

In the following configurations, we intentionally kept the same storage configuration for all R3 instance types to give you the flexibility to scale up from r3.2xlarge or r3.4xlarge to r3.8xlarge without having to reconfigure your storage.

`gp2`-based storage configuration for SAP HANA data and logs:

Instance type	RAM (GiB)	General Purpose SSD (`gp2`) storage (SAP HANA data and log volumes) – striped with LVM	Total maximum throughput (MiB/s)*	Total baseline IOPS	Total burst IOPS
x1.32xlarge	1952	3 x 1024 GiB	480	9000	N/A
r3.8xlarge	244	3 x 300 GiB	480	2700	9000
r3.4xlarge	122	3 x 300 GiB	480	2700	9000
r3.2xlarge	61	3 x 300 GiB	480	2700	9000

`io1`-based storage configuration for SAP HANA data and logs:

Instance type	RAM (GiB)	Provisioned IOPS SSD (`io1`) storage (SAP HANA data and log volumes) – striped with LVM	Total maximum throughput (MiB/s)*	Total provisioned IOPS
x1.32xlarge	1952	3 x 1024 GiB	960	15000
r3.8xlarge	244	2 x 450 GiB	640	9000
r3.4xlarge	122	2 x 450 GiB	640	9000
r3.2xlarge	61	2 x 450 GiB	640	9000

> ***** This value represents the maximum throughput that could be achieved when striping multiple EBS volumes. Actual throughput depends on the instance type. Every instance type has its own Amazon EBS throughput limit. For details, see Amazon EBS-Optimized Instances in the AWS documentation.

In addition to the SAP HANA data and log volumes, all instances deployed by this Quick Start will have the following storage configuration for root, SAP binaries, and SAP HANA shared and backup volumes:

Instance type	RAM (GiB)	Root volume (gp2)	SAP binaries (gp2)	SAP HANA shared (gp2)	SAP HANA backup (st1)
x1.32xlarge	1952	1 x 50 GiB	1 x 50 GiB	1 x 1024 GiB	1 x 4096 GiB
r3.8xlarge	244	1 x 50 GiB	1 x 50 GiB	1 x 300 GiB	1 x 1024 GiB
r3.4xlarge	122	1 x 50 GiB	1 x 50 GiB	1 x 300 GiB	1 x 512 GiB
r3.2xlarge	61	1 x 50 GiB	1 x 50 GiB	1 x 300 GiB	1 x 512 GiB

Storage for SAP HANA backup is configured with Amazon EBS Throughput Optimized HDD (st1) volumes. This volume type provides low-cost magnetic storage designed for large sequential workloads. SAP HANA uses sequential I/O with large blocks to back up the database, so st1 volumes provide a low-cost, high-performance option for this scenario. To learn more about st1 volumes, see Amazon EBS Volume Types in the AWS documentation.

The SAP HANA backup volume size is designed to provide optimal baseline and burst throughput as well as the ability to hold several backup sets. Holding multiple backup sets in the backup volume makes it easier to recover your database if necessary. You may resize your SAP HANA backup volume after initial setup if needed. To learn more about resizing your Amazon EBS volumes, see Expanding the Storage Size of an EBS Volume on Linux in the AWS documentation.

Memory Sizing for Deployment

Before you begin deployment, please consult the SAP documentation listed in this section to determine memory sizing for your needs. This evaluation will help you choose Amazon EC2 instances during deployment. (Note that the links in this section require SAP support portal credentials.)

> **Note**   Starting with SAP HANA Platform SPS 11, you can take full advantage of the 1,952 GiB of memory that the X1 instance offers for your production OLAP workloads, such as SAP Business Warehouse on SAP HANA or data mart on SAP HANA. You can utilize the entire memory that X1 offers with any SAP HANA version for OLTP workloads, such as SAP Business Suite on HANA (SoH), S/4HANA, and any non-production OLAP workloads.

- To obtain sizing information for a system that has not yet been implemented, use the SAP QuickSizer. The SAP QuickSizer provides information on both the SAP HANA in-memory database and the SAP NetWeaver application server where applicable.

- To migrate an existing SAP NetWeaver Business Warehouse system from any database platform to SAP HANA, SAP strongly recommends the new ABAP sizing report for SAP NetWeaver BW, which is described in SAP Note 1736976.

  Further sizing information is also available in the SAP HANA Administration Guide and in the following SAP HANA notes:

SAP note	Description
1736976	Sizing Report for BW on SAP HANA
1637145	SAP BW on SAP HANA: Sizing SAP In-Memory Database
1702409	HANA DB: Optimal number of scale-out nodes for BW on SAP HANA
1855041	Sizing Recommendation for Master Node in BW-on-HANA
1793345	Sizing for SAP Business Suite on SAP HANA
1872170	Business Suite on SAP HANA memory sizing

Operating System for Deployment

This reference deployment supports the following operating systems and versions for your SAP HANA instance:

- SLES: SLES 11 SP4, SLES 12, SLES 12 SP1

- RHEL: RHEL 6.6

> **Important**   X1 instances require a minimum Linux kernel version of 3.10. Only SLES 12 and later versions meet this requirement, so SLES 11 and RHEL 6 are not supported with the X1 instance type. Deployment will fail on X1 if you choose a Linux operating system based on an earlier version of the kernel.

Deployment Scenarios

This guide serves as a reference for customers who are interested in deploying SAP HANA on AWS in a self-service fashion. The AWS CloudFormation template provided with this Quick Start bootstraps the AWS infrastructure and automates the deployment of SAP HANA on the AWS cloud from scratch. The Quick Start provides two deployment scenarios:

- **Deployment of SAP HANA into a new Amazon VPC** (end-to-end deployment) builds the Amazon VPC, subnets, NAT gateway, security groups, bastion host, and optional Windows Server and SAP HANA server(s) with Amazon EBS.

- **Deployment of SAP HANA into an existing Amazon VPC** provisions SAP HANA servers in your existing infrastructure.

This Quick Start supports the deployment of SAP HANA on two operating systems: You can choose either SLES or RHEL. (See the previous section for supported versions.) The Quick Start also lets you configure additional deployment settings such as CIDR blocks, number of SAP HANA nodes, and Availability Zone, as discussed later in this guide.

Deployment Steps

The procedure for an end-to-end deployment of SAP HANA on AWS consists of the following steps. For detailed instructions, follow the links for each step.

- Step 1. Prepare an AWS account

  This involves signing up for an AWS account, choosing a region, creating a key pair, and requesting increases for account limits, if necessary.

- Step 2. Subscribe to RHEL for SAP HANA (skip this step if you're planning to use the SLES operating system)

  The RHEL AMI is available from the AWS MarketPlace via subscription.

- Step 3. Download the SAP HANA software (skip this step if you don't want to install SAP HANA software with the deployment)

  This step involves downloading the SAP HANA software from SAP and placing the files in an S3 bucket.

- Step 4. Launch the stack

  In this step, you'll launch the AWS CloudFormation template into your AWS account, specify parameter values, and create the stack. The Quick Start provides separate templates for end-to-end deployment and deployment into an existing Amazon VPC.

- Step 5. Access SAP HANA nodes to verify your deployment

  You can access nodes by using SAP HANA Studio or through SSH and the bastion host.

- Step 6. Complete any post-deployment tasks

  Before you start using SAP HANA on AWS, we recommend that you check for the latest updates and make sure that your system is backed up and configured correctly.

Step 1. Prepare Your AWS Account

If you already have an AWS account, skip to step 2.

1. If you don't already have an AWS account, create one at http://aws.amazon.com by following the on-screen instructions. Part of the sign-up process involves receiving a phone call and entering a PIN using the phone keypad.

2. Use the region selector in the navigation bar to choose the Amazon EC2 Region where you want to deploy SAP HANA on AWS.

   Amazon EC2 locations are composed of *Regions* and *Availability Zones*. Regions are dispersed and located in separate geographic areas.
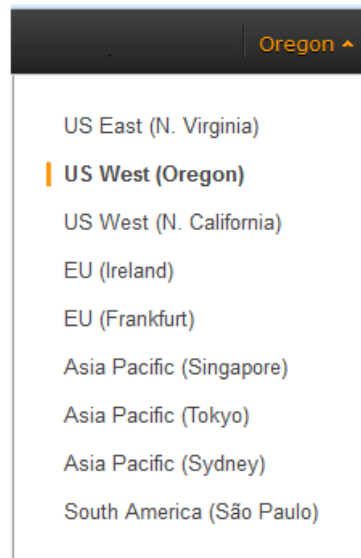
Figure 3: Choosing an Amazon EC2 Region

Consider choosing a region closest to your data center or corporate network to reduce network latency between systems running on AWS and the systems and users on your corporate network.

> **Important**   This Quick Start supports the R3 or X1 instance type for the SAP HANA part of the deployment. At the time of this writing, R3 instances are available in all AWS regions except China (Beijing) and South America (São Paulo). For more information about X1, see the Amazon EC2 X1 Instances webpage. We recommend that you check the availability of AWS services before you choose a region. Otherwise, deployment will fail.

