

SAP HANA on the AWS Cloud

Quick Start Reference Deployment

*Sabareesan Radhakrishnan, Guilherme Felix, and Somckit Khemmanivanh
Solutions Architects, AWS Partner Program*

July 2014

Last updated: May 2020 (see [revisions](#))

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

*Supports SAP HANA Platform Edition 1 SPS 9–12
and SAP HANA Platform Edition 2 SPS 0-4*



Contents

Overview	4
SAP HANA on AWS	4
Quick Links	5
Cost and Licenses	5
AWS Services	7
Architecture	9
Deployment Scenarios	10
Single-AZ, Single-Node Architecture	11
Single-AZ, Multi-Node Architecture	12
Multi-AZ (HA), Single-Node Architecture	13
Planning the Deployment	14
AWS Instance Type for SAP HANA	14
Storage Configuration for SAP HANA	15
Memory Sizing for Deployment	22
Operating System for Deployment	24
Requirements for Multi-AZ, Single-Node HA Scenarios	25
Deployment Steps	26
Step 1. Prepare Your AWS Account	26
Step 2. Subscribe to the Operating System Image	29
Step 3. Download the SAP HANA Software	29
Step 4. Launch the Quick Start	34
Single-AZ Deployment	35
Multi-AZ Deployment	40
Step 5. Access SAP HANA Nodes	45
Using SAP HANA Studio	46
Using OS-Level Access	47
Step 6. Perform Post-Deployment Tasks	49

Troubleshooting	51
Support	56
Security	56
Network Security.....	56
Identity and Access Management (IAM).....	56
Operating System Security	57
Security Groups.....	57
Additional Resources	57
Appendix A: Subscribing to RHEL for SAP	59
Appendix B: Subscribing to SLES for SAP.....	60
Appendix C: Security Groups.....	61
Send Us Feedback	63
Document Revisions.....	63

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 AWS.

[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 EC2, Amazon EBS, and the operating system—SUSE Linux Enterprise Server (SLES) or Red Hat Enterprise Linux (RHEL)—are optimally configured to achieve the best performance for your SAP HANA system right out of the box.

This Quick Start supports the R3, R4, R5, and X1 instance families as well as Amazon EC2 High Memory instances with 6 TB, 9 TB, 12 TB, 18 TB, and 24 TB of memory, to boost the performance of large-scale, memory-intensive SAP HANA workloads on AWS. To learn more about High Memory instances, see the [Amazon EC2 High Memory Instances](#) webpage. To learn more about X1 instances, see [Amazon EC2 X1 Instances](#) and the [Amazon EC2 User Guide](#).

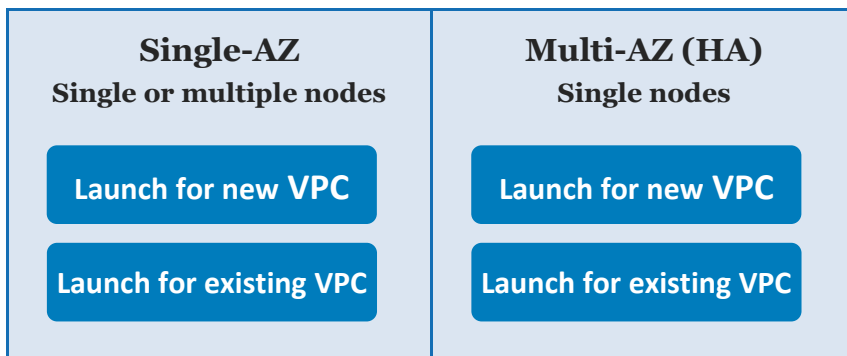
Note In this Quick Start, we use the terms *SAP HANA platform*, *system*, *database*, *server*, *node*, and *instance*, depending on context, to refer to the SAP HANA deployment on AWS.

This Quick Start currently supports SAP HANA Platform Edition 1 SPS 9x–12x, and SAP HANA Platform Edition 2 SPS 0–4.

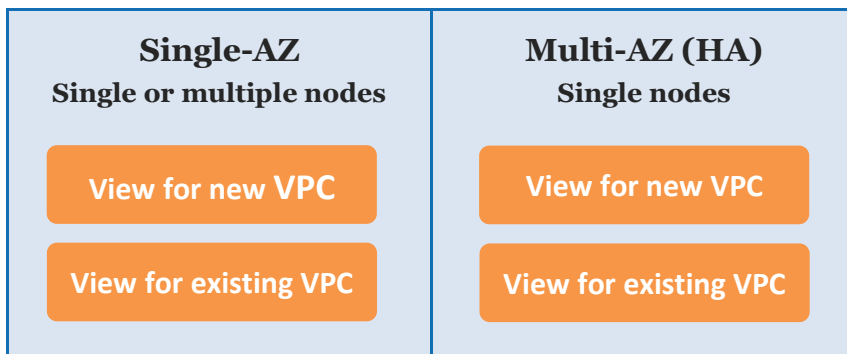
Quick Links

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 one of the environments shown in the [Architecture](#) section in a new or existing virtual private cloud (VPC). Each deployment takes 25-60 minutes to complete, depending on the selected configuration. If you're new to AWS or SAP HANA, please review the architectures, read the implementation details, and follow the step-by-step instructions provided in this guide.



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



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 r4.16xlarge and x1.32xlarge instances.

Note The prices in the table cover four of many possible configuration options and do not include NAT gateway or data transfer costs. You can use the [AWS Pricing Calculator](#) 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 AWS Region in which the AWS resources are launched.

SAP HANA instance	Scenario	Includes
r4.16xlarge	Single-AZ, single-node (scale-up)	Bastion host: t2.small RDP instance: c4.xlarge SAP HANA: Single instance, SLES, r4.16xlarge Storage: General Purpose (SSD) for SAP HANA data and logs
x1.32xlarge	Single-AZ, single-node (scale-up)	Bastion host: t2.small RDP instance: c4.xlarge SAP HANA: Single instance, SLES, x1.32xlarge Storage: Provisioned IOPS (SSD) for SAP HANA data and logs
r4.16xlarge	Multi-AZ, single-node, high availability setup	Bastion host: t2.small RDP instance: c4.xlarge SAP HANA: Two instances, Linux**, r4.16xlarge Storage: General Purpose (SSD) for SAP HANA data and logs
x1.32xlarge	Multi-AZ, single-node, high availability setup	Bastion host: t2.small RDP instance: c4.xlarge SAP HANA: Two instances, Linux**, x1.32xlarge Storage: Provisioned IOPS (SSD) for SAP HANA data and logs

* Cost estimates shown in the AWS Simple Monthly Calculator are based on pricing for On-Demand Instances. You could achieve a discount of up to 75% for your EC2 instances if you use Reserved Instances. For details, see the [Amazon EC2 Reserved Instances](#) page.

** The cost estimates for Multi-AZ, single-node, high availability setup are based on the Linux operating system, because SLES for SAP is available with the BYOS (Bring Your Own Subscription) model or must be purchased from [AWS Marketplace](#).

These deployments use 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.

The Single-AZ deployment option launches the Amazon Machine Image (AMI) for the operating system you choose—SLES, SLES for SAP, SLES for SAP BYOS, RHEL for SAP HANA, or RHEL for SAP with High Availability (HA) and Update Services (US). The license cost for the SLES or RHEL for SAP HANA operating system is included in the Amazon EC2 hourly price. There is an additional software cost for SLES for SAP and RHEL for SAP with HA and US AMI subscriptions in AWS Marketplace. For SLES for SAP BYOS AMIs, you can use your existing valid subscriptions.

The Multi-AZ (HA) deployment option supports operating systems SLES for SAP, SLES for SAP BYOS, and RHEL for SAP with HA and US, which provide packages for configuring pacemaker-based cluster for High Availability solutions. You must have subscribed to the AWS Marketplace offering for [SLES for SAP](#) and [RHEL for SAP with HA and US](#), or have a registration code (if you choose SLES for SAP BYOS) before you start deployment.

AWS Services

The core AWS components used by this Quick Start include the following AWS services. (If you are new to AWS, see [Getting Started](#) on the AWS website.)

- [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 EC2 instances in the AWS Cloud. Each EBS volume is automatically replicated within its Availability Zone to protect you from component failure, offering high availability and durability. 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., 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.
- [AWS CloudTrail](#) – AWS CloudTrail enables governance, compliance, operational auditing, and risk auditing of your AWS account. With CloudTrail, you can log, continuously monitor, and retain account activity related to actions across your AWS infrastructure. You can use [CloudTrail trails](#) to deliver CloudTrail events to an S3 bucket.
- [AWS Config](#) – AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. With AWS Config, you can review changes in configurations and relationships between AWS resources, dive into detailed resource configuration histories, and determine your overall compliance against the configurations specified in your internal guidelines.
- [AWS Lambda](#) – AWS Lambda is a serverless compute service that runs your code in response to events and automatically manages the underlying compute resources for

you. You can use AWS Lambda to extend other AWS services with custom logic, or create your own back-end services that operate at AWS scale, performance, and security. AWS Lambda can automatically run code in response to multiple events, such as HTTP requests via Amazon API Gateway, modifications to objects in Amazon S3 buckets, table updates in Amazon DynamoDB, and state transitions in AWS Step Functions.

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.

The Quick Start deploys and configures the following components:

- A virtual private cloud (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.
- EC2 instance(s) to host the SAP HANA database, with the appropriate EBS volumes configured to meet or exceed SAP HANA storage key performance indicators (KPIs). The Single-AZ configuration gives you a choice of Linux operating system (SLES, SLES for SAP, SLES for SAP BYOS, RHEL for SAP HANA, or RHEL for SAP with HA and US). The Multi-AZ configuration supports SLES for SAP, SLES for SAP BYOS, and RHEL for SAP with HA and US.
- A bastion host in the public subnet with an Elastic IP address to allow inbound SSH (Secure Shell) access to the EC2 instance(s) that host the SAP HANA database.
- A managed NAT gateway to allow outbound internet access for resources in the private subnet.
- IAM instance roles with fine-grained permissions for access to AWS services.
- 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 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.

- An optional configuration of AWS CloudTrail trails and AWS Config to help with compliance, audit, and change management requirements.
- The Multi-AZ, high availability Quick Start deploys two SAP HANA nodes in two different subnets in separate Availability Zones, and configures the SAP HANA instances for high availability. In this scenario, both SAP HANA nodes share the same EC2 instance size, SAP HANA SID, SAP HANA system number, and other characteristics.

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 templates provided with this Quick Start bootstrap the AWS infrastructure and automate the deployment of SAP HANA on the AWS Cloud from scratch. The Quick Start provides these deployment scenarios:

- **Single-AZ deployment (single-node and multi-node):** This scenario uses a single Availability Zone for the deployment. The single-node option provisions a single EC2 instance for SAP HANA in the private subnet of the Availability Zone. The multi-node option provisions up to five EC2 instances for SAP HANA in the private subnet. You can choose from two additional deployment options:
 - Deploying SAP HANA into a new VPC (end-to-end deployment) builds the VPC, subnets, NAT gateway, security groups, bastion host, and optional Windows Server and SAP HANA server(s) with Amazon EBS.
 - Deploying SAP HANA into an existing VPC provisions SAP HANA servers in your existing infrastructure.
- **Multi-node, high-availability deployment (single-node):** This scenario uses two Availability Zones for the SAP HANA environment. It provisions two SAP HANA servers (primary and secondary) in separate private subnets in those zones and configures them for high availability. You can choose from two additional deployment options:
 - Deploying SAP HANA into a new VPC (end-to-end deployment) builds the VPC, subnets, NAT gateway, security groups, bastion host, optional Windows Server, and two SAP HANA servers (primary and secondary) with Amazon EBS.
 - Deploying SAP HANA into an existing VPC provisions two SAP HANA servers in your existing infrastructure and configures them for high availability.

The following sections discuss the Quick Start architecture for each scenario.

Single-AZ, Single-Node Architecture

The Single-AZ, single-node deployment option provisions a single 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 a 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.

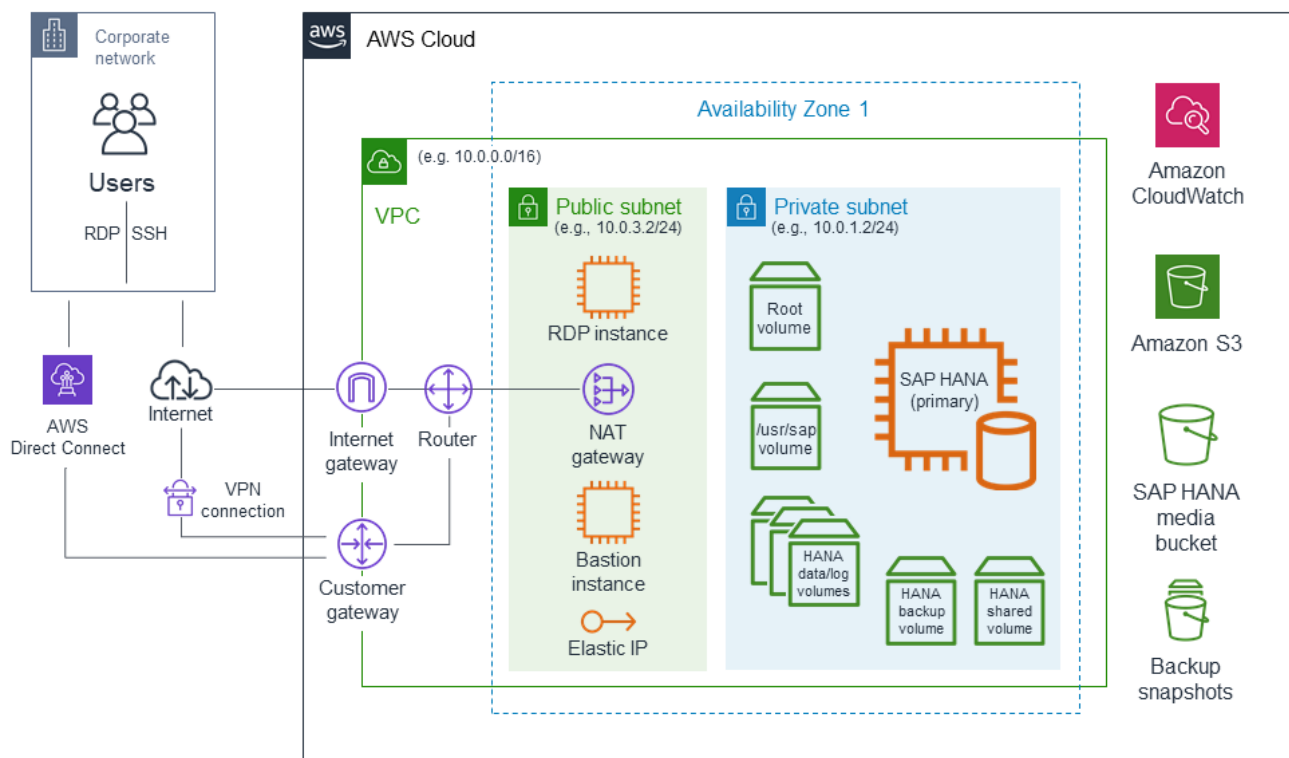


Figure 1: Single-AZ, single-node architecture for SAP HANA on AWS

Single-AZ, Multi-Node Architecture

The following Amazon EC2 instances are certified by SAP to deploy scale-out SAP HANA clusters for online analytic processing (OLAP) workloads.

Instance type	Number of scale-out nodes
u-6tb1.metal	16 nodes (general availability)
x1.16xlarge	7 nodes (general availability)
x1.32xlarge	25 nodes (general availability)
x1e.32xlarge	25 nodes (general availability)
r3.8xlarge	5 nodes (general availability) 17 nodes (controlled availability)
r5.24xlarge	16 nodes (general availability)

The following EC2 instances are certified by SAP to deploy scale-out SAP HANA clusters for online transaction processing (OLTP) workloads.

Instance type	Number of scale-out nodes
u-6tb1.metal	4 nodes (general availability)
u-12tb1.metal	4 nodes (general availability)

The Single-AZ, multi-node deployment option provisions up to five EC2 instances (see note) 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.

Note If you would like to deploy SAP HANA scale-out clusters that are larger than 5 nodes, contact us at saphana@amazon.com. We will provide you with an AWS CloudFormation template that is designed to help deploy SAP HANA clusters that are larger than 5 nodes.

As in single-node deployment, access to all SAP HANA servers is provided either through AWS Systems Manager [Session Manager](#), 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](#).

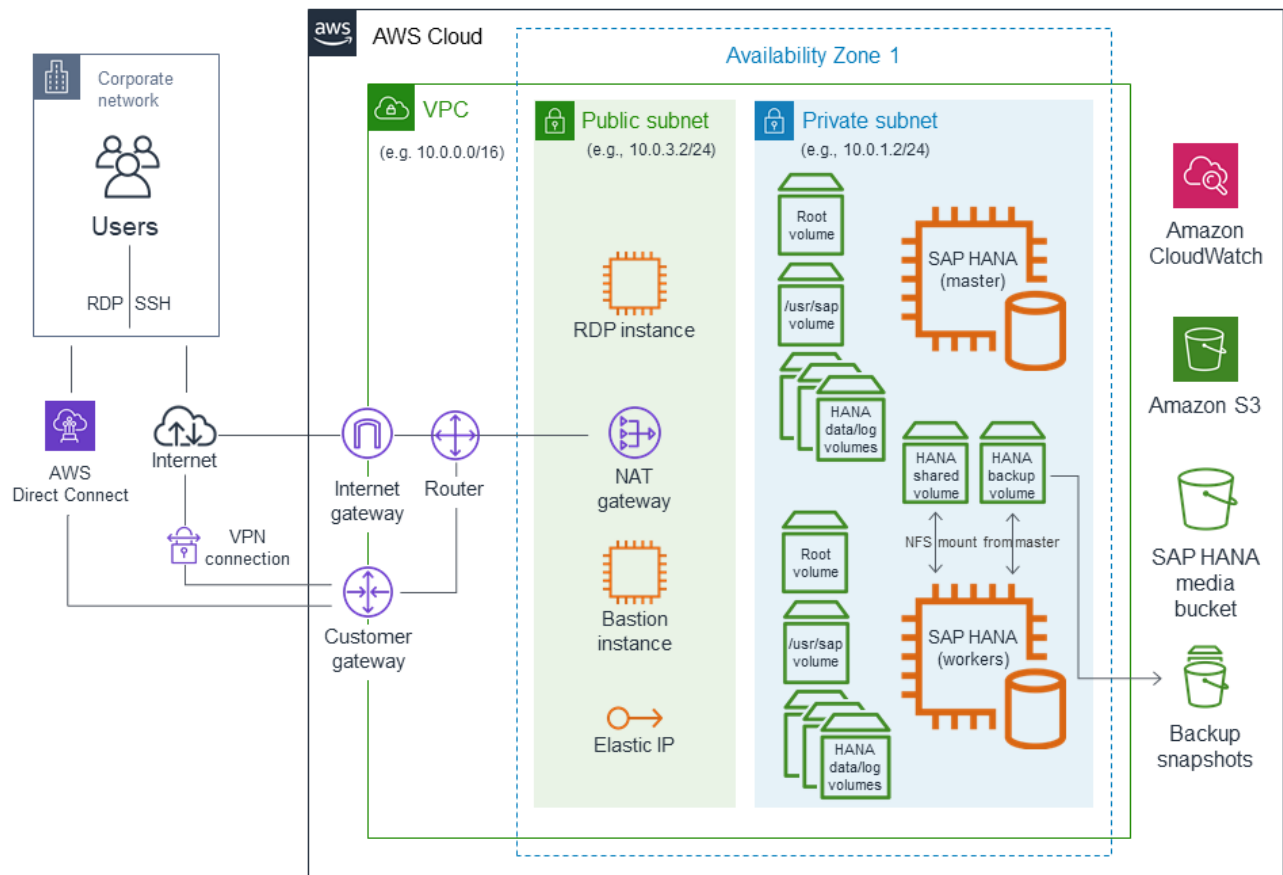


Figure 2: Single-AZ, multi-node architecture for SAP HANA on AWS

Multi-AZ (HA), Single-Node Architecture

The Multi-AZ (HA), single-node deployment option provisions two EC2 instances in private subnets in two different Availability Zones, with your choice of Amazon EBS storage. For this architecture, you can choose SLES High Availability Extension (HAE) or RHEL for SAP High Availability solutions. There are two main scenarios for an SAP HANA high availability deployment:

- **Performance-optimized:** In this scenario, both primary and secondary nodes have the same size and use synchronous replication mode for SAP HANA System Replication (HSR). The secondary instance is configured to preload tables.
- **Cost-optimized:** In this scenario, the secondary node can be used for production as well as non-production (for example, QA) work. Before takeover by production, the non-production instance running on the secondary node has to be stopped to provide resources for production. In this scenario, you can configure synchronous replication mode; however, table preload needs to be switched off on the secondary node.

This Quick Start supports the performance-optimized scenario.

The Quick Start follows security best practices by establishing a VPC with public and private subnets. For increased SUSE cluster resiliency, a secondary corosync communication ring (using the redundant ring protocol) is created by adding a secondary IP address (on the same subnet) to the existing Amazon EC2 [elastic network interfaces](#). Refer to section 4.1 to 4.3 of this SUSE [document](#) for information about the redundant ring protocol. 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 Session Manager, the bastion host, or an SSH client on the optional Windows Server instance.

The Multi-AZ, single-node, high availability setup is based on [SAP Note 2309342 – SUSE Linux Enterprise High Availability Extension on AWS for SAP HANA](#).

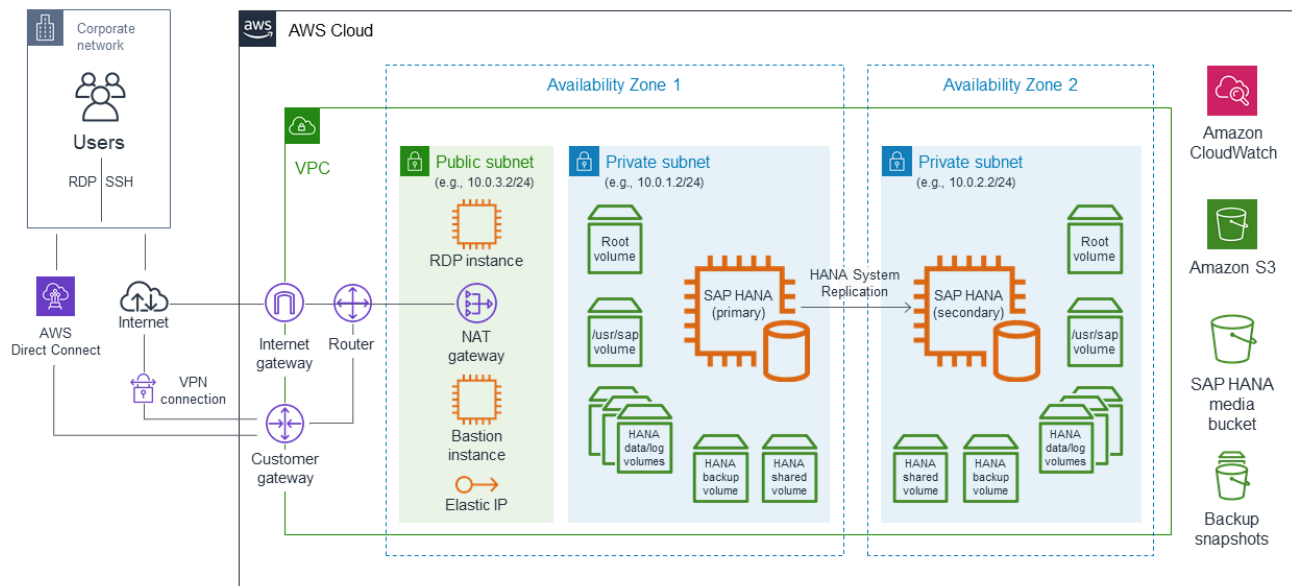


Figure 3: Multi-AZ, single-node, high availability architecture for SAP HANA on AWS

Planning the Deployment

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

AWS Instance Type for SAP HANA

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

- u-6tb1.metal, u-9tb1.metal, u-12tb1.metal, u-18tb1.metal, and u-24tb1.metal
- x1.16xlarge, x1.32xlarge, x1e.xlarge*, x1e.2xlarge*, x1e.4xlarge*, and x1e.32xlarge
- r3.2xlarge*, r3.4xlarge*, and r3.8xlarge
- r4.2xlarge* (default), r4.4xlarge*, r4.8xlarge, and r4.16xlarge
- r5.2xlarge*, r5.4xlarge*, r5.12xlarge, r5.24xlarge, and r5.metal

* These instance types aren't supported for SAP HANA production workloads.

The default instance type is r5.2xlarge, but you can switch to one of the other types during deployment. The r3.8xlarge, r4.8xlarge, r4.16xlarge, r5.12xlarge, r5.24xlarge, r5.metal, x1.16xlarge, x1.32xlarge, x1e.32xlarge, u-6tb1.metal, u-9tb1.metal, u-12tb1.metal, u-18tb1.metal, and u-24tb1.metal instance types are supported by SAP for production use. For more information about different instance types and their use cases, see the [Amazon EC2 Instances](#) webpage. For a list of SAP-certified EC2 instances, see the [SAP HANA hardware directory](#).

The Amazon EC2 High Memory instance types u-6tb1.metal, u-9tb1.metal, u-12tb1.metal, u-18tb1.metal, and u-24tb1.metal are offered as bare-metal instances on Amazon EC2 Dedicated Hosts. If you're planning to use these instances, you must have a Dedicated Host for Amazon EC2 High Memory Instances allocated to your account by AWS before proceeding with your deployment. For more information, see the [Amazon EC2 High Memory 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.