3. Create a key pair in your preferred region. To do this, in the navigation pane of the Amazon EC2 console, choose **Key Pairs**, **Create Key Pair**, type a name, and then choose **Create**.
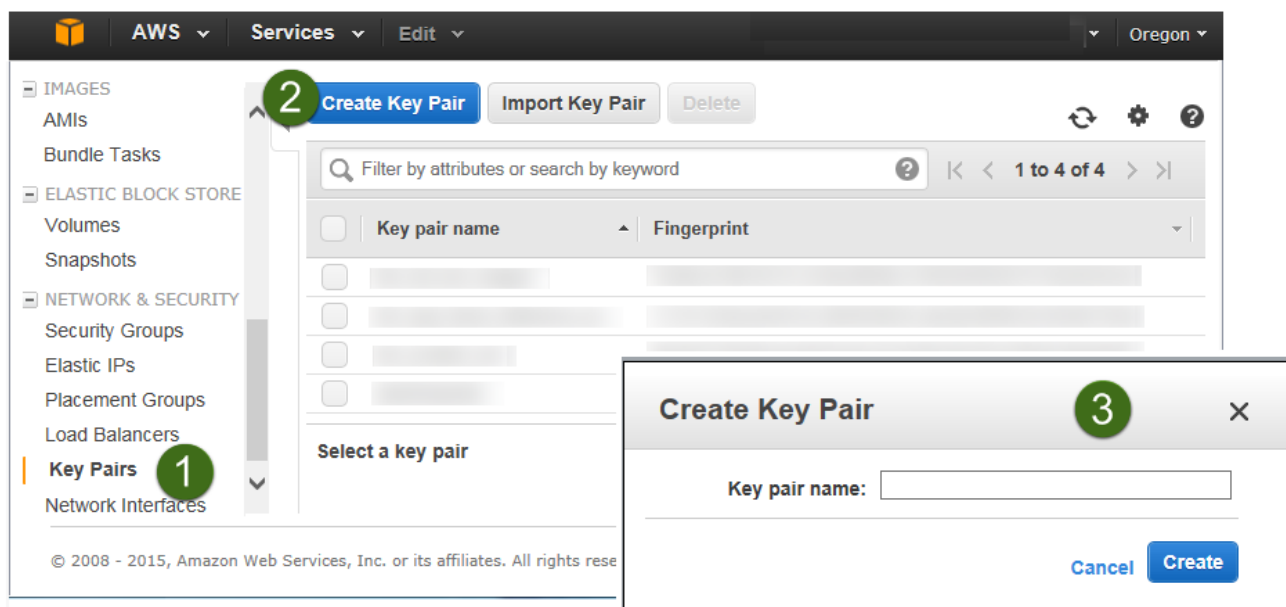
Figure 4: Creating a Key Pair

Amazon EC2 uses public-key cryptography to encrypt and decrypt login information. To be able to log in to your instances, you must create a key pair. With Windows instances, we use the key pair to obtain the administrator password via the Amazon EC2 console and then log in using Remote Desktop Protocol (RDP) as explained in the step-by-step instructions in the *Amazon Elastic Compute Cloud User Guide*. On Linux, we use the key pair to authenticate SSH login.

4. (Production deployments only) If necessary, request a service limit increase for the instance type you're using. If you already have an existing deployment that uses this instance type, and you think you might exceed the default limit with this reference deployment, you will need to request an increase. To do this, in the AWS Support Center, choose **Create Case**, **Service Limit Increase**, **EC2 instances**, and then

complete the fields in the limit increase form. It might take a few days for the new service limit to become effective. For more information, see Amazon EC2 Service Limits in the AWS documentation.

Figure 5: Requesting a Service Limit Increase

Step 2. Subscribe to RHEL for SAP HANA

If you're planning to use the SLES operating system, no subscription is required, and you can skip ahead to step 3.

1. Log in to the AWS Marketplace at http://aws.amazon.com/marketplace.

2. Subscribe to **Red Hat Enterprise Linux for SAP HANA**.

   This involves accepting the terms of the license agreement and receiving confirmation email. For detailed instructions, see Appendix A.

Step 3. Download the SAP HANA Software

Skip this step if you already have the unextracted SAP HANA software in an Amazon S3 bucket **or** if you don't want to install SAP HANA during this deployment.

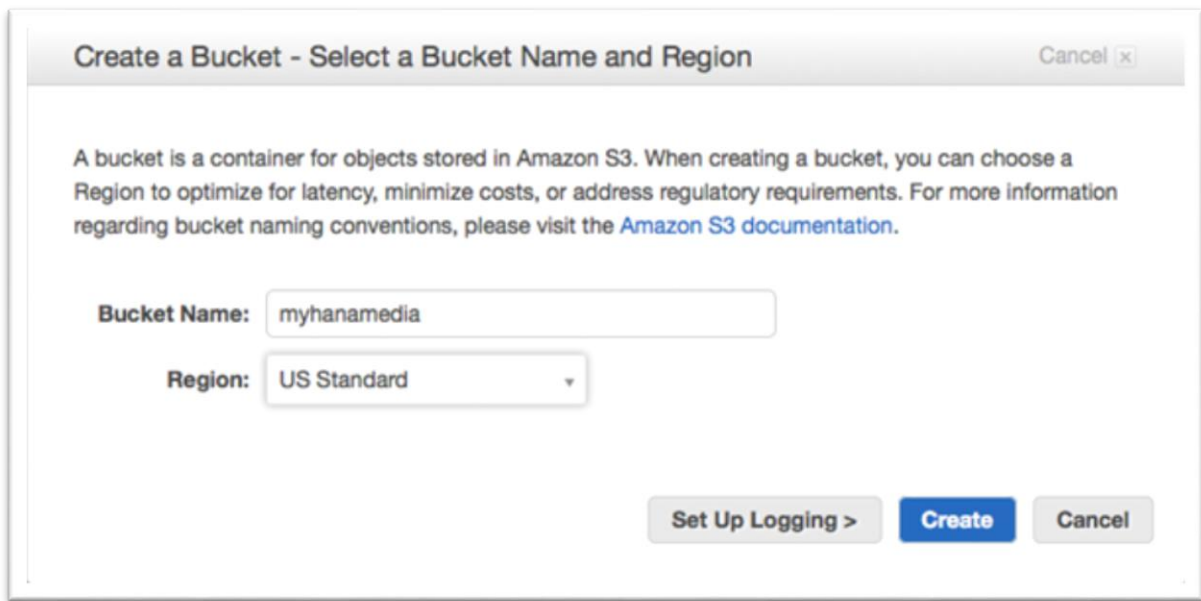1. Go to http://support.sap.com/swdc and log in.

   Tip If your SAP Support Portal account doesn't allow access to the software and you believe that you should already be entitled to the software, contact the SAP Global Support Customer Interaction Center through the web form available at http://support.sap.com/contactus.

2. In the left navigation pane, choose **Installation and Upgrades**, and then choose **A – Z index**.

3. In the **Installations and Upgrades** window, choose **H**, and then choose **SAP HANA Platform Edition** from the list.

4. Choose **SAP HANA Platform Edit. 1.0**, and then choose **Installation**.

5. In the **Downloads** windows, find the revision you wish to download and download each file directly to your local drive. Note that this Quick Start currently supports SAP HANA Platform SPS 9*x* through 12*x*.

   Next, you'll create an S3 bucket for storing the SAP HANA installation files.

   Important Do not extract the downloaded HANA software. Just stage the files in an Amazon S3 bucket as is. The Quick Start will automatically extract the media and install the software for you.

6. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3.

7. Choose **Create Bucket**.

8. In the **Create a Bucket** dialog box, provide a name for your new bucket, choose the region where you want to create your bucket (this should be a region that is close to your location), and then choose **Create**. For detailed information about bucket names and region selection, see the Amazon S3 documentation.