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 highly recommend using the io1 configuration for your mission-critical SAP HANA workloads for production use.

Note that only u-6tb1.metal, u-9tb1.metal, u-12tb1.metal, , u-18tb1.metal, u-24tb1.metal, x1e.32xlarge, x1.32xlarge, x1.16xlarge, r5.24xlarge, r5.12xlarge, r5.metal, r4.16xlarge,

r4.8xlarge, and r3.8xlarge instances are certified for production use. You can use all the instance types supported by this Quick Start for non-production use.

For multi-node deployments, storage volumes for SAP HANA data and logs are provisioned in the master node and every worker node.

In the following configurations, we intentionally kept the same storage configuration for SAP HANA data and log volumes for all R3, certain R4 and R5, and smaller X1e instance types so you can scale up from smaller instances to larger instances without having to reconfigure your storage.

Configuration tables:

- [gp2-based storage configuration for SAP HANA data](#)
- [gp2-based storage configuration for SAP HANA logs](#)
- [io1-based storage configuration for SAP HANA data](#)
- [io1-based storage configuration for SAP HANA logs](#)

Footnotes for the following tables:

* Each logical processor offered by Amazon EC2 High Memory Instances is a hyperthread on a physical CPU core.

** This value represents the maximum throughput that could be achieved with the EBS volume configuration. Actual throughput depends on the instance type. Every instance type has its own Amazon EBS throughput maximum. For details, see [Amazon EBS-Optimized Instances](#) in the AWS documentation.

gp2-based storage configuration for SAP HANA data:

Instance type	Memory (GiB)	vCPUs / logical processors*	General Purpose SSD (gp2) storage for SAP HANA data (striped with LVM)	Total maximum throughput (MiB/s)	Total baseline IOPS	Total burst IOPS
Certified for production use						
u-24tb1.metal	24,576	448	6 x 4,800 GiB	1,500	86,400	N/A
u-18tb1.metal	18,432	448	6 x 3,600 GiB	1,500	64,800	N/A
u-12tb1.metal	12,288	448	4 x 3,600 GiB	1,000	43,200	N/A
u-9tb1.metal	9,216	448	4 x 2,700 GiB	1,000	32,400	N/A
u-6tb1.metal	6,144	448	4 x 1,800 GiB	1,000	21,600	N/A

x1e.32xlarge	3,904	128	3 x 1,600 GiB	750	14,400	N/A
x1.32xlarge	1,952	128	3 x 800 GiB	750	7,200	9,000
x1.16xlarge	976	64	3 x 400 GiB	750	3,600	9,000
r5.metal	768	96	3 x 400 GiB	750	3,600	9,000
r5.24xlarge	768	96	3 x 400 GiB	750	3,600	9,000
r4.16xlarge	488	64	3 x 225 GiB	750	2,025	9,000
r5.12xlarge	384	48	3 x 225 GiB	750	2,025	9,000
r4.8xlarge	244	32	3 x 225 GiB	750	2,025	9,000
r3.8xlarge						
Supported for non-production use only						
x1e.4xlarge	488	16	3 x 225 GiB	750**	2,025	9,000
x1e.2xlarge	244	8	3 x 225 GiB	750**	2,025	9,000
x1e.xlarge	122	4	3 x 225 GiB	750**	2,025	9,000
r5.4xlarge	128	16	3 x 225 GiB	750**	2,025	9,000
r5.2xlarge	64	8	3 x 225 GiB	750**	2,025	9,000
r4.4xlarge	122	16	3 x 225 GiB	750**	2,025	9,000
r3.4xlarge						
r4.2xlarge	61	8	3 x 225 GiB	750**	2,025	9,000
r3.2xlarge						

gp2-based storage configuration for SAP HANA logs:

Instance type	Memory (GiB)	vCPUs / logical processors*	General Purpose SSD (gp2) storage for SAP HANA logs (striped with LVM)	Total maximum throughput (MiB/s)	Total baseline IOPS	Total burst IOPS
Certified for production use						
u-24tb1.metal	24,576	448	2 x 300 GiB	500	1,800	6,000
u-18tb1.metal	18,432	448	2 x 300 GiB	500	1,800	6,000
u-12tb1.metal	12,288	448	2 x 300 GiB	500	1,800	6,000
u-9tb1.metal	9,216	448	2 x 300 GiB	500	1,800	6,000
u-6tb1.metal	6,144	448	2 x 300 GiB	500	1,800	6,000
x1e.32xlarge	3,904	128	2 x 300 GiB	500	1,800	6,000
x1.32xlarge	1,952	128	2 x 300 GiB	500	1,800	6,000
x1.16xlarge	976	64	2 x 300 GiB	500	1,800	6,000
r5.metal	768	96	2 x 300 GiB	500	1,800	6,000

r5.24xlarge	768	96	2 x 300 GiB	500	1,800	6,000
r4.16xlarge	488	64	2 x 300 GiB	500	1,800	6,000
r5.12xlarge	384	48	2 x 300 GiB	500	1,800	6,000
r4.8xlarge	244	32	2 x 300 GiB	500	1,800	6,000
r3.8xlarge						
Supported for non-production use only						
x1e.4xlarge	488	16	2 x 175 GiB	500**	1,050	6,000
x1e.2xlarge	244	8	2 x 175 GiB	500**	1,050	6,000
x1e.xlarge	122	4	2 x 175 GiB	500**	1,050	6,000
r5.4xlarge	128	16	2 x 175 GiB	500**	1,050	6,000
r5.2xlarge	64	8	2 x 175 GiB	500**	1,050	6,000
r4.4xlarge	122	16	2 x 175 GiB	500**	1,050	6,000
r3.4xlarge						
r4.2xlarge	61	8	2 x 175 GiB	500**	1,050	6,000
r3.2xlarge						

io1-based storage configuration for SAP HANA data:

Instance type	Memory (GiB)	vCPUs / logical processors *	Provisioned IOPS SSD (io1) storage for SAP HANA data (striped with LVM)	Total maximum throughput (MiB/s)	Total provisioned IOPS
Certified for production use					
u-24tb1.metal	24,576	448	6 x 4,800 GiB	3,000	18,000
u-18tb1.metal	18,432	448	6 x 3,600 GiB	3,000	18,000
u-12tb1.metal	12,288	448	4 x 3,600 GiB	2,000	9,000
u-9tb1.metal	9,216	448	4 x 2,700 GiB	2,000	9,000
u-6tb1.metal	6,144	448	4 x 1,800 GiB	2,000	9,000
x1e.32xlarge	3,904	128	3 x 1,600 GiB	1,500	9,000
x1.32xlarge	1,952	128	3 x 800 GiB	1,500	9,000
x1.16xlarge	976	64	1 x 1,200 GiB	500	7,500
r5.metal	768	96	1 x 1,200 GiB	500	7,500
r5.24xlarge	768	96	1 x 1,200 GiB	500	7,500
r4.16xlarge	488	64	1 x 600 GiB	500	7,500
r5.12xlarge	384	48	1 x 600 GiB	500	7,500

r4.8xlarge r3.8xlarge	244	32	1 x 300 GiB	500	7,500
Supported for non-production use only					
x1e.4xlarge	488	16	1 x 600 GiB	500**	2,000
x1e.2xlarge	244	8	1 x 300 GiB	500**	2,000
x1e.xlarge	122	4	1 x 300 GiB	500**	2,000
r5.4xlarge	128	16	1 x 300 GiB	500**	2,000
r5.2xlarge	64	8	1 x 300 GiB	500**	2,000
r4.4xlarge r3.4xlarge	122	16	1 x 300 GiB	500**	2,000
r4.2xlarge r3.2xlarge	61	8	1 x 300 GiB	500**	2,000

io1-based storage configuration for SAP HANA logs:

Instance type	Memory (GiB)	vCPUs / logical processors *	Provisioned IOPS SSD (io1) storage for SAP HANA logs (striped with LVM)	Total maximum throughput (MiB/s)	Total provisioned IOPS
Certified for production use					
u-24tb1.metal	24,576	448	1 x 525 GiB	500	2,000
u-18tb1.metal	18,432	448	1 x 525 GiB	500	2,000
u-12tb1.metal	12,288	448	1 x 525 GiB	500	2,000
u-9tb1.metal	9,216	448	1 x 525 GiB	500	2,000
u-6tb1.metal	6,144	448	1 x 525 GiB	500	2,000
x1e.32xlarge	3,904	128	1 x 525 GiB	500	2,000
x1.32xlarge	1,952	128	1 x 525 GiB	500	2,000
x1.16xlarge	976	64	1 x 525 GiB	500	2,000
r5.metal	768	96	1 x 525 GiB	500	2,000
r5.24xlarge	768	96	1 x 525 GiB	500	2,000
r4.16xlarge	488	64	1 x 260 GiB	500	2,000
r5.12xlarge	384	48	1 x 260 GiB	500	2,000
r4.8xlarge r3.8xlarge	244	32	1 x 260 GiB	500	2,000
Supported for non-production use only					
x1e.4xlarge	488	16	1 x 260 GiB	500**	1,000

x1e.2xlarge	244	8	1 x 260 GiB	500**	1,000
x1e.xlarge	122	4	1 x 260 GiB	500**	1,000
r5.4xlarge	128	16	1 x 260 GiB	500**	1,000
r5.2xlarge	64	8	1 x 260 GiB	500**	1,000
r4.4xlarge	122	16	1 x 260 GiB	500**	1,000
r3.4xlarge					
r4.2xlarge	61	8	1 x 260 GiB	500**	1,000
r3.2xlarge					

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	Memory (GiB)	vCPUs / logical processors*	Root volume (gp2)	SAP binaries (gp2)	SAP HANA shared** (gp2)	SAP HANA backup*** (st1)
Certified for production use						
u-24tb1.metal	24,576	448	1 x 50 GiB	1 x 50 GiB	1 x 1,024 GiB	2 x 16,384 GiB
u-18tb1.metal	18,432	448	1 x 50 GiB	1 x 50 GiB	1 x 1,024 GiB	2 x 16,384 GiB
u-12tb1.metal	12,288	448	1 x 50 GiB	1 x 50 GiB	1 x 1,024 GiB	1 x 16,384 GiB
u-9tb1.metal	9,216	448	1 x 50 GiB	1 x 50 GiB	1 x 1,024 GiB	1 x 16,384 GiB
u-6tb1.metal	6,144	448	1 x 50 GiB	1 x 50 GiB	1 x 1,024 GiB	1 x 12,288 GiB
x1e.32xlarge	3,904	128	1 x 50 GiB	1 x 50 GiB	1 x 1,024 GiB	1 x 8,192 GiB
x1.32xlarge	1,952	128	1 x 50 GiB	1 x 50 GiB	1 x 1,024 GiB	1 x 4,096 GiB
x1.16xlarge	976	64	1 x 50 GiB	1 x 50 GiB	1 x 1,024 GiB	1 x 2,048 GiB
r5.metal	768	96	1 x 50 GiB	1 x 50 GiB	1 x 1,024 GiB	1 x 2,048 GiB
r5.24xlarge	768	96	1 x 50 GiB	1 x 50 GiB	1 x 1,024 GiB	1 x 2,048 GiB
r4.16xlarge	488	64	1 x 50 GiB	1 x 50 GiB	1 x 512 GiB	1 x 1,024 GiB
r5.12xlarge	384	48	1 x 50 GiB	1 x 50 GiB	1 x 512 GiB	1 x 1024 GiB
r4.8xlarge	244	32	1 x 50 GiB	1 x 50 GiB	1 x 300 GiB	1 x 1,024 GiB
r3.8xlarge						

Supported for non-production use only						
x1e.4xlarge	488	16	1 x 50 GiB	1 x 50 GiB	1 x 512 GiB	1 x 1,024 GiB
x1e.2xlarge	244	8	1 x 50 GiB	1 x 50 GiB	1 x 300 GiB	1 x 512 GiB
x1e.xlarge	122	4	1 x 50 GiB	1 x 50 GiB	1 x 300 GiB	1 x 512 GiB
r5.4xlarge	128	16	1 x 50 GiB	1 x 50 GiB	1 x 300 GiB	1 x 512 GiB
r5.2xlarge	64	8	1 x 50 GiB	1 x 50 GiB	1 x 300 GiB	1 x 512 GiB
r4.4xlarge	122	16	1 x 50 GiB	1 x 50 GiB	1 x 300 GiB	1 x 512 GiB
r3.4xlarge						
r4.2xlarge	61	8	1 x 50 GiB	1 x 50 GiB	1 x 300 GiB	1 x 512 GiB
r3.2xlarge						

* Each logical processor offered by Amazon EC2 High Memory Instances is a hyperthread on a physical CPU core.

** In a multi-node architecture, the SAP HANA NFS shared volume is provisioned only once on the master node.

*** In a multi-node architecture, the SAP HANA backup volume can be deployed as NFS or Amazon EFS. The size of the SAP HANA NFS backup volume is multiplied by the number of nodes. The SAP HANA backup volume is provisioned only once on the master node, and NFS is mounted on the worker nodes. There is no provision needed for [Amazon EFS](#) as it is built to scale on demand, growing and shrinking automatically as files are added and removed.

Note General Purpose SSD (gp2) volumes created or modified after 12/03/2018 have a throughput maximum between 128 MiB/s and 250 MiB/s depending on volume size. Volumes greater than 170 GiB and below 334 GiB deliver a maximum throughput of 250 MiB/s if burst credits are available. Volumes with 334 GiB and above deliver 250 MiB/s, irrespective of burst credits. For details, see [Amazon EBS Volume Types](#) in the AWS documentation.

For single-node deployment, 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.

For multi-node deployment, SAP HANA backup deployed with [Amazon EFS](#) can support performance over 10 GB/sec and over 500,000 IOPS.

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.)

- 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:

- SUSE Linux Enterprise Server for SAP (SLES for SAP)
- SUSE Linux Enterprise Server (SLES)
- Red Hat Linux Enterprise Server for SAP with High Availability and Update Services (RHEL for SAP with HA and US)
- Red Hat Enterprise Linux for SAP HANA (RHEL for SAP HANA)

The following table summarizes the list of supported operating systems and versions.

Operating system	Supports Single-AZ deployment?	Supports Multi-AZ deployment?
<ul style="list-style-type: none"> • SLES 11 SP4 • SLES 12 • SLES 12 SP1 • SLES 12 SP3 • SLES 12 SP4 • SLES 15 	Yes	No
<ul style="list-style-type: none"> • SLES 12 SP2 for SAP* • SLES 12 SP3 for SAP* 	Yes	Yes
<ul style="list-style-type: none"> • SLES 12 SP4 for SAP* • SLES 12 SP5 for SAP* • SLES 15 for SAP 	Yes	Yes
<ul style="list-style-type: none"> • RHEL 7.5 for SAP HANA 	Yes	No
<ul style="list-style-type: none"> • RHEL 7.4 for SAP with HA and US • RHEL 7.5 for SAP with HA and US • RHEL 7.6 for SAP with HA and US 	Yes	Yes

* Both AWS Marketplace and Bring Your Own Subscription (BYOS) are supported. You must provide a registration key for the BYOS option.

We strongly recommend using SLES for SAP or RHEL for SAP with HA and US for your deployments. These operating systems provide the software required to set up high availability and offer an extended support life cycle.

To learn more about the benefits of using SLES for SAP, see the [SUSE product page](#). To learn more about the benefits of using RHEL for SAP with HA and US, see [Red Hat Enterprise Linux for SAP Offerings on AWS FAQ](#) on the Red Hat website.

Note Before you start your deployment, see the [SAP HANA hardware directory](#) to find the list of operating systems supported on a specific Amazon EC2 instance type for SAP HANA workloads.

SAP maintains a support matrix that details dependencies between operating system versions and SAP HANA versions. Review [SAP Note 2235581 – SAP HANA: Supported Operating Systems](#) to make sure that you are choosing a supported combination for your deployment.

Requirements for Multi-AZ, Single-Node HA Scenarios

If you choose the Multi-AZ, single-node, high availability deployment option, you must follow a few additional requirements. Otherwise, the Quick Start will display an error during parameter validation.

- Subnets where primary and secondary SAP HANA instances will be deployed must share a common route table. Deploying the Quick Start into a new VPC will ensure this automatically. However, if you're deploying the Quick Start into an existing VPC, you must ensure that your existing subnets have a common route table.
- SLES HAE and RHEL High Availability agents require that the Pacemaker tag and the overlay IP address you provide by setting deployment parameters can be uniquely identified. Therefore, you need to ensure the following:
 - The value you provide for the **PaceMakerTag** parameter isn't being used by any other EC2 instances in your account, in the AWS Region where you are deploying the Quick Start.
 - The IP address you provide for the **VirtualIPAddress** parameter is outside the VPC CIDR and isn't being used in the route table associated with the subnets where primary and secondary HANA instances will be deployed.
- After deployment, we strongly recommend that you validate the setup of your environment before you use the HA cluster for production. For testing scenarios, see section 6 of the SUSE whitepaper, [SAP HANA SR Performance Optimized Scenario](#).

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 quotas, if necessary.

- [Step 2. Subscribe to the operating system image](#) (skip this step if you're planning to use the SLES operating system)

- [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 Single-AZ and Multi-AZ deployment scenarios as well as end-to-end deployment and deployment into an existing 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 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 <https://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 AWS 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.

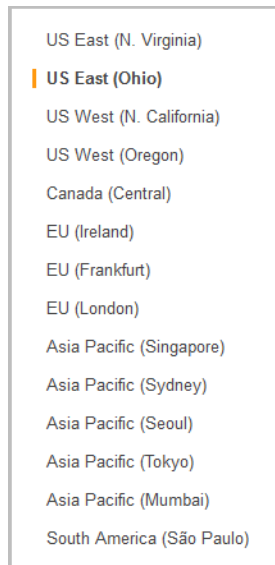


Figure 4: Choosing an AWS Region

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, R4, R5, X1, or Amazon EC2 High Memory instance type for the SAP HANA part of the deployment. 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**.

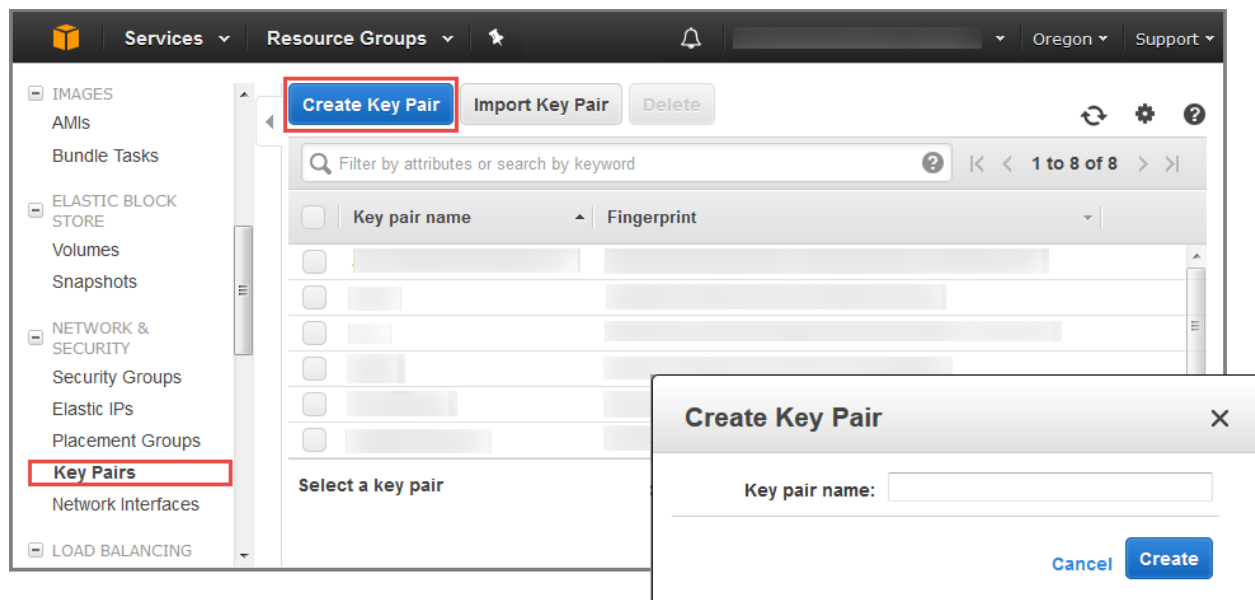


Figure 5: Creating a key pair

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. If you want to enable AWS CloudTrail trails and AWS Config, create an S3 bucket to store CloudTrail trail logs. In the [Amazon S3 console](#), choose **Create Bucket**, provide a name for your new bucket, choose the AWS Region where you want to create the bucket, and then choose **Create**. For detailed information about bucket names and region selection, see the [Amazon S3 documentation](#). You can skip this step if you'd like to use your existing bucket to store the CloudTrail trail logs.
5. (Production deployments only) If necessary, [request a service quota 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 quota with this reference deployment, you will need to request an increase. To do this, on the [Service Quotas](#) console, for each instance type that you want a service quota increase, choose the instance type, choose **Request quota increase**, and then complete the fields in the quota increase form. It might take a few days for the new service quota to become effective. For more information, see the [AWS documentation](#).

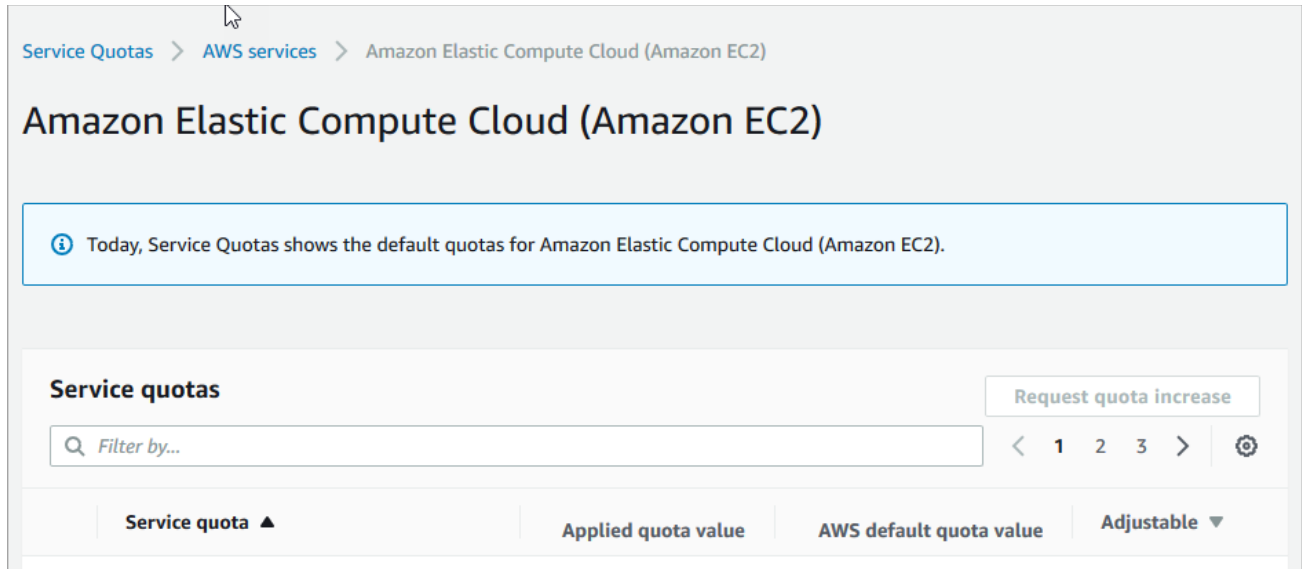


Figure 6: Requesting a service quota increase

Step 2. Subscribe to the Operating System Image

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 <https://aws.amazon.com/marketplace>.
2. Subscribe to the **SLES for SAP**, **RHEL for SAP with HA and US**, or **RHEL for SAP HANA** image.

This involves accepting the terms of the license agreement and receiving confirmation email. For detailed instructions, see [Appendix A](#) for RHEL for SAP, or [Appendix B](#) for SLES for SAP.

Note that the Multi-AZ (HA) deployment scenario requires SLES for SAP or RHEL for SAP with HA and US; see the [Operating System for Deployment](#) section earlier in this guide.

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 <https://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. Under **Installation and Upgrades**, choose **Access Downloads**, 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.**, 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.

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 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 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 follow the wizard to configure permissions and other settings and to create the bucket. For detailed information about bucket names and region selection, see the [Amazon S3 documentation](#).



The screenshot shows the 'Create bucket' wizard in the AWS console. The title bar is blue with a close button (X) on the right. Below the title bar is a progress indicator with four steps: 1. Name and region (highlighted in blue), 2. Configure options, 3. Set permissions, and 4. Review. The main content area is dark blue and contains the following fields:

- Name and region**: A section header.
- Bucket name**: A text input field containing 'myhanamedia'.
- Region**: A dropdown menu showing 'US East (N. Virginia)'.
- Copy settings from an existing bucket**: A dropdown menu with the text 'Select bucket (optional) 61 Buckets'.

At the bottom of the form, there are three buttons: 'Create' (highlighted in blue), 'Cancel', and 'Next'.

Figure 7: Creating an S3 bucket

9. 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.