Figure 6: Creating an S3 Bucket

9. In the upper right, choose **Properties**, **Permissions**, and set permissions to ensure that only you and authorized personnel from your organization have access to this bucket. You can also set up an IAM or bucket policy to provide fine-grained access. For details, see Managing Access Permissions to Your Amazon S3 Resources in the Amazon S3 documentation.
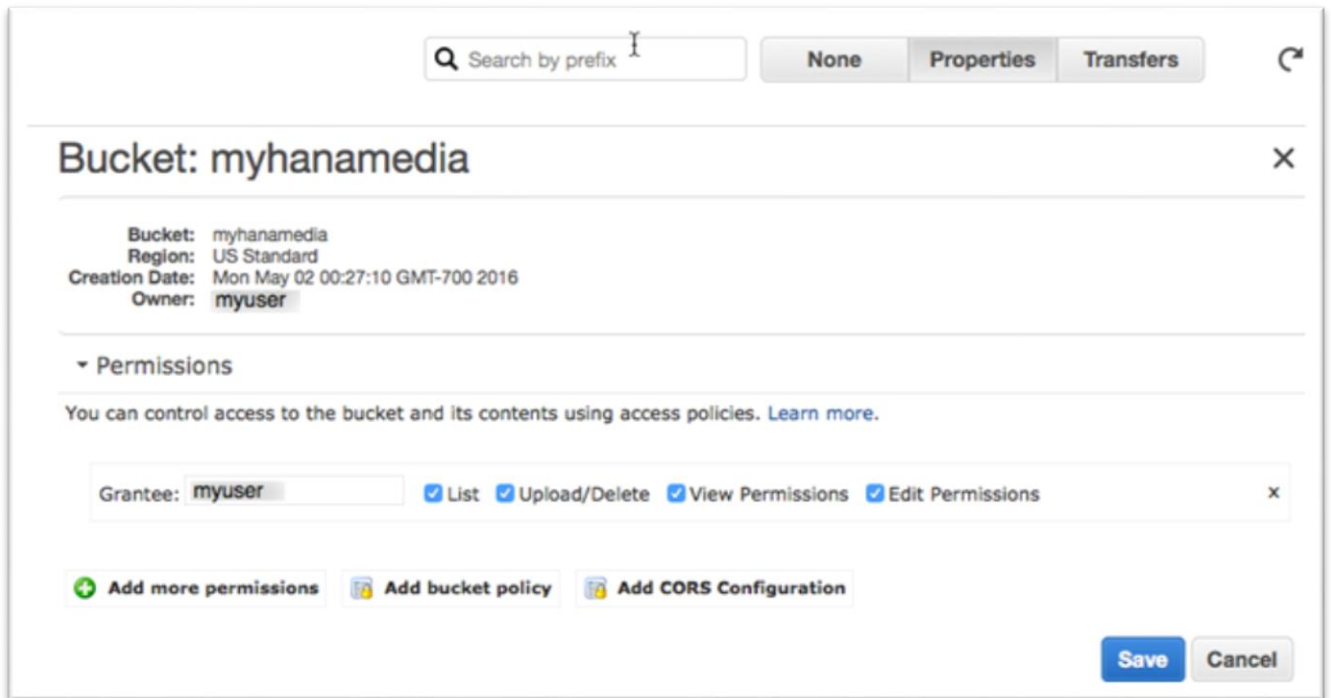
Figure 7: Setting Bucket Permissions

10. Choose the bucket that you created, and add folders to organize your SAP HANA downloads. We recommend that you create a folder for each version of SAP HANA for ease of maintenance and identification.
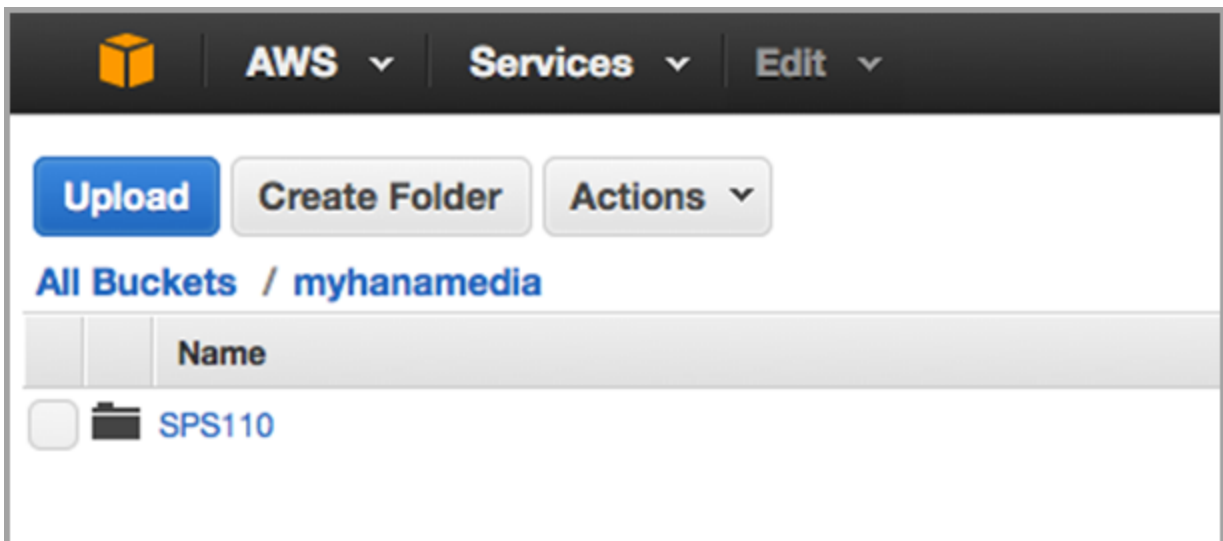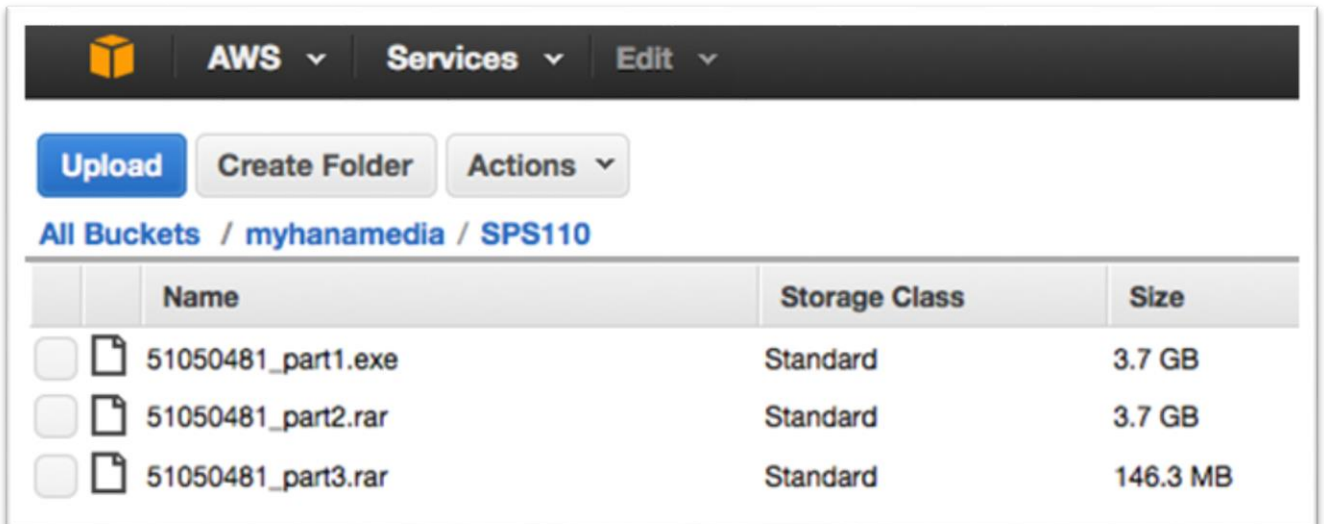

Figure 8: Adding Subfolders for SAP HANA Versions

11. Choose **Upload** to place the unextracted SAP HANA software in the appropriate folder, as illustrated in Figure 9.

Figure 9: Staging the SAP HANA Files

In the example shown in Figure 9, the path for this specific version of SAP HANA software will be s3://myhanamedia/SPS110/. Use this path in the SAP HANA media parameter field in the next step.
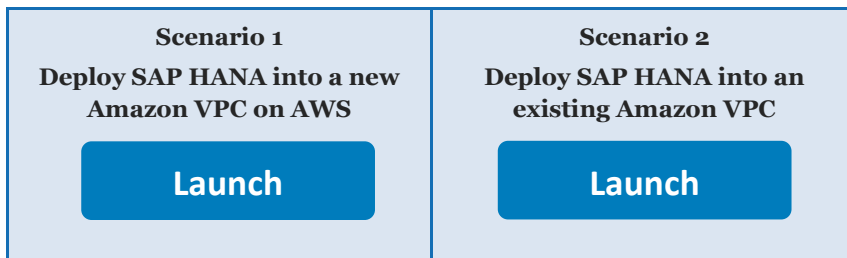
> **Tip**   Place only the main SAP HANA installation files in the S3 bucket (as shown in Figure 9). Do not place multiple SAP HANA versions in the same folder.

Step 4. Launch the Quick Start

> **Important**   With this reference deployment, you can choose between SUSE Linux Enterprise Server (SLES) and Red Hat Enterprise Linux (RHEL) as the operating system for SAP HANA. If you decide to use RHEL, make sure to complete step 2 first to subscribe to the AMI in the AWS Marketplace. If you decide to choose SLES as the operating system for SAP HANA, no subscriptions are required.

In this section, we've provided general instructions for deploying the templates in the AWS CloudFormation console.

1. Choose one of the following options to launch the AWS CloudFormation template into your AWS account. For help choosing an option, see the discussion of [deployment scenarios](#) earlier in this guide.

Scenario 1	Scenario 2
Deploy SAP HANA into a new Amazon VPC on AWS	**Deploy SAP HANA into an existing Amazon VPC**
Launch	**Launch**

Each stack takes approximately 35-60 minutes to complete, depending on the number of SAP HANA nodes you choose to deploy.

> **Note**   You are responsible for the cost of the AWS services used while running this Quick Start reference deployment. There is no additional cost for using this Quick Start. See the [Cost and Licenses](#) section for cost estimates. Prices are subject to change. For full details, see the pricing pages for each AWS service you will be using in this Quick Start or the AWS Simple Monthly Calculator.

2. Check the region that's displayed in the upper-right corner of the navigation bar, and change it if necessary. This is where the network infrastructure for SAP HANA will be built. The template is launched in the US East (N. Virginia) region by default.