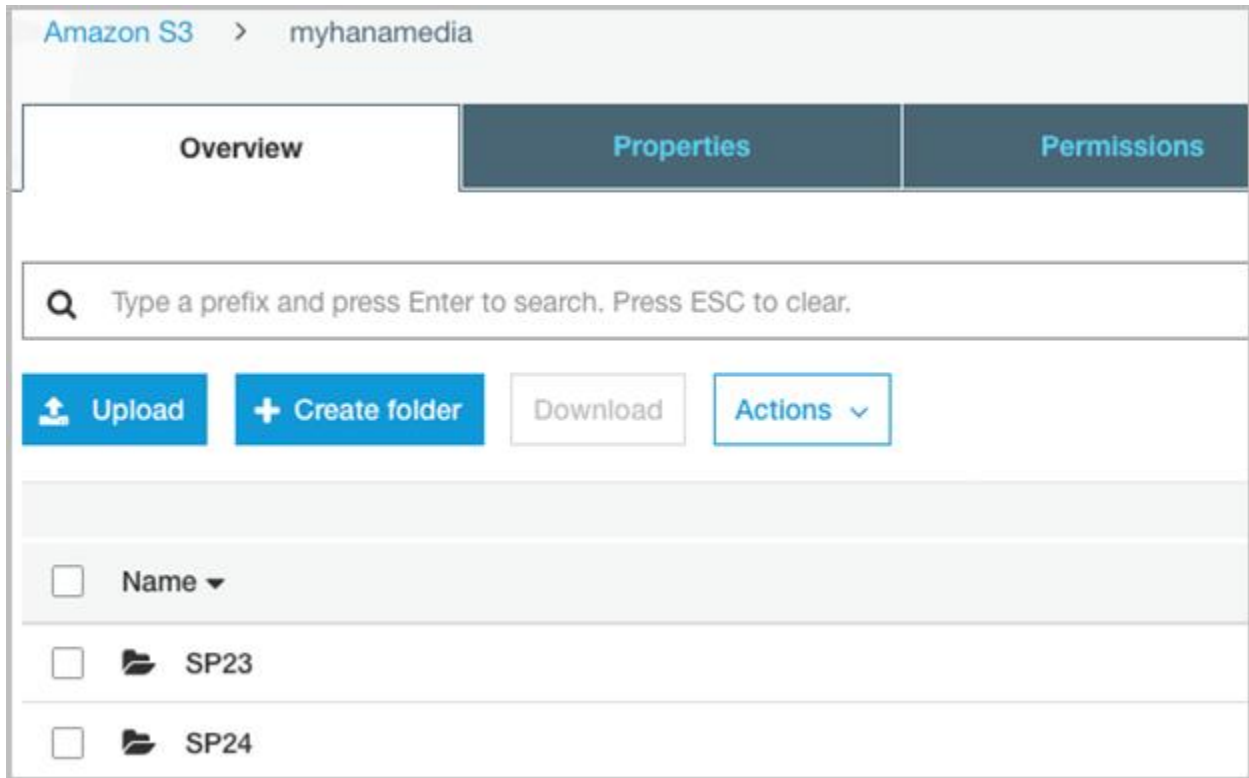


Figure 8: Adding subfolders for SAP HANA versions

10. Choose **Upload** to place the unextracted SAP HANA software in the appropriate folder, as illustrated in Figures 9 and 10, depending on the version of SAP HANA you've downloaded.

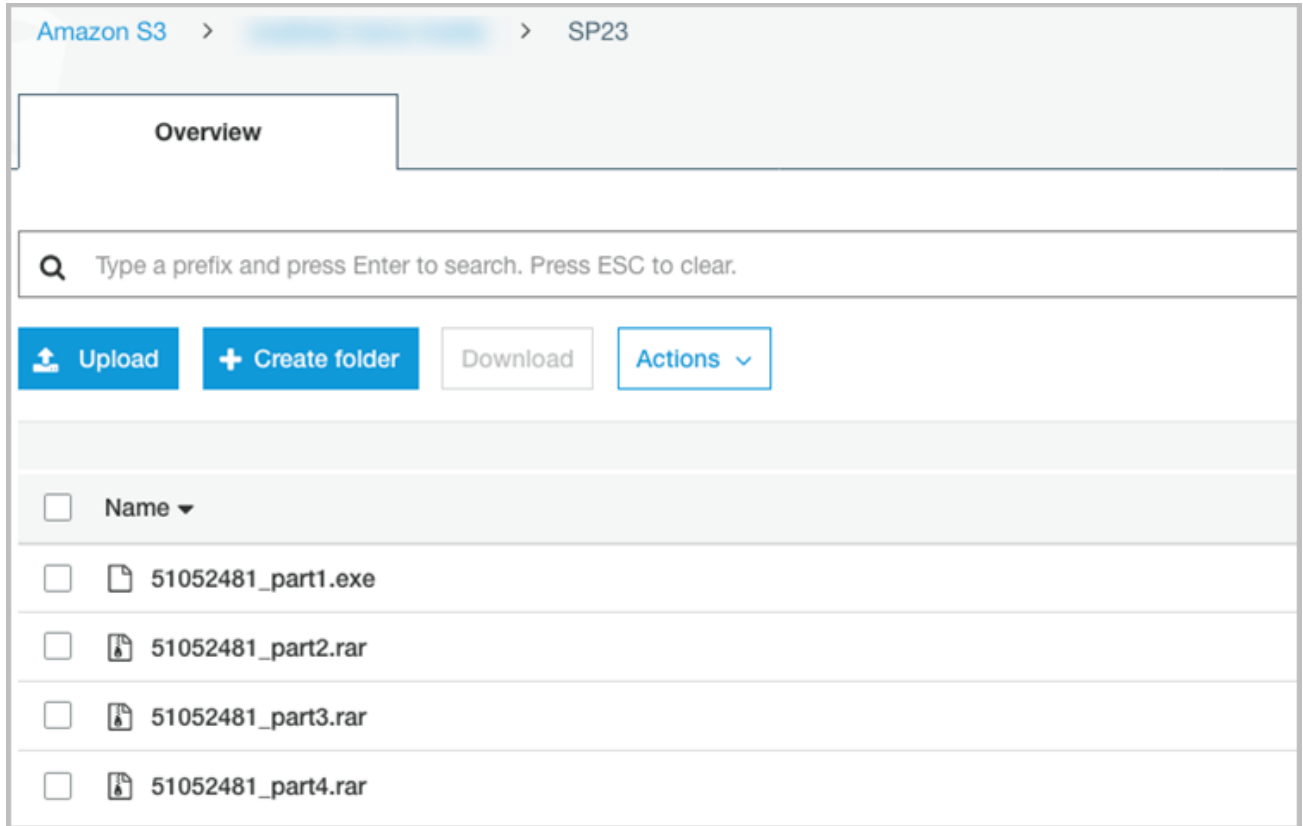


Figure 9: Staging multi-part SAP HANA installation files

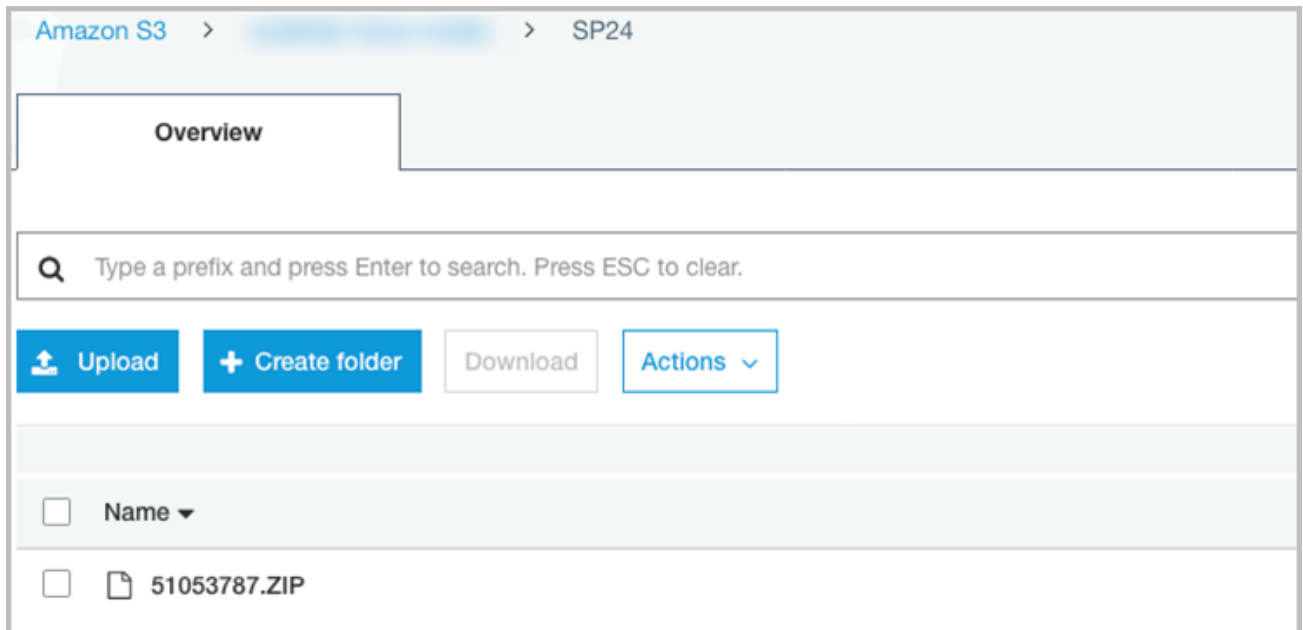


Figure 10: Staging a single SAP HANA installation file

In the examples shown in Figures 9 and 10, the path for the specific version of SAP HANA software will be `s3://myhanamedia/SP23/` or `s3://myhanamedia/SP24/`. Use this path in the **Amazon S3 URL for SAP HANA software** (HANAINstallMedia) parameter field in the next step.

Tip Place only the main SAP HANA installation files in the S3 bucket (as shown in Figures 9 and 10). 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 SLES, SLES for SAP, RHEL for SAP with HA and US, or RHEL for SAP HANA as the operating system. You will have to first complete the subscription process detailed in [step 2](#) if you choose to use RHEL for SAP or SLES for SAP images. If you decide to choose SLES as the operating system, 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.

Single-AZ Single or multiple nodes	Multi-AZ (HA) Single nodes
Launch for new VPC	Launch for new VPC
Launch for existing VPC	Launch for existing VPC

The deployment takes approximately 25 minutes for the Single-AZ, single-node configuration; 35-60 minutes for Single-AZ, multi-node deployment; and around 35 minutes for Multi-AZ (HA), single-node deployment.

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. By default, the template is launched in the US East (N. Virginia) Region.
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.

In the following sections and tables, parameters are listed and described separately for the two deployment scenarios:

- [Single-AZ deployment](#)
- [Multi-AZ deployment](#)

When you finish reviewing and customizing the parameters, choose **Next**.

Single-AZ Deployment

You can choose to deploy the Quick Start into a new or an existing VPC. These two options provide the same parameters except for network infrastructure settings.

[View the template for new VPC](#)

[View the template for existing VPC](#)

- **New VPC – network infrastructure parameters**

Parameter label (name)	Default	Description
CIDR block of Amazon VPC (VPCCIDR)	10.0.0.0/16	CIDR block of the VPC to create for the SAP HANA deployment.
CIDR block of private subnet (PrivSubCIDR)	10.0.1.0/24	CIDR block of the private subnet in the new VPC where SAP HANA will be deployed.

CIDR block of public subnet (DMZCIDR)	10.0.2.0/24	CIDR block of the public subnet in the new VPC where the managed NAT gateway, bastion host, and RDP instance will be deployed.
CIDR block for SSH/RDP access (RemoteAccessCIDR)	0.0.0.0/0	CIDR block from where you are likely to access your bastion and RDP instance.
Instance type for bastion host (BASTIONInstanceType)	t2.small	EC2 instance type for the bastion host.
Availability Zone for subnet creation (AvailabilityZone)	<i>Depends on region selected (e.g., us-east-1a)</i>	Availability Zone where you want to create your subnets.

- **Existing VPC – network infrastructure parameters**

Parameter label (name)	Default	Description
VPC ID (VPCID)	<i>Requires input</i>	ID of the existing VPC where you want to deploy your SAP HANA resources.
Private subnet (HANASubnet)	<i>Requires input</i>	ID of the private subnet in your existing VPC where you want to deploy SAP HANA instances.
Public subnet (DMZSubnet)	<i>Requires input</i>	ID of the public subnet in your existing VPC where you want to deploy the optional RDP instance.
CIDR block of private subnet (PrivSubCIDR)	10.0.1.0/24	CIDR block of the private subnet in the existing VPC where SAP HANA will be deployed. You can retrieve this value from the previous two parameters (private and public subnet IDs), as shown in Figure 11.
CIDR block of public subnet (DMZCIDR)	10.0.2.0/24	CIDR block of the public subnet in the existing VPC where the bastion host and NAT gateway exist. You can retrieve this value from the previous two parameters (private and public subnet IDs), as shown in Figure 11.
CIDR block for RDP access (RemoteAccessCIDR)	0.0.0.0/0	CIDR block from where you want to access the RDP instance.

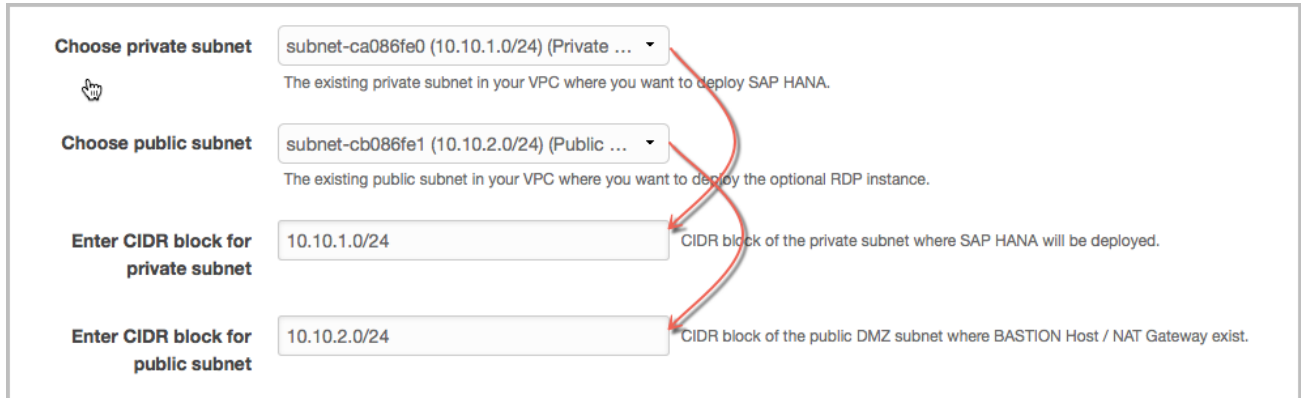


Figure 11: Finding the values for the CIDR block parameters

• **Shared parameters for Single-AZ deployments**

The following parameters apply to both new VPC and existing VPC deployments with Single-AZ configuration.

Server and storage configuration:

Parameter label (name)	Default	Description
Operating system for SAP HANA (MyOS)	SuSELinux12SP4 ForSAP	Operating system and version to be used for SAP HANA master/worker nodes (see Operating System for Deployment section).
SUSE Registration Code (SLESBYOSRegCode)	—	Registration code for the SLES for SAP BYOS image (required only if you're using the BYOS image).
Instance type for SAP HANA (MyInstanceType)	r5.2xlarge	EC2 instance type for SAP HANA nodes (see AWS Instance Type for SAP HANA section).
Number of SAP HANA hosts (HostCount)	1	Total number of nodes you want to deploy in the SAP HANA cluster.
Dedicated Host ID (DedicatedHostId)	—	The IDs of the dedicated hosts on which you want to deploy your SAP HANA instances. Use comma-separated values to specify multiple host IDs for Multi-AZ or scale-out deployments. Make sure that your host reservations are in your target Availability Zone. This parameter is required for Amazon EC2 High Memory instance types (u-6tb1.metal, u-9tb1.metal, u-12tb1.metal, u-18tb1.metal, and u-24tb1.metal).
Automatic recovery (AutoRecovery)	Yes	Set to No to disable the automatic recovery feature on your SAP HANA nodes. This parameter isn't supported for deployments on Dedicated Hosts.

Key pair (KeyName)	<i>Requires input</i>	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 EC2 instances launched by the Quick Start.
Storage volume type for SAP HANA data (VolumeTypeHanaData)	gp2	Amazon EBS storage type to be used for SAP HANA data volumes (see Storage Configuration for SAP HANA section).
Storage volume type for SAP HANA log (VolumeTypeHanaLog)	gp2	Amazon EBS storage type to be used for SAP HANA log volumes (see Storage Configuration for SAP HANA section).
EFS (AWSEFS)	Yes	Set to Yes to use Amazon EFS for shared file systems (/hana/shared, /backup) in a multi-node deployment.
Encryption (Encryption)	No	Set to Yes to enable encryption for all volumes (except root) created for SAP HANA nodes.

SAP HANA database configuration:

Parameter label (name)	Default	Description
Domain name (DomainName)	local	Name to use for fully qualified domain names.
SAP HANA master host name (HANAMasterHostname)	imdbmaster	Host name to use for the SAP HANA master node (DNS short name).
SAP HANA worker host name (HANAWorkerHostname)	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 (SID)	HDB	System ID to be used for HANA installation and setup.
SAP HANA instance number (SAPInstanceNum)	00	Instance number to be used for SAP HANA installation and setup and to open up specific ports for security groups.
SAP HANA password (HANAMasterPass)	<i>Requires input</i>	Password to be used for SAP HANA installation.
SAP HANA Server Timezone (SAPTZ)	UC	Time zone settings to use on the SAP HANA server.
Amazon S3 URL for SAP HANA software (HANAInstallMedia)	<i>Requires input</i>	Full path to Amazon S3 location where you've placed the SAP HANA software in step 3 .

Install SAP HANA software? (InstallHANA)	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 label (name)	Default	Description
Placement group name (PlacementGroupName)	—	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 (InstallRDPInstance)	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 (RDPInstanceType)	c5.large	EC2 instance type for Windows RDP instance.
CIDR block for RDP access (RemoteAccessCIDR)	0.0.0.0/0	CIDR block from where you want to access the RDP instance.
CIDR block of SAP Application Server (ApplicationCIDR)	0.0.0.0/0	CIDR block of the subnet where SAP application servers are deployed.
URL for Proxy Server (Proxy)	—	Address of proxy server for HTTP/HTTPS access if you route your internet traffic via a proxy server (e.g., http://xyz.abc.com:8080 or http://10.x.x.x:8080). This parameter is used only with RHEL for SAP operating system deployments.
AWS CloudTrail Trail & AWS Config (EnableLogging)	No	Set to Yes to enable AWS CloudTrail trails and AWS Config, to help with compliance, audit, and change management requirements.
Amazon S3 bucket name (CloudTrailS3Bucket)	<i>Requires input</i>	Name of the S3 bucket where you want to store AWS CloudTrail trails and AWS Config logs, from step 1 . This parameter is ignored if the EnableLogging parameter is set to No .

Advanced configuration:

Parameter label (name)	Default	Description
Private bucket (PrivateBucket)	—	Name of the main build bucket where installation scripts and templates are located. Don't use this parameter unless directed by AWS Support.

Custom Storage Configuration file location (CustomStorageConfig)	—	Location of the JSON file for customized storage configuration. Don't use this parameter unless directed by AWS Support.
---	---	--

Multi-AZ Deployment

You can choose to deploy the Quick Start into a new or an existing VPC. These two options provide the same parameters except for network infrastructure settings.

[View the template for new VPC](#)

[View the template for existing VPC](#)

• New VPC – network infrastructure parameters

Parameter label (name)	Default	Description
VPC CIDR (VPCCIDR)	10.0.0.0/16	CIDR block of the VPC to create for the SAP HANA deployment.
Availability Zones for subnet creation (AvailabilityZones)	<i>Depends on region selected (e.g., us-east-1a, us-east-1b)</i>	Availability Zones where you want to create your subnets. This field displays the available zones within your selected region. You must select two Availability Zones from this list. If you choose more than two, the Quick Start uses the first two zones from your list.
CIDR block for private subnet 1 (PrivateSubnet1CIDR)	10.0.1.0/24	CIDR block of the private subnet to create in the first Availability Zone. This is where the primary SAP HANA server will be deployed.
CIDR block for private subnet 2 (PrivateSubnet2CIDR)	10.0.2.0/24	CIDR block of the private subnet to create in the second Availability Zone. This is where the secondary SAP HANA server will be deployed.
CIDR block for public subnet 1 (PublicSubnet1CIDR)	10.0.3.0/24	CIDR block of the public subnet to create in the first Availability Zone. This is where the managed NAT gateway, bastion host, and RDP instance will be deployed.
CIDR block for public subnet 2 (PublicSubnet2CIDR)	10.0.4.0/24	CIDR block of the public subnet to create in the second Availability Zone. No resource will be deployed in this subnet.

• Existing VPC – network infrastructure parameters

Parameter label (name)	Default	Description
Choose VPC ID (VPCID)	<i>Requires input</i>	ID of the existing VPC where you want to deploy your SAP HANA resources.
Private subnet ID for HANA Primary Instance (PrivateSubnet1ID)	<i>Requires input</i>	ID of the private subnet in the first Availability Zone in your existing VPC. This is where the primary SAP HANA instance will be deployed.

Private subnet ID for HANA Secondary instance (PrivateSubnet2ID)	<i>Requires input</i>	ID of the private subnet in the second Availability Zone in your existing VPC. This is where the secondary SAP HANA instance will be deployed.
Public subnet ID for RDP instance (PublicSubnet1ID)	<i>Requires input</i>	CIDR block of the public subnet in your existing VPC where the bastion host and NAT gateway exist.
CIDR block for private subnet 1 (PrivateSubnet1CIDR)	<i>Requires input</i>	CIDR block of the private subnet in the first Availability Zone of your existing VPC. This is where the primary SAP HANA server will be deployed. You can retrieve this value from the private and public subnet IDs, as shown previously in Figure 11.
CIDR block for private subnet 2 (PrivateSubnet2CIDR)	<i>Requires input</i>	CIDR block of the private subnet in the second Availability Zone of your existing VPC. This is where the secondary SAP HANA server will be deployed. You can retrieve this value from the private and public subnet IDs, as shown previously in Figure 11.
CIDR block for public subnet (PublicSubnet1CIDR)	<i>Requires input</i>	CIDR block of the public subnet in your existing VPC where the managed NAT gateway, bastion host, and RDP instance exist. You can retrieve this value from the private and public subnet IDs, as shown previously in Figure 10.

• Shared parameters for Multi-AZ deployments

The following parameters apply to both new VPC and existing VPC deployments with a Multi-AZ configuration.

Server and storage configuration:

Parameter label (name)	Default	Description
Operating system for SAP HANA (MyOS)	SuSELinux12SP4 ForSAP	Operating system and version to be used for SAP HANA master/worker nodes (see Operating System for Deployment section).
SUSE BYOS Registration Code (SLESBYOSRegCode)	—	Registration code for the SLES for SAP BYOS image (required only if you're using the BYOS image).
Instance type for SAP HANA (MyInstanceType)	r5.2xlarge	EC2 instance type for SAP HANA nodes (see AWS Instance Type for SAP HANA section).
Dedicated Host ID (DedicatedHostId)	—	The IDs of the dedicated hosts on which you want to deploy your SAP HANA instances. Use comma-separated values to specify multiple host IDs for multi-AZ deployments. Make

		sure that your host reservations are in your target Availability Zones. This parameter is required for Amazon EC2 High Memory instance types (u-6tb1.metal, u-9tb1.metal, u-12tb1.metal, u-18tb1.metal, and u-24tb1.metal).
Key pair (KeyName)	<i>Requires input</i>	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 EC2 instances launched by the Quick Start.
Storage volume type for SAP HANA data (VolumeTypeHanaData)	gp2	Amazon EBS storage type to be used for SAP HANA data volumes (see Storage Configuration for SAP HANA section).
Storage volume type for SAP HANA log (VolumeTypeHanaLog)	gp2	Amazon EBS storage type to be used for SAP HANA log volumes (see Storage Configuration for SAP HANA section).
Turn on encryption? (Encryption)	No	Set to Yes to enable encryption for all volumes (except root) created for SAP HANA nodes.

SAP HANA database configuration:

Parameter label (name)	Default	Description
Domain name (DomainName)	local	Name to use for fully qualified domain names.
SAP HANA primary host name (HANAHostnamePrimary)	prihana	Host name to use for the SAP HANA primary instance (DNS short name).
SAP HANA secondary host name (HANAHostnameSecondary)	sechana	Host name to use for the SAP HANA secondary instance (DNS short name).
SAP HANA system ID (SID)	HDB	System ID to be used for HANA installation and setup.
SAP HANA instance number (SAPInstanceNum)	00	Instance number to be used for SAP HANA installation and setup and to open up specific ports for security groups.
SAP HANA password (HANAMasterPass)	<i>Requires input</i>	Password to be used for SAP HANA installation. This password must follow SAP HANA password restrictions and will be used for both SAP HANA instances.
SAP HANA Server Timezone (SAPTZ)	UC	Time zone settings to use on the SAP HANA server.
Amazon S3 URL for SAP HANA software (HANAInstallMedia)	<i>Requires input</i>	Full path to Amazon S3 location where you've placed the SAP HANA software in step 3 .

Parameter label (name)	Default	Description
Install SAP HANA software? (InstallHANA)	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.