3. On the **Select Template** page, keep the default URL for the AWS CloudFormation template, and then choose **Next**.

4. On the **Specify Details** page, change the stack name if needed. Review the parameters for the template. Provide values for the parameters that require input. For all other parameters, review the default settings and customize them as necessary. When you finish reviewing and customizing the parameters, choose **Next**.

   Parameters are grouped in five major categories. In the following tables, parameters are listed and described separately for the two deployment scenarios:

   – [Parameters for deployment into a new Amazon VPC](#)

   – [Parameters for deployment into an existing Amazon VPC](#)

Scenario 1: Parameters for deployment into a new Amazon VPC

View template

Network infrastructure configuration:

Parameter	Default	Description
CIDR block of Amazon VPC	10.0.0.0/16	CIDR block of the Amazon VPC to create for the SAP HANA deployment.
CIDR block of private subnet	10.0.1.0/24	CIDR block of the private subnet in the new Amazon VPC where SAP HANA will be deployed.
CIDR block of public subnet	10.0.2.0/24	CIDR block of the public subnet in the new Amazon VPC where the managed NAT gateway, bastion host, and RDP instance will be deployed.
CIDR block for SSH/RDP access	0.0.0.0/0	CIDR block from where you are likely to access your bastion and RDP instance.
Instance type for bastion host	t2.small	Amazon EC2 instance type for the bastion host.
Availability Zone for subnet creation	*Depends on region selected (e.g., us-east-1a)*	Availability Zone where you want to create your subnets.

Server and storage configuration:

Parameter	Default	Description
Operating system and version for SAP HANA	SLES12SP1	Operating system and version to be used for SAP HANA master/worker nodes (see Operating System for Deployment section).
Instance type for SAP HANA	r3.2xlarge	Amazon EC2 instance type for SAP HANA nodes (see AWS Instance Type for SAP HANA section).
Number of SAP HANA hosts	1	Total number of nodes you want to deploy in the SAP HANA cluster.
Automatic recovery	Yes	Set to No to disable the automatic recovery feature on your SAP HANA nodes.
Key pair	*Requires input*	An existing public/private key pair, which enables you to connect securely to your instance after it launches. When you created an AWS account, this is the key pair you created in your preferred region. This key pair can be used with all Amazon EC2 instances launched by the Quick Start.
Storage volume type	gp2	Amazon EBS storage type to be used for SAP HANA data and log volumes (see Storage Configuration for SAP HANA section).
Encryption	No	Set to Yes to enable encryption for all volumes (except root) created for SAP HANA nodes.

SAP HANA database configuration:

Parameter	Default	Description
Domain name	local	Name to use for fully qualified domain names.
SAP HANA master host name	imdbmaster	Host name to use for the SAP HANA master node (DNS short name).
SAP HANA worker host name	imdbworker	Host name to use for the SAP HANA worker nodes (DNS short name). **Note**: Depending on the number of nodes, each worker node will have a numeric value appended to it automatically (for example, imdbworker01, imdbworker02, etc.).
SAP HANA system ID	HDB	System ID to be used for HANA installation and setup.
SAP HANA instance number	00	Instance number to be used for SAP HANA installation and setup and to open up specific ports for security groups.
SAP HANA password	*Requires input*	Password to be used for SAP HANA installation. This password must follow SAP HANA password restrictions.
Amazon S3 URL for SAP HANA software	*Requires input*	Full path to Amazon S3 location where you've placed the SAP HANA software in step 3.
Install SAP HANA software?	Yes	Set to No if you don't want the Quick Start to install SAP HANA during deployment. This (installation) parameter works in conjunction with the previous (URL) parameter. If you leave the URL parameter blank, the Quick Start won't install the software even if the installation parameter is set to Yes. If you provide a URL but set the installation parameter to No, the Quick Start will stage the SAP HANA software but won't install it.

Optional configuration:

Parameter	Default	Description
Placement group name	—	Name of existing placement group where you want to deploy your SAP HANA cluster. This is required for SAP HANA scale-out deployments only.
Windows RDP instance	No	Set to Yes to enable the deployment of an Amazon EC2 instance with Windows Server to host SAP HANA Studio.
Instance type for RDP host	c4.large	Amazon EC2 instance type for Windows RDP instance.

Advanced configuration:

Parameter	Default	Description
Private bucket	—	Name of the main build bucket where installation scripts and templates are located. Don't change this unless directed by AWS Support.

Scenario 2: Parameters for deployment into an existing Amazon VPC

[View template](#)

Network infrastructure configuration:

Parameter	Default	Description
VPC ID	*Requires input*	ID of the existing Amazon VPC where you want to deploy your SAP HANA resources.
Private subnet	*Requires input*	ID of the private subnet in your existing Amazon VPC where you want to deploy SAP HANA instances.
Public subnet	*Requires input*	ID of the public subnet in your existing Amazon VPC where you want to deploy the optional RDP instance.
CIDR block of private subnet	10.0.1.0/24*	CIDR block of the private subnet in the existing Amazon VPC where SAP HANA will be deployed.
CIDR block of public subnet	10.0.2.0/24*	CIDR block of the public subnet in the existing Amazon VPC where the bastion host and NAT gateway exist.

> * You can retrieve these values from the previous two parameters (private and public subnet IDs), as shown in Figure 10.

Figure 10: Finding the Values for the CIDR Block Parameters

Server and storage configuration:

Parameter	Default	Description
Operating system for SAP HANA	SLES12SP1	Operating system and version to be used for SAP HANA master/worker nodes (see Operating System for Deployment section).
Instance type for SAP HANA	r3.2xlarge	Amazon EC2 instance type for SAP HANA nodes (see AWS Instance Type for SAP HANA section).
Number of SAP HANA hosts	1	Total number of nodes you want to deploy in the SAP HANA cluster.
Automatic recovery	Yes	Set to No to disable the automatic recovery feature on your SAP HANA nodes.
Key pair	*Requires input*	An existing public/private key pair, which enables you to connect securely to your instance after it launches. When you created an AWS account, this is the key pair you created in your preferred region. This key pair can be used with all Amazon EC2 instances launched by the Quick Start.
Storage volume type	gp2	Amazon EBS storage type to be used for SAP HANA data and log volumes (see Storage Configuration for SAP HANA section).
Encryption	No	Set to Yes to enable encryption for all volumes (except root) created for SAP HANA nodes.

SAP HANA database configuration:

Parameter	Default	Description
Domain name	local	Name to use for fully qualified domain names.
SAP HANA master host name	imdbmaster	Host name to use for the SAP HANA master node (DNS short name).
SAP HANA worker host name	imdbworker	Host name to use for the SAP HANA worker nodes (DNS short name). **Note**: Depending on the number of nodes, each worker node will have a numeric value appended to it automatically (for example, imdbworker01, imdbworker02, etc.).
SAP HANA system ID	HDB	System ID to be used for HANA installation and setup.
SAP HANA instance number	00	Instance number to be used for SAP HANA installation and setup and to open up specific ports for security groups.
SAP HANA password	*Requires input*	Password to be used for SAP HANA installation.
Amazon S3 URL for SAP HANA software	*Requires input*	Full path to Amazon S3 location where you've placed the SAP HANA software in step 3.
Install SAP HANA software?	Yes	Set to No if you don't want the Quick Start to install SAP HANA during deployment. This (installation) parameter works in conjunction with the previous (URL) parameter. If you leave the URL parameter blank, the Quick Start won't

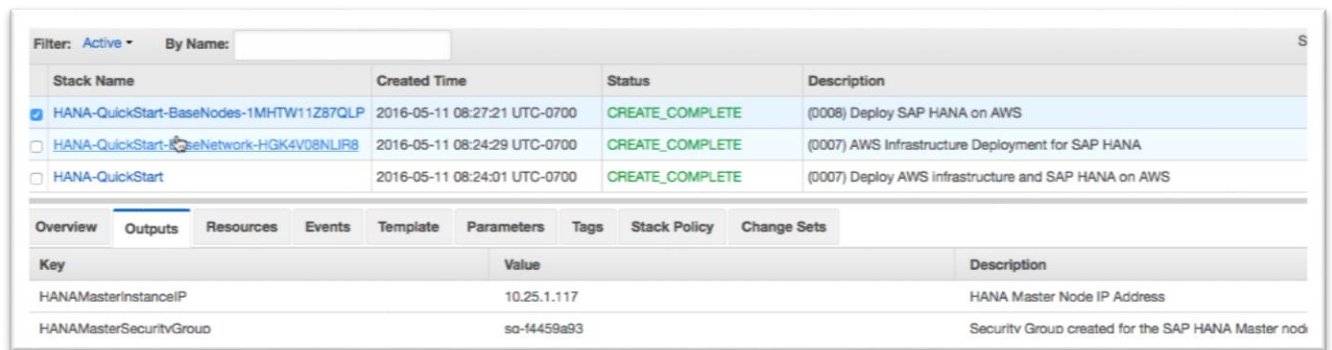| Parameter	Default	Description
		install the software even if the installation parameter is set to Yes. If you provide a URL but set the installation parameter to No, the Quick Start will stage the SAP HANA software but won't install it.

Optional configuration:

Parameter	Default	Description
Placement group name	—	Name of existing placement group where you want to deploy your SAP HANA cluster. This is required for SAP HANA scale-out deployments only.
Windows RDP instance	No	Set to Yes to enable the deployment of an Amazon EC2 instance with Windows Server to host SAP HANA Studio.
Instance type for RDP host	c4.large	Amazon EC2 instance type for Windows RDP instance.
CIDR block for RDP access	0.0.0.0/0	CIDR block from where you want to access the RDP instance.

Advanced configuration:

Parameter	Default	Description
Private bucket	—	Name of the main build bucket where installation scripts and templates are located. Don't change this unless directed by AWS Support.

5. On the **Options** page, you can specify tags (key-value pairs) for resources in your stack and set additional options. When you're done, choose **Next**.

6. On the **Review** page, review and confirm the template settings. Under **Capabilities**, select the checkbox to acknowledge that the template will create IAM resources.

7. Choose **Create** to deploy the stack.

8. Monitor the status of the stack. When the status is **CREATE_COMPLETE**, you can continue to the next step to verify your deployment.

Figure 11: Successful Completion of Template Deployment

The Quick Start creates an Amazon DynamoDB table to track the status of deployment. You can also view the status of your deployment by looking at the contents of the Amazon DynamoDB table from the AWS Management Console.

Step 5. Access SAP HANA Nodes

The default network security setup for this solution follows AWS security best practices. The provisioning logic creates the solution architecture described in the [Architecture](#) section, with the SAP HANA instances in a private subnet to restrict direct exposure to the Internet. As such, the SAP HANA instances can be accessed only through instances that are placed in the public subnet or DMZ layer.

You can access SAP HANA nodes through this DMZ layer in two ways:

- **Access with SAP HANA Studio**: To access your SAP HANA database with SAP HANA Studio, you will need to use a remote desktop client to connect to the Windows Server instance. Once connected, you can manually install SAP HANA Studio and start accessing your SAP HANA database.

- **OS-level access:** Connect to the bastion host and then to the SAP HANA instance(s) by using an SSH client of your choice.

These two methods are discussed in the following sections.

> **Tip**   To connect directly to the SAP HANA systems from a corporate network, you can provision an encrypted IPSec hardware VPN connection between your corporate data center and your Amazon VPC. For details, see [Amazon VPC](#) on the AWS website.

You can also set up AWS Direct Connect between your data center and AWS to gain direct access to your AWS resources. See [Amazon Direct Connect](#) on the AWS website for details.

Using SAP HANA Studio

To install SAP HANA Studio, establish a connection to the Windows Server instance.

1. Open the Amazon EC2 console at https://console.aws.amazon.com/ec2/.

2. From the console dashboard, choose **Running Instances** to find the RDP instance.
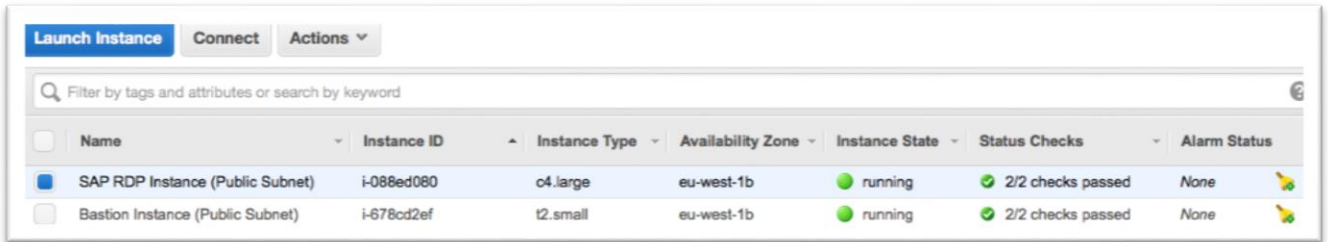
Figure 12: Amazon EC2 Running Instances with RDP Instance Selected

3. Select your RDP instance and choose **Connect**.

4. Get the Windows administrator password from the Amazon EC2 console:

   a. In the **Connect to Your Instance** dialog box, choose **Get Password**.

   b. Paste the contents of your private key in the space provided, or choose **Browse** and navigate to your private key file, select the file, and choose **Open** to copy the entire contents of the file into the contents box.

   The password will be decrypted and displayed.

5. In the **Connect to Your Instance** dialog box, choose **Download Remote Desktop File**, or connect by using an RDP client of your choice.

6. Install SAP HANA Studio. You can do this in two ways:

   – Download the SAP HANA Studio installation files from SAP Service Marketplace.
     —or—
   – Download and extract the SAP HANA software from your S3 bucket to install SAP HANA Studio.

7. When the installation is complete, start SAP HANA Studio and add a system with the following parameters:

   – IP address: address of master node

   – Instance number: 00

   – User: SYSTEM

   – Password: the master password you entered during step 4

> **Note**   At this point, we recommend that you make a backup of your newly installed SAP HANA instance by using SAP HANA Studio. You can also use the Amazon EC2 console to make a complete system image (Amazon Machine Image, or AMI) that can be used for recovery or for additional system builds. Keep in mind that this image is only a point-in-time snapshot.

Using OS-Level Access

You can also connect to the bastion host to establish a remote SSH connection to any of the SAP HANA master or worker nodes.

1. On the Amazon EC2 console, choose **Running Instances**.

2. Select your bastion host, and note the public Elastic IP address displayed below your running instances.
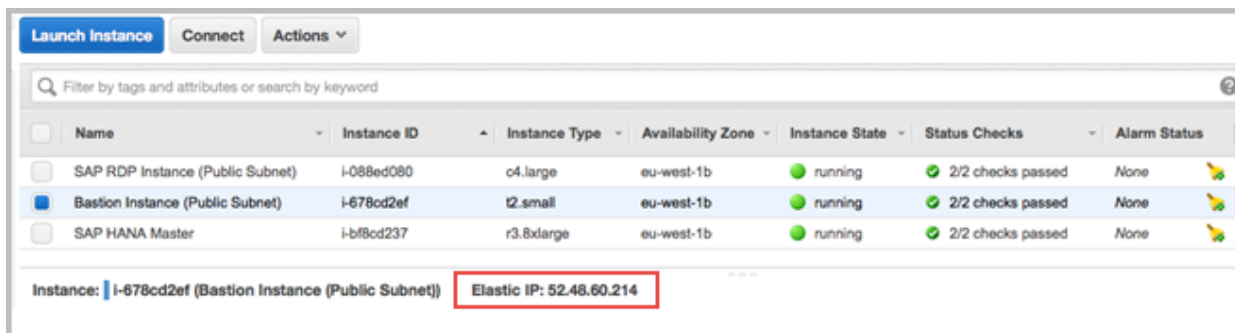
Figure 13: Elastic IP Address for Bastion Host

3. Using an SSH client of your choice (for example, PuTTY or iTerm), connect to the bastion host and use the key pair you specified during the deployment process.

> **Note**   If your connection times out, you might need to adjust the security group rules for the bastion host to allow access from your computer's IP address or proxy server. For more information, see Security Group Rules in the *Amazon EC2 User Guide*.

iTerm Example

1. Add the private key to the authentication agent (`ssh-add`).

2. Connect to the bastion host by using SSH, with the `-A` option to forward the key, specifying the username `ec2-user`.

3. Connect to the SAP HANA server by IP address using SSH.

Figure 14: iTerm Example for SSH Connection

PuTTY Example

1. Download PuTTY (putty.exe), PuTTY Key Generator (puttygen.exe), and Pageant (pageant.exe).

2. Load your private key into PuTTY Key Generator and save it as a .ppk file that PuTTY can use.

3. Run Pageant.exe, and add your new. ppk key. The Pageant process must be running in order for agent forwarding to work.

4. Configure PuTTY with the private key and select **Allow agent forwarding**.
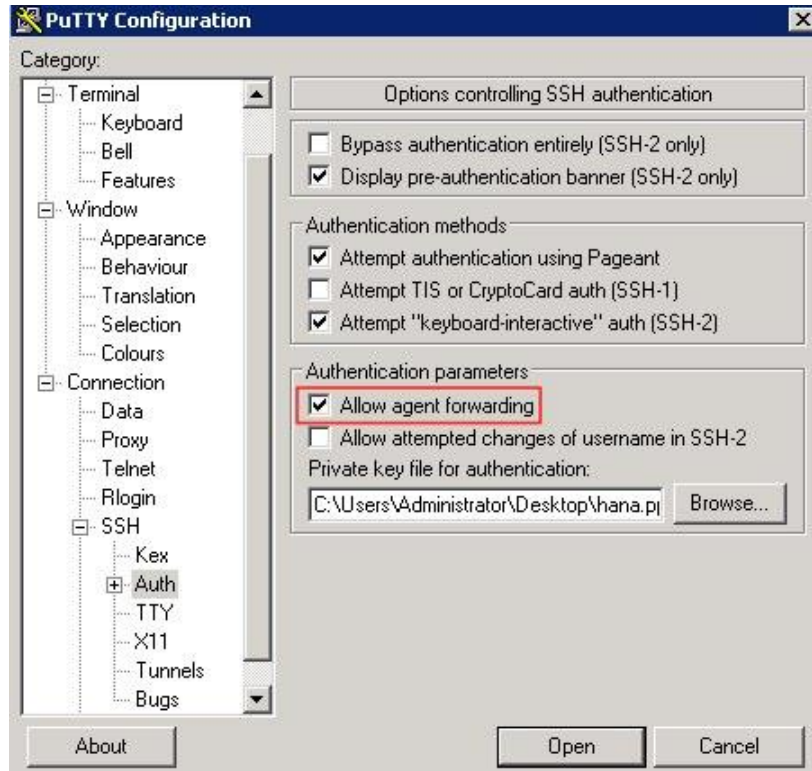
Figure 15: PuTTY Example for SSH Connection

5. Save the configuration.

6. Open up the connection to the bastion host by using SSH with the `ec2-user` user ID.

7. Connect to the SAP HANA server by using SSH.

Step 6. Perform Post-Deployment Tasks

Before you start using your SAP HANA instance, we recommend that you perform the following post-deployment steps. For details, refer to the SAP HANA Installation and Update Guide on the SAP website.