SAP HANA high availability configuration:

Parameter label (name)	Default	Description
HANA primary site name (PrimaryHANASite)	HAP	SAP HANA primary site name for SAP HANA System Replication.
HANA secondary site name (SecondaryHANASite)	HAS	SAP HANA secondary site name for SAP HANA System Replication.
Overlay IP address (VirtualIPAddress)	192.168.1.99	Overlay IP address to be assigned to the active node. This IP address should be outside the VPC CIDR and must not be used by any other HA cluster.
Pacemaker Tag Name (PaceMakerTag)	PaceTag	Tag to be assigned to each EC2 instance. This tag is used by the Pacemaker component of SLES HAE and RHEL for SAP High Availability solutions. It must be unique for your account and should not be used with any other EC2 instance (cluster or single node).

Optional configuration:

Parameter label (name)	Default	Description
Windows RDP instance (InstallRDPInstance)	No	Set to Yes to enable the deployment of an EC2 instance with Windows Server to host SAP HANA Studio.
Instance type for RDP host (RDPIInstanceType)	c4.large	EC2 instance type for Windows RDP instance.
CIDR block for RDP access (RemoteAccessCIDR)	0.0.0.0/0	CIDR block from where you want to access the RDP instance.
AWS CloudTrail & AWS Config (EnableLogging)	No	Set to Yes to enable AWS CloudTrail trails and AWS Config to help with compliance, audit, and change management requirements.
S3 bucket name for AWS CloudTrail and AWS Config (CloudTrailS3Bucket)	<i>Requires input</i>	Name of the S3 bucket where you want to store AWS CloudTrail trails and AWS Config logs, from step 1 . This parameter is ignored if the EnableLogging parameter is set to No .

Parameter label (name)	Default	Description
CIDR for application instances (ApplicationCIDR)	0.0.0.0/0	CIDR block of the subnet where SAP application servers are deployed.
Proxy server address (Proxy)	—	Address of proxy server for HTTP/HTTPS access if you route your internet traffic via a proxy server (e.g., http://xyz.abc.com:8080 or http://10.x.x.x:8080). This parameter is used only with RHEL for SAP operating system deployments.

Advanced configuration:

Parameter label (name)	Default	Description
S3 bucket containing SAP HANA assets (PrivateBucket)	aws-quickstart/quickstart-sap-hana	Name of the main build bucket where installation scripts and templates are located. Don't use this parameter unless directed by AWS Support.
Custom storage configuration location (CustomStorageConfig)	aws-quickstart/quickstart-sap-hana/scripts	Location of the JSON file for customized storage configuration. Don't use this parameter unless directed by AWS Support.
S3 bucket containing HANA HA Assets (QSS3BucketName)	aws-quickstart	Name of the main build bucket where the SAP HANA HA Lambda functions and templates are located. Don't change this parameter unless directed by AWS Support.
HANA HA bucket key prefix (QSS3KeyPrefix)	quickstart-sap-hana-ha	Key prefix to use for the HANA HA assets bucket. Don't change this parameter unless directed by AWS Support.

- On the **Options** page, you can [specify tags](#) (key-value pairs) for resources in your stack and [set additional options](#).
- (Optional) On the **Options** page, under **Advanced**, enable **Termination Protection**. We strongly recommend that you set this option if you are creating the stack for production or other critical workloads. When you're done, choose **Next**.
- On the **Review** page, review and confirm the template settings. Under **Capabilities**, select the checkbox to acknowledge that the template will create IAM resources.
- Choose **Create** to deploy the stack.
- Monitor the status of the stack. When the status is **CREATE_COMPLETE**, you can continue to the next step to verify your deployment.

Stack Name	Created Time	Status	Description
<input checked="" type="checkbox"/> HANA-QuickStart-BaseNodes-1MHTW11Z87QLP	2016-05-11 08:27:21 UTC-0700	CREATE_COMPLETE	(0008) Deploy SAP HANA on AWS
<input type="checkbox"/> HANA-QuickStart-BaseNetwork-HGK4V08NLRB	2016-05-11 08:24:29 UTC-0700	CREATE_COMPLETE	(0007) AWS Infrastructure Deployment for SAP HANA
<input type="checkbox"/> HANA-QuickStart	2016-05-11 08:24:01 UTC-0700	CREATE_COMPLETE	(0007) Deploy AWS infrastructure and SAP HANA on AWS

Key	Value	Description
HANAMasterInstanceIP	10.25.1.117	HANA Master Node IP Address
HANAMasterSecurityGroup	sg-f4459a93	Security Group created for the SAP HANA Master node

Figure 12: Successful completion of template deployment (Single-AZ scenarios)

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 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 13: 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.

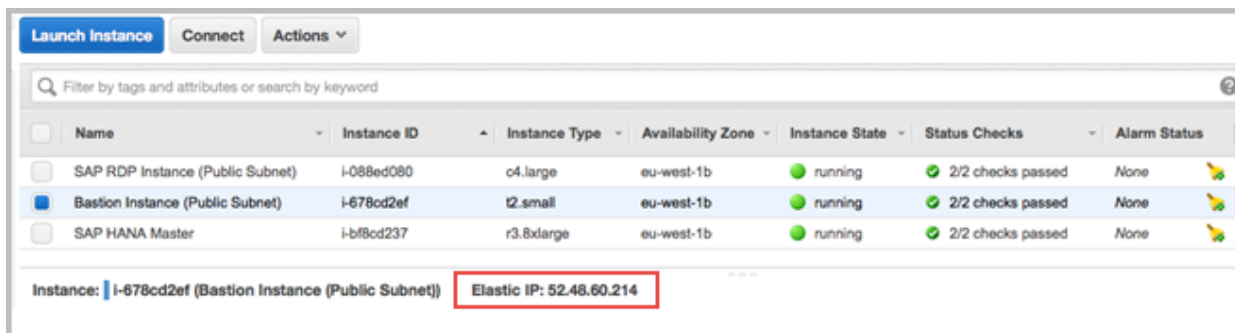


Figure 14: Elastic IP address for bastion host

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.

```
ssh-add mykey.pem
Identity added: mykey.pem (mykey.pem)
ssh -A 52.48.60.214 -l ec2-user
Last login: Tue May 3 23:00:04 2016 from 54-240-197-225.amazon.com

  _| _|_ )
  _| (  /  Amazon Linux AMI
  _|\_|_|

https://aws.amazon.com/amazon-linux-ami/2016.03-release-notes/
7 package(s) needed for security, out of 23 available
Run "sudo yum update" to apply all updates.
[ec2-user@ip-10-30-2-76 ~]$ ssh 10.30.1.251
Last login: Tue May 3 23:00:29 2016 from ip-10-30-2-76.eu-west-1.compute.internal
SUSE Linux Enterprise Server 12 SP1 x86_64 (64-bit)

As "root" (sudo or sudo -i) use the:
- zypper command for package management
- yast command for configuration management

Management and Config: https://www.suse.com/suse-in-the-cloud-basics
Documentation: http://www.suse.com/documentation/sles12/
Forum: https://forums.suse.com/forumdisplay.php?93-SUSE-Public-Cloud

Have a lot of fun...
ec2-user@imdbmaster:~> sudo su -
imdbmaster:~ #
```

Figure 15: 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**.

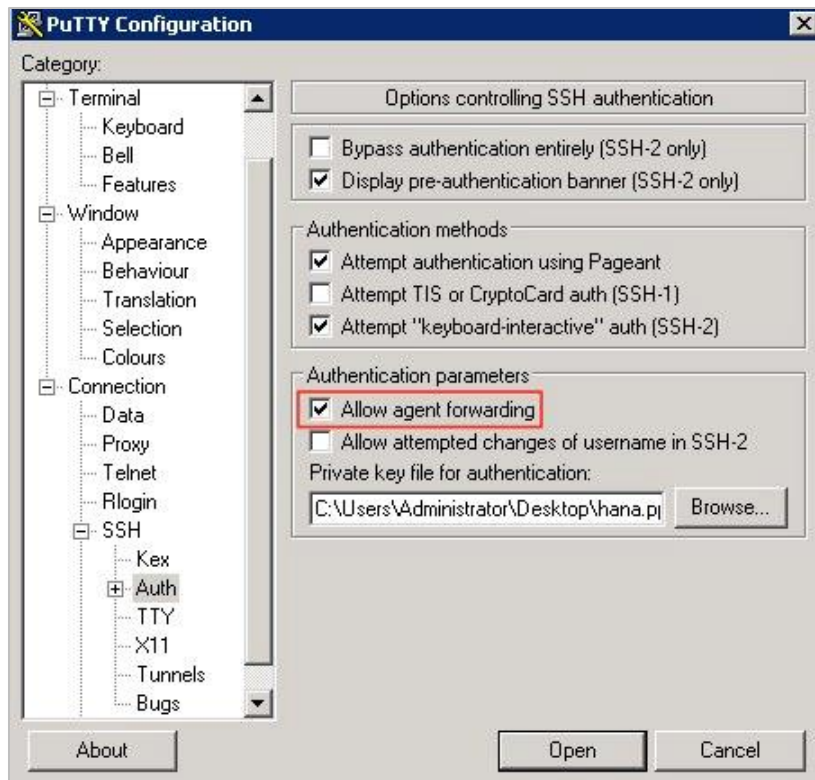


Figure 16: PuTTY example for SSH connection

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.

- Make sure that you are running on the latest supported kernel. For details, see SAP OSS Note [2205917](#) or [2240716](#) for SLES, and [2247020](#) or [2292690](#) for RHEL.
- If you are using RHEL for SAP or SLES for SAP as your operating system, C-states and P-states are automatically set by using **saptune** (for SLES) and **tuned-adm** (for RHEL) on all [supported EC2 instances](#). You need to reboot your instance after deployment to enable this. If you are using SLES as your operating system, configure C-states and P-states for lower latency by following the instructions provided in the SAP OSS Note [2205917](#) for SLES 12 or above.

- Update your SAP HANA software with the latest patches.
- Install any additional components such as Application Function Libraries (AFL) or Smart Data Access (SDA).
- If you have deployed an SAP HANA scale-out cluster, consider adding additional elastic network interfaces (ENIs) and security groups to logically separate network traffic for client, inter-node, and optional SAP HANA System Replication (HSR) communications. For details, see the [SAP HANA Operations Guide](#).
- 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.
- If you have chosen to enable AWS CloudTrail trails and AWS Config, review the list of resources being monitored by the AWS Config service from the [AWS Config console](#). By default, this Quick Start enables monitoring of configuration changes for the following resources: S3 buckets, CloudTrail trails, EC2 instances, EC2 volumes, security groups, subnets, VPCs, route tables, network access control lists (ACLs), and network interfaces. If you'd like to add or remove resources, you can adjust this list as needed in the AWS Config console.
- For Multi-AZ, high-availability deployments, we strongly recommend that you validate the setup of your environment before you use the HA cluster for production. After the successful completion of the Quick Start and before you deploy an application on SAP HANA, run a round of tests to ensure that failover and failback are working as desired. For testing scenarios, see section 6 of the SUSE whitepaper, [SAP HANA SR Performance Optimized Scenario](#).

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 that generates a new VPC for a Single-AZ, single-node or multi-node setup, and I see two additional templates being launched in the AWS CloudFormation console. Why?

When you launch the SAP HANA Quick Start for a Single-AZ configuration and a new VPC, it launches two templates: one template to set up your network infrastructure (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.

I launched the SAP HANA Quick Start template for a Multi-AZ, single-node, high-availability deployment, and I see multiple additional templates being launched in the AWS CloudFormation console. Why?

The Multi-AZ, single-node, high availability template has been modularized for logical separation of tasks. For example, it includes nested templates for the following tasks, and you will see additional outputs for these:

- VPC creation
- Parameter validation
- Deployment of SAP HANA instances
- Pre-HA configuration
- HA configuration

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

If you're deploying SAP HANA Platform Edition 2 SPS 1 or a later version, the Quick Start uses MDC by default. If you're deploying an earlier version of SAP HANA, the Quick Start installs SAP HANA software with the single-tenant option. If you need MDC configured for an earlier version, we recommend that you skip the SAP HANA software installation during

deployment (set the **InstallHANA** 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
In order to use this AWS Marketplace product you need to accept terms and subscribe. To do so please visit http://aws.amazon.com/marketplace/	You are trying to launch your stack with a RHEL or SLES AMI from AWS Marketplace without a subscription.	SLES for SAP, RHEL for SAP with HA and US, and RHEL for SAP HANA require a subscription. Follow the instructions in step 2 and launch your stack again.
API: ec2: RunInstances Not authorized for images: <i>ami-ID</i>	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 requires SLES 12 and above, or RHEL 7.x and above for SAP workloads.
Signal-failure function not implemented.	Deployment failed for an unknown reason.	Contact the support team and open a support case.
Your requested instance type (<i>type</i>) is not supported in your requested	The instance type you selected is not available in	If you are trying to deploy SAP HANA into an existing VPC, create a new private

Error message	Possible cause	What to do
Availability Zone (zone). Please retry your request by not specifying an Availability Zone or choosing (zone).	the chosen Availability Zone. (Your subnet is associated with a single Availability Zone.)	<p>subnet in the suggested Availability Zone and associate a route table with it. Retry the deployment by choosing the newly created private subnet.</p> <p>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.</p>
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 HANA instance is not routed through a NAT instance or NAT gateway.	<p>See if it is possible to temporarily route the internet traffic by using a NAT instance or NAT gateway.</p> <p>If your internet traffic has to go through your internal proxy, contact your network team for access to the SUSE or RHEL update repository. For further assistance, open a support case in the AWS Support Center.</p>
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 quota increase by completing the request form in the Service Quotas console.
WaitCondition timed out. Received 0 conditions when expecting 1.	No or slow internet access is causing delays when downloading SAP HANA software.	<p>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.</p> <p>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.</p>

Error message	Possible cause	What to do
The HANA installation did not succeed. Please check installation media.	SAP HANA installation failed or SAP HANA services didn't start up successfully.	<p>Verify that you have staged the SAP HANA software properly in the S3 bucket with correct permissions. (See step 3 for details.)</p> <p>Another reason could be that SAP HANA services did not start up after the installation.</p> <p>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.</p>
System Administrator password must contain at least 8 characters.	The SAP HANA master password contains \$ or other special characters.	<p>Change the SAP HANA master password and then relaunch the Quick Start.</p> <p>The password must be at least 8 characters, consisting of uppercase and lowercase letters and numbers. Avoid using special characters such as @ or \$.</p>
No such s3 bucket with name <i>name</i>	The S3 bucket name you have provided to store AWS CloudTrail trails and AWS Config log files doesn't exist.	Create a new S3 bucket where you want to store your AWS CloudTrail trails and AWS Config log files, or specify a bucket that already exists. (See step 1 for details.)
Failed to put delivery channel <i>stackname-DeliveryChannel-ID</i> because the maximum number of delivery channels: 1 is reached.	You have already enabled AWS Config in your target region.	Launch your stack again, and make sure that the EnableLogging parameter is set to No .
SUSE BYOS did not succeed. Check SUSE registration code or internet connection	You have used an invalid SUSE registration code, or your instance could not connect to the SUSE registration server.	Verify your SUSE registration code and internet connectivity. Launch your stack again with a valid SUSE registration code for SLES for SAP images.
storage.json file not found or found empty. If custom storage.json is used, check file permission.	You are using a custom storage configuration file, but it is either inaccessible or empty.	Verify that your JSON file for custom storage configuration has the appropriate permission settings, and that the location you specified with the CustomStorageConfig (Custom Storage Configuration file location) parameter is correct.
Tag not unique	You have provided a value for the PaceMakerTag parameter that is already in usage.	The Pacemaker tag name should be unique and not assigned to any other EC2 instances in your account, in the AWS Region where you are deploying SAP

Error message	Possible cause	What to do
		HANA. Choose a different value for this parameter and try again.
Virtual IP address should not be in VPC CIDR	You have selected a virtual IP address that is within the VPC CIDR.	The virtual IP address you specify must be outside the VPC CIDR. For example, if the VPC CIDR is 10.1.0.0/16, you cannot provide a virtual IP address that starts with 10.1.
Primary and Secondary Subnet must have same route table	The route tables for the primary and secondary subnets are different.	The subnets where the primary and secondary SAP HANA hosts will be deployed must share the same route table. Adjust the route tables, after appropriate impact assessment, and try again.
Virtual IP is already being used (in Route Table of Subnet)	There is already an existing route to the virtual IP address provided.	The route table for the subnets where the primary and secondary SAP HANA instances will be deployed already has an entry for the virtual IP address you specified. The virtual IP address should be unique, so select a new IP address and try again.
The EC2 'Recover' Action is not valid for the associated instance	You are deploying your instance in a Dedicated Host, but the automatic recovery option is enabled.	Automatic recovery isn't supported on Dedicated Hosts. Redeploy the Quick Start and set the Automatic Recovery parameter to No .
Invalid value 'aes' for zone. Invalid format (Service: AmazonEC2; Status Code: 400; Error Code: InvalidParameterValue)	You selected u-6tb1.metal, u-9tb1.metal, u-12tb1.metal, u-18tb1.metal, or u-24tb1.metal as your instance type for SAP HANA, but you didn't enter a Dedicated Host ID.	Amazon EC2 High Memory instance types (6 TB, 9 TB, 12 TB, 18 TB, and 24 TB) require a Dedicated Host. Try your deployment again, and enter a Dedicated Host ID.

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.

Operating System 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
<https://aws.amazon.com/documentation/cloudformation/>
- Amazon EBS
 - User guide
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html>
 - Volume types
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html>
 - Optimized instances
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSOptimized.html>
- Amazon EC2
 - User guide for Microsoft Windows
<https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/>
 - User guide for Linux
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/>

- X1 instances

<https://aws.amazon.com/ec2/instance-types/x1/>

- Amazon VPC

<https://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

https://do.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

<https://do.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

<https://do.awsstatic.com/enterprise-marketing/SAP/SAP-HANA-on-AWS-Manual-Setup-Guide.pdf>

- Migrating SAP HANA Systems to X1 Instances on AWS

<https://do.awsstatic.com/enterprise-marketing/SAP/migrating-sap-hana-to-x1-on-aws.pdf>

- Additional information about SAP solutions on AWS

<https://aws.amazon.com/sap/whitepapers/>

Quick Start reference deployments

- SAP Business One, version for SAP HANA

<https://docs.aws.amazon.com/quickstart/latest/sap-b1-hana/>

- SAP NetWeaver for AS ABAP and SAP HANA

<https://aws-quickstart.s3.amazonaws.com/quickstart-sap-netweaver-abap/doc/sap-netweaver-abap-on-the-aws-cloud.pdf>

- Additional reference deployments

<https://aws.amazon.com/quickstart/>

Appendix A: Subscribing to RHEL for SAP

1. Log in to AWS Marketplace at <https://aws.amazon.com/marketplace>.
2. Search for **RHEL for SAP HANA**, and then choose **Red Hat Enterprise Linux for SAP with HA and US** or **Red Hat Enterprise Linux for SAP HANA** from the results.

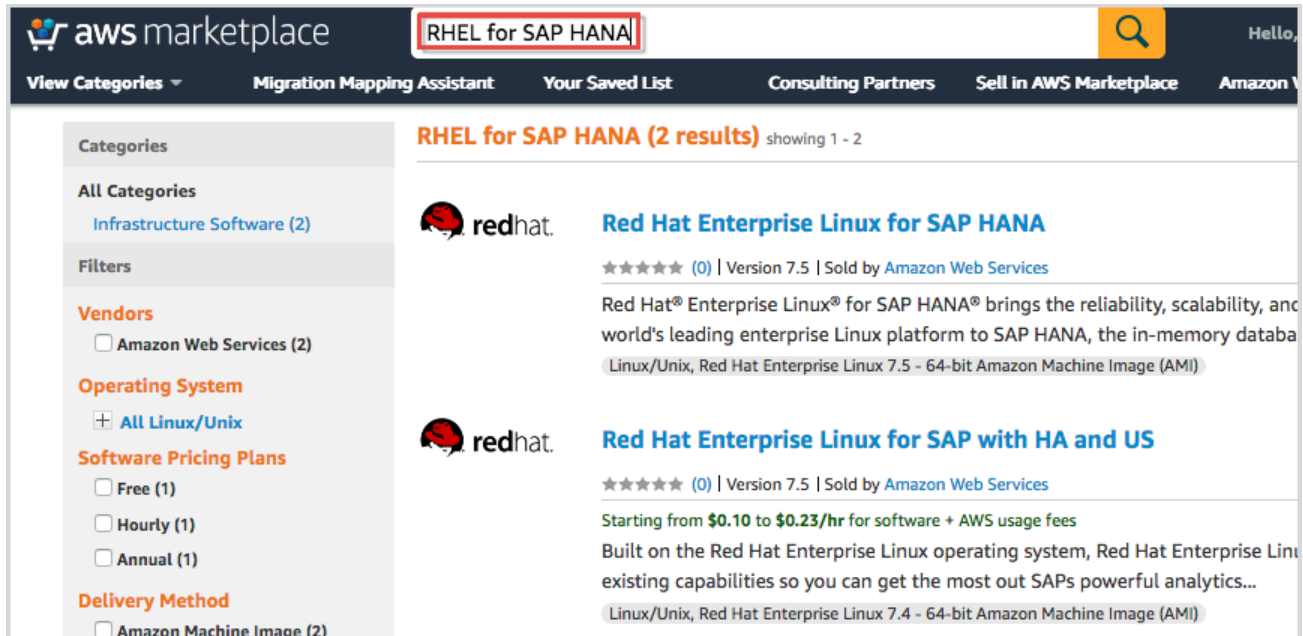


Figure 17: RHEL for SAP in AWS Marketplace

3. Choose **Continue to Subscribe**.
4. 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.

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

Appendix B: Subscribing to SLES for SAP

1. Log in to AWS Marketplace at <https://aws.amazon.com/marketplace>.
2. Type **SLES for SAP** in the search box. From the results, choose the version of SUSE Linux Enterprise Server for SAP that you want to use.



Figure 19: SLES for SAP in AWS Marketplace

3. Choose **Continue to Subscribe**.
4. Read the terms and conditions of software usage, carefully review the additional software pricing information, and then choose **Accept Terms**.




Figure 20: 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.

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

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.
0.0.0.0/0	TCP	80 (HTTP)	Allow outbound HTTP access from instances deployed in the VPC to anywhere.
0.0.0.0/0	TCP	443 (HTTPS)	Allow outbound HTTPS access from instances deployed in the 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##13	Database client access to system database.
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	3##41 – 3##44	Database client access to tenant database.
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 send your feedback to quickstart@amazon.com.

Document Revisions

Date	Change	Location
May 2020	Updates to maximum throughput and IOPS for u-6tb1.metal, u-9tb1.metal, and u-12tb1.metal.	Storage Configuration
March 2020	Support for u-6tb1.metal instances.	Changes throughout templates and guide
October 2019	Support for 18 TB and 24 TB Amazon EC2 High Memory instance sizes.	Changes throughout templates and guide
August 2019	Removed deprecated OS version SLES 12 SP1 for SAP and RHEL 7.3 for SAP.	Operating System for Deployment
July 2019	Multi-AZ support for SLES 12 SP4 for SAP, SLES 15 for SAP, RHEL 7.4 for SAP with HA and US, RHEL 7.5 for SAP with HA and US, and RHEL 7.6 for SAP with HA and US.	Operating System for Deployment
June 2019	Support for SLES 12 SP4, SLES 15, RHEL 7.6 and additional instance types.	Changes throughout templates and guide
December 2018	Support for R5 instances.	Changes throughout templates and guide
November 2018	Support for RHEL for SAP with HA and US, and newer version of RHEL for SAP HANA.	Changes throughout templates and guide
September 2018	Support for Amazon Ec2 High Memory instance types.	Changes throughout templates and guide
August 2018	Support for x1e.2xlarge and x1e.xlarge instance types for cost-optimized, non-production deployments. Support for dedicated host deployments.	Changes throughout templates and guide
June 2018	Support for Multi-AZ, high-availability configurations.	Changes throughout templates and guide. Specifically, see sections for Multi-AZ configuration in Architecture and parameter tables .
April 2018	Support for SLES 12 SP3, separate volumes for SAP HANA data and logs, additional outputs, SSM agent installation for SLES to support future enhancements, and other minor enhancements and bug fixes.	Changes throughout templates and guide

Date	Change	Location
January 2018	Support for RHEL 7.3 and SLES for SAP BYOS, time zone configuration, and custom storage configuration	Changes throughout templates and guide
November 2017	Support for SLES 12 SP2 for SAP OS, x1e.4xlarge instance, new regions, and minor bug fixes. Support for AWS CloudTrail trails and AWS Config.	Changes throughout templates and guide
September 2017	Support for new x1e.32xlarge instance type, SLES for SAP images, and minor bug fixes.	Changes throughout templates and guide
July 2017	Support for routing internet traffic via proxy for RHEL-based deployments, adjusted security groups to access multi-tenant database, and other minor bug fixes.	Changes throughout templates and guide
December 2016	Support for R4 instance type, HANA Platform Edition 2, and RHEL 7.2.	Changes throughout templates and guide
November 2016	Support for x1.16xlarge, RHEL 6.7 support with X1, regional expansions, C-state/P-state instructions and other minor enhancements.	Changes throughout templates and guide
August 2016	Support for X1 scale-out deployment.	Changes throughout templates and guide
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 multiple enhancements.	Changes throughout templates and 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 templates and guide
July 2014	Initial publication	—

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

The software included with this guide is licensed under the Apache License, Version 2.0 (the "License"). You may not use this file except in compliance with the License. A copy of the License is located at <http://aws.amazon.com/apache2.0/> or in the "license" file accompanying this file. This code is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.