- If you have chosen SLES as the operating system for your SAP HANA instance, make sure that the Linux kernel version is at least **3.12.51** to avoid possible XFS corruption. If your kernel version is less than 3.12.51, update the kernel to the minimum required version. For details, see SAP OSS Notes 2205917 and 2246163.

- Update your SAP HANA software with the latest patches.

- Install any additional components such as Application Function Libraries (AFL) or Smart Data Access (SDA).

- Configure and back up your newly installed SAP HANA database.
- Consider creating a scheduled snapshot for your backup volume. For instructions, see Take Scheduled EBS Snapshots in the Amazon CloudWatch documentation.

Troubleshooting

If you run into any problems deploying this Quick Start, review the following FAQ for troubleshooting tips and guidance.

Where are the logs that monitor the Quick Start deployment progress?

You can find the deployment log in the /root/install/ folder of the SAP HANA instance. The name of the log file is install.log. You can log in to the SAP HANA instance as soon as you see that it's in the running state, and the instance passes the status checks in the Amazon EC2 console.

I launched the SAP HANA Quick Start template for a new Amazon VPC, and I see two additional templates being launched in the AWS CloudFormation console. Why?

When you launch the SAP HANA Quick Start for a new VPC, it launches two templates: one template to set up your network infrastructure (Amazon VPC, subnets, managed NAT gateway, and so on), and a second template to deploy and configure your SAP HANA instance(s), once the network infrastructure is in place.

Can I enable Multitenant Database Containers (MDC) as part of this SAP HANA Quick Start?

By default, the Quick Start installs SAP HANA software with the single-tenant option. If you need MDC configured, we recommend that you skip the SAP HANA software installation during deployment (set the SAP HANA installation parameter to No). Once the deployment is complete, you can manually install your SAP HANA software with MDC or other advanced options.

Where is my SAP HANA software staged when downloaded from the S3 bucket?

The unextracted SAP HANA software is downloaded to the /media/compressed/ folder, and the extracted files are stored in the /media/extracted/ folder.

My deployment status says ROLLBACK_COMPLETE and my deployment failed with a CREATE_FAILED error. What should I do?

If you encounter a **CREATE_FAILED** error, refer to the following table for known issues and solutions.

Error message	Possible cause	What to do
API: ec2: RunInstances Not authorized for images: *ami-ID*	The template is referencing an AMI that has expired.	We refresh AMIs on a regular basis, but our schedule isn't always synchronized with AWS AMI updates. If you get this error message, notify us, and we'll update the template with the new AMI ID. If you'd like to fix the template yourself, you can download it and update the `Mappings` section with the latest AMI ID for your region.
The instance configuration for this AWS Marketplace product is not supported. Please see *link* for more information about supported instance types, regions, and operating systems.	You are trying to launch a RHEL/SLES Marketplace AMI with an instance type that isn't supported.	Check your instance type and try to relaunch it with a supported instance type. If you want to extend the support for your desired instance type, contact the support team and open a support case.
X1 instance type requires a minimum Linux kernel version of 3.10. Choose the right operating system and try again.	You are trying to launch the X1 instance type with an unsupported version of the Linux operating system.	The X1 instance type for SAP HANA supports SLES version 12 and above. You can choose SLES 12 or SLES 12 SP1 for your deployment with the X1 instance type. If you're using R3 instances, you have two additional operating system choices: SLES 11 SP4 and RHEL 6.6.
Signal-failure function not implemented.	Deployment failed for an unknown reason.	Contact the [support team](#) and open a support case.
Your requested instance type (*type*) is not supported in your requested Availability Zone (*zone*). Please retry your request by not specifying an Availability Zone or choosing (*zone*).	The instance type you selected is not available in the chosen Availability Zone. (Your subnet is associated with a single Availability Zone.)	If you are trying to deploy SAP HANA into an existing VPC, create a new private subnet in the suggested Availability Zone and associate a route table with it. Retry the deployment by choosing the newly created private subnet. If you are trying to deploy SAP HANA in a new VPC and you receive this message, choose the suggested Availability Zone when you launch the template.
Not able to access SUSE (or Red Hat) update repository, package installation may fail.	The SAP HANA instance is unable to access the SUSE or RHEL update repository to download OS packages. The possible cause could be that Internet traffic for the SAP	See if it is possible to temporarily route the Internet traffic by using a NAT instance or NAT gateway. If your Internet traffic has to go through your internal proxy, contact your network team for access to the SUSE or RHEL update repository.

Error message	Possible cause	What to do
	HANA instance is not routed through a NAT instance or NAT gateway.	For further assistance, open a support case in the AWS Support Center.
We currently do not have sufficient *instance-type* capacity in the AZ you requested.	The Availability Zone you're deploying into doesn't have enough capacity or doesn't support the instance type.	Retry the deployment with a different instance type, or choose a subnet in a different Availability Zone.
Instance *ID* did not stabilize.	You have exceeded your IOPS for the region.	Request a limit increase by completing the request form in the AWS Support Center.
WaitCondition timed out. Received 0 conditions when expecting 1.	No or slow Internet access is causing delays when downloading SAP HANA software.	The Quick Start requires Internet access to download the setup scripts from an S3 bucket. This process will hang and eventually time out if Internet access isn't available. Make sure that you are deploying your SAP HANA instance in the private subnet and that you have access to the Internet either via a NAT gateway or a NAT instance. If you are routing your Internet traffic via your corporate proxy, check with your network team and ensure that firewall settings are set properly to allow such access.

You might also get this timeout error if you have an existing Amazon DynamoDB table from a previous Quick Start deployment. In this case, the table might contain duplicate entries with the same host name but different IP addresses, and the deployment script might fail to read it properly. Delete the table and retry the deployment. |
| **The HANA installation did not succeed. Please check installation media.** | SAP HANA installation failed or SAP HANA services didn't start up successfully. | Verify that you have staged the SAP HANA software properly in the S3 bucket with correct permissions. (See step 3 for details.)

Another reason could be that SAP HANA services did not start up after the installation.

In either case, consider redeploying your instance with the SAP HANA installation parameter set to No. The Quick Start redeployment will skip the SAP HANA installation, and you can manually install the SAP HANA software to troubleshoot the issue. |
| **System Administrator password must contain at least 8 characters.** | The SAP HANA master password contains $ or other special characters. | Change the SAP HANA master password and then relaunch the Quick Start.

The password must be at least 8 characters, consisting of uppercase and lowercase letters and numbers. Avoid using special characters such as @ or $. |

If the problem you encounter isn't covered in the previous table, we recommend that you relaunch the template with **Rollback on failure** set to **No** (this setting is under **Advanced** in the AWS CloudFormation console, **Options** page) and open a support case in the AWS Support Center for further troubleshooting. When rollback is disabled, the stack's state will be retained and the instance will be left running, so the support team can help troubleshoot the issue.

> **Important** When you set **Rollback on failure** to **No**, you'll continue to incur AWS charges for this stack. Please make sure to delete the stack when you've finished troubleshooting.

Support

If you encounter an issue deploying this Quick Start, check the Troubleshooting section first to see if the issue is covered. If it isn't, or the suggested solution doesn't resolve the issue, open a support case in the AWS Support Center. Assistance with SAP HANA deployment issues requires a subscription to the Business support plan.

If you're opening a support case, please attach the install.log file from the SAP HANA master instance (this is the log file that is located in the /root/install/ folder) to the ticket.

Security

The AWS cloud provides a scalable, highly reliable platform that helps enable customers to deploy applications and data quickly and securely.

When you build systems on the AWS infrastructure, security responsibilities are shared between you and AWS. This shared model can reduce your operational burden as AWS operates, manages, and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the services operate. In turn, you assume responsibility and management of the guest operating system (including updates and security patches), other associated application software such as SAP HANA, as well as the configuration of the AWS-provided security group firewall. For more information about security on AWS, visit the AWS Security Center.

Network Security

The default network security setup of this solution follows security best practices of AWS. The provisioned SAP HANA instances can be accessed only in three ways:

- By connecting to either the SAP HANA Studio Windows instance by using a remote desktop client, or to the bastion host by using SSH.

- From the CIDR block specified as `RemoteAccessCIDR` during the provisioning process.

- Alternatively, access can be restricted to a known CIDR block if a provisioned VPN tunnel exists between your own data center and AWS.

Identity and Access Management (IAM)

This solution leverages an IAM role with least privileged access. It is not necessary or recommended to store SSH keys or secret keys or access keys on the provisioned instances.

OS Security

The root user on Linux or the administrator on the Windows RDP instance can be accessed only by using the SSH key specified during the deployment process. AWS does not store these SSH keys, so if you lose your SSH key, you can lose access to these instances. Operating system patches are your responsibility and should be performed on a periodic basis.

Security Groups

A *security group* acts as a 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.

The security groups created and assigned to the individual instances as part of this solution are restricted as much as possible while allowing access to the various functions of SAP HANA. See Appendix C for a complete list of ports and protocols configured as part of this solution.

Additional Resources

AWS services

- AWS CloudFormation
  http://aws.amazon.com/documentation/cloudformation/

- Amazon EBS

- – User guide
  http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html

- – Volume types
  http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html

- – Optimized instances
  http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSOptimized.html

- Amazon EC2
  - – User guide for Microsoft Windows
    http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/

  - – User guide for Linux
    http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/

  - – X1 instances
    http://aws.amazon.com/ec2/instance-types/x1/

- Amazon VPC
  http://aws.amazon.com/documentation/vpc/

SAP HANA documentation

- SAPPartnerEdge
  https://partneredge.sap.com

- SAP notes and Knowledge Base articles
  https://support.sap.com/notes

SAP HANA on AWS

- SAP HANA on AWS Implementation and Operations Guide
  http://d0.awsstatic.com/enterprise-marketing/SAP/SAP_HANA_on_AWS_Implementation_and_Operations_Guide.pdf

- High Availability and Disaster Recovery Options for SAP HANA on AWS
  http://d0.awsstatic.com/enterprise-marketing/SAP/sap-hana-on-aws-high-availability-disaster-recovery-guide.pdf

- Setting up AWS Resources and SLES for SAP HANA Installation
  http://d0.awsstatic.com/enterprise-marketing/SAP/SAP-HANA-on-AWS-Manual-Setup-Guide.pdf

- Migrating SAP HANA Systems to X1 Instances on AWS
  http://do.awsstatic.com/enterprise-marketing/SAP/migrating-sap-hana-to-x1-on-aws.pdf

- Additional information about SAP solutions on AWS
  http://aws.amazon.com/sap/whitepapers/

Quick Start reference deployments

- SAP Business One, version for SAP HANA, Quick Start
  https://docs.aws.amazon.com/quickstart/latest/sap-b1-hana/

- Additional reference deployments
  https://aws.amazon.com/quickstart/

Appendix A: Subscribing to RHEL for SAP HANA

1. Log in to the AWS Marketplace at http://aws.amazon.com/marketplace.

2. Type **RHEL for SAP HANA** in the search box, and then choose **Red Hat Enterprise Linux for SAP HANA** from the results.
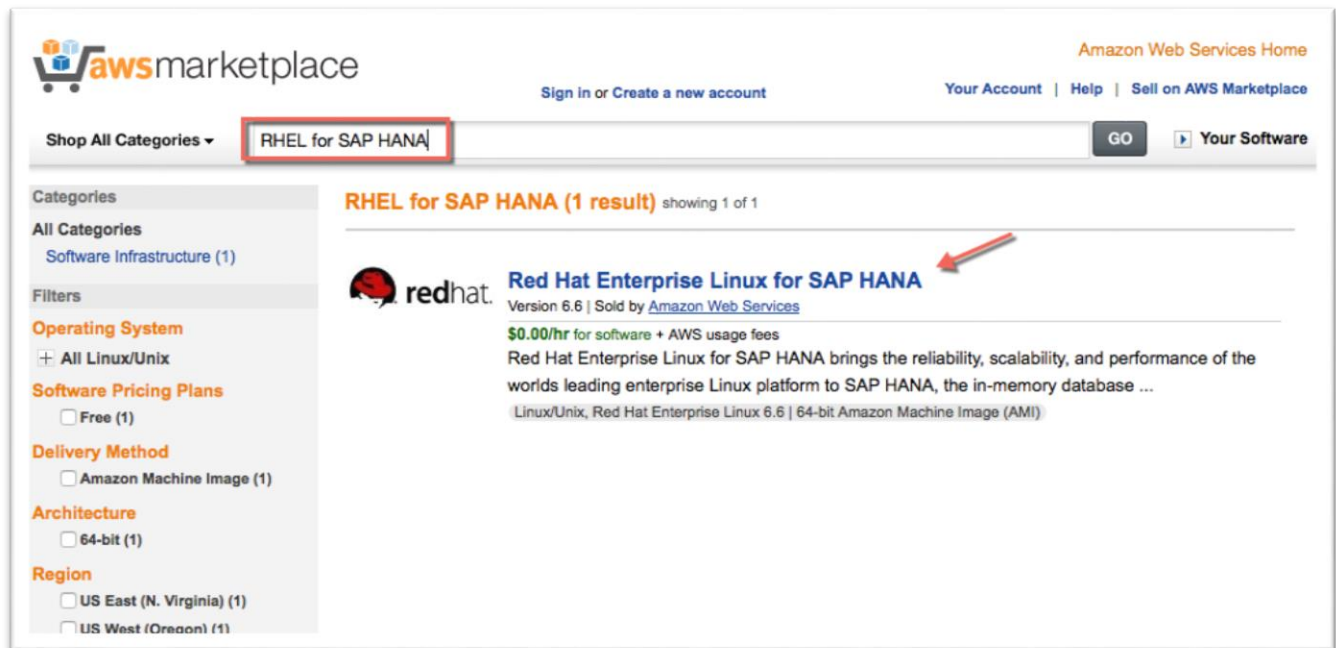
Figure 16: RHEL for SAP HANA in the AWS Marketplace

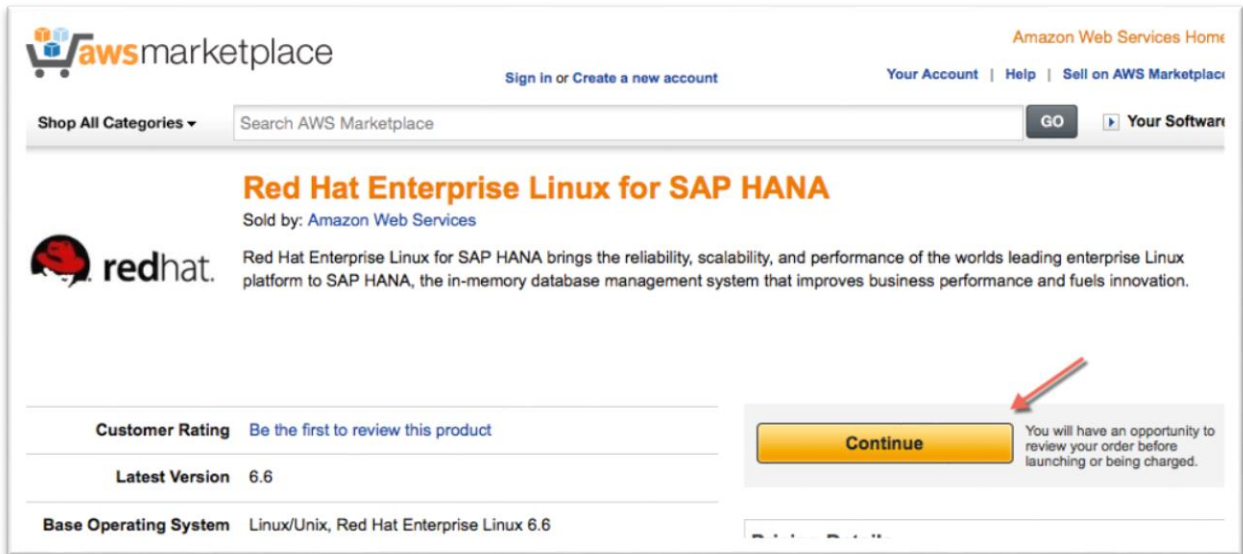3. From the **Red Hat Enterprise Linux for SAP HANA** page, choose **Continue**.

Figure 17: Choosing RHEL for SAP HANA

4. Choose the **Manual Launch** tab. Read the terms and conditions of software usage, and then choose **Accept Terms**.
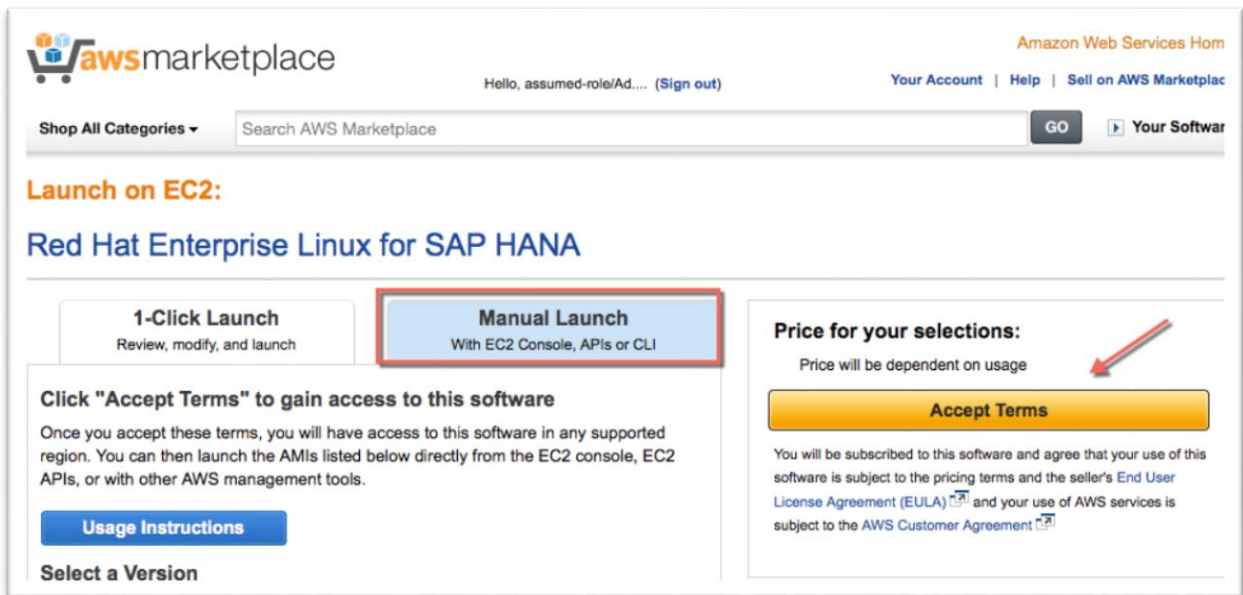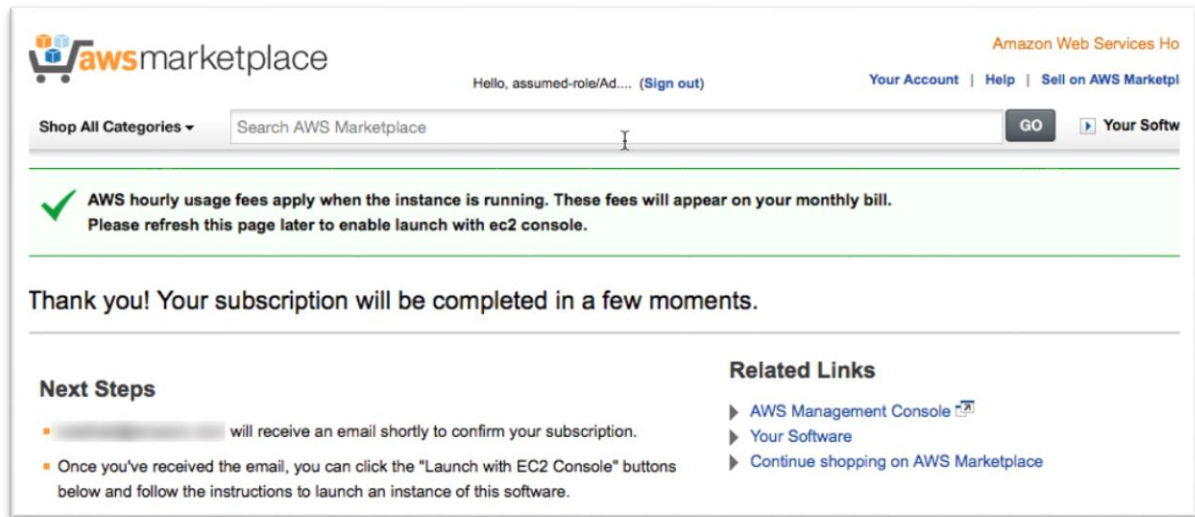

Figure 18: Reviewing the Terms of the License Agreement

You will get a confirmation page confirming your subscription, and an email confirmation will be sent to the account owner.

Figure 19: Completing the Subscription Request

5. When you receive the confirmation email for your subscription, proceed with the SAP HANA reference deployment.

Appendix B: Enhancements in This Release

AWS is constantly working to improve the design and ease of use of these Quick Start solutions, and to incorporate suggestions and new requirements from AWS customers and partners. This version of the Quick Start simplifies the overall process of deploying fully functional SAP HANA systems on the AWS cloud while providing more flexibility. This latest version includes the following changes and enhancements.

Performance and security:

- Support for the latest X1 instance type.

- Provisioning of managed NAT gateway in place of NAT instance.

- Bastion host for SSH access with tightened security.

High availability and disaster recovery:

- Automatic recovery option, including Amazon CloudWatch alarms for the Amazon EC2 instance hosting your SAP HANA database.

Ease of use:

- Better user experience with categories and friendly labels for AWS CloudFormation input parameters.

- Simpler deployment process with no extraction of media files and no snapshots.

- Ability to use the SAP HANA media staged in your S3 bucket rather than an Amazon EBS snapshot.

New deployment options:

- Ability to provision SAP-certified infrastructure (Amazon EC2 and Amazon EBS) without installing the SAP HANA software. (This gives you the flexibility to install the SAP HANA software on your own if you prefer a custom installation.)

- Ability to choose the specific Linux operating system and version for your SAP HANA system without having to specify the Amazon Machine Image (AMI) IDs.

New configuration options:

- Ability to choose different storage options (`gp2` or `io1`) for your SAP HANA data and log volumes.

- Ability to encrypt your Amazon EBS volumes.

- Ability to provision the SAP HANA worker node name during a scale-out deployment.

- Optional provisioning of the Amazon EC2 instance with Windows Server as an RDP system.

Appendix C: Security Groups

The following tables show the configured inbound and outbound protocols and ports allowed for the various instances deployed by this Quick Start.

RDP Security Group			
Inbound			
Source	Protocol	Port Range (Service)	Comments
Restricted to CIDR block specified during deployment	TCP	3389 (RDP)	Allow inbound RDP access to Windows instances from your network (over the Internet gateway).
Outbound			
Destination	Protocol	Port Range	Comments
0.0.0.0/0	TCP	1 - 65535	Allow outbound access from RDP server to anywhere.

Bastion Host Security Group			
Inbound			
Source	Protocol	Port Range (Service)	Comments
Restricted to CIDR block specified during deployment	TCP	22 (SSH)	Allow inbound SSH access to Linux instances from your network (over the Internet gateway).
Outbound			
Destination	Protocol	Port Range	Comments
10.0.1.0/24	TCP	22 (SSH)	Allow SSH access from bastion host to the 10.0.1.0 subnet.

Outbound, cont'd			
Destination	Protocol	Port Range	Comments
0.0.0.0/0	TCP	80 (HTTP)	Allow outbound HTTP access from instances deployed in the Amazon VPC to anywhere.
0.0.0.0/0	TCP	443 (HTTPS)	Allow outbound HTTPS access from instances deployed in the Amazon VPC to anywhere.

SAP HANA Master and Worker** Security Groups			
Inbound *(## represents the SAP instance number)*			
Source	Protocol	Port Range (Service)	Comments
10.0.1.0/24	TCP	1 - 65535	Communication between instances within the private subnet.
10.0.1.0/24	TCP/UDP	111, 2049, 4000-4002	Ports used for NFS communication.
10.0.1.0/24	TCP	3##00 − 3##10	Database internal communication and SAP support access.
**10.0.1.0/24	TCP	22 (SSH)	Allow SSH access from other SAP HANA nodes.
10.0.2.0/24	TCP	22 (SSH)	Allow SSH access from bastion host placed in the public subnet.
10.0.2.0/24	TCP	1128 - 1129	Host agent access.
10.0.2.0/24	TCP	43##	Access to XSEngine (HTTPS) from the 10.0.2.0 subnet.
10.0.2.0/24	TCP	80##	Access to XSEngine (HTTP) from the 10.0.2.0 subnet.
10.0.2.0/24	TCP	8080 (HTTP)	Software Update Manager (SUM) access (HTTP).
10.0.2.0/24	TCP	8443 (HTTPS)	Software Update Manager (SUM) access (HTTPS).
10.0.2.0/24	TCP	3##15	Database client access.
10.0.2.0/24	TCP	3##17	Database client access.
10.0.2.0/24	TCP	5##13 − 5##14	Allow access for HANA Studio from RDP instance.
Outbound			
0.0.0.0/0	TCP	1 - 65535	Outbound access from SAP HANA master allowed to anywhere.

Send Us Feedback

We welcome your questions and comments. Please post your feedback on the AWS Quick Start Discussion Forum.

Document Revisions

Date	Change	Location
June 2016	Removed sign-up instructions for X1 instances. X1 instances are now available to all users.	Changes throughout guide
May 2016	Updated templates with the enhancements listed in Appendix B.	Changes throughout guide
December 2015	Added troubleshooting information	Troubleshooting section
June 2015	• You can now choose between the RHEL and SLES operating systems for SAP HANA deployment. • The Quick Start now supports multi-node configurations.	New parameter for RHEL support, new appendixes, and adjustments to storage size information
April 2015	Replaced SAP Business One notes with a link to the SAP Business One, version for SAP HANA, Quick Start deployment guide.	—
December 2014	• Updated approximate pricing information. • In the first template, changed the default type for NATInstanceType to t2.small to support the EU (Frankfurt) region. • Added clarifications for deploying SAP Business One, version for SAP HANA.	Changes throughout guide
July 2014	Initial publication	—

Notices