SAP HANA on AWS
Implementation and Operations Guide

Created by: Amazon Web Services, Inc.
sap-on-aws@amazon.com

Version: 1.0 – February 2014

Contents

About this Guide

This guide provides best practice guidelines for implementing and operating the SAP HANA Infrastructure Subscription offering on Amazon Web Services (AWS). The intended audience of this guide is SAP customers and partners. This guide is not intended to replace any of the standard SAP HANA documentation. SAP Administration installation guides and notes can be found at:

- SAP Library (help.sap.com) - SAP HANA Administration Guide
- SAP Installation Guides
- SAP Notes

This guide assumes that you have a basic knowledge of Amazon Web Services. If you are new to AWS please read the following guides before continuing with this guide.

- Getting Started with AWS
- What is Amazon EC2?
- SAP on AWS Implementation Guide

Additional SAP on AWS Documentation

In addition to this guide the following SAP on AWS guides can be found at http://aws.amazon.com/sap > Resources

SAP on AWS Operations Guide

The *SAP on AWS Operations Guide* provides guidelines on the special considerations that must be taken into account when operating SAP environments on AWS.

SAP on AWS High Availability Guide

The *SAP on AWS High Availability Guide* provides guidelines on how to configure SAP systems on Amazon EC2 in such a way as to be able to protect the application from various single points of failure.

SAP on AWS Backup and Recovery Guide

The *SAP on AWS Backup and Recovery Guide* provides guidelines on how to backup SAP systems running on AWS. The guide focuses on the essential differences in backing up SAP systems on AWS compared to traditional infrastructure.

Overview of SAP HANA on AWS

AWS and SAP have worked together closely over the past couple of years to make SAP HANA available on the flexible AWS platform. Today there are multiple SAP HANA offerings available on AWS. An overview of the different offerings is provided in the following section.

SAP HANA Developer Edition

Description	Fully featured SAP HANA virtual appliance on AWS for individual developers
Use Cases	• Non-production only • Develop, test and demo applications • Learning environment
HANA Licensing	Free license provided by SAP for developers
Available from	SAP SCN
EC2 instance types	m2.xlarge / m2.2xlarge / m2.4xlarge
Number of nodes	1
HANA memory	17.1 GiB / 34.2 GiB / 68.4 GiB

SAP HANA One

Description	Fully featured SAP HANA virtual appliance on AWS
Use Cases	• Production and non-production • Analytics acceleration • Data merging • Temporary event based analytics • Self-service BI • Prototypes and proofs-of-concept • No connection to SAP-licensed products other than Lumira permitted
HANA Licensing	$0.99 p/hour on-demand license from SAP via the AWS Marketplace
Available from	AWS Marketplace
EC2 instance types	cc2.8xlarge
Number of nodes	1
HANA Memory	60.5 GiB

SAP HANA Infrastructure Subscription

Description	Fully featured SAP HANA virtual appliance on AWS
Use Cases	• Production, non-production, POC's, DR • All SAP HANA use cases supported for non-production scenarios on single node and multi-node HANA virtual appliances. • SAP BW supported in production on single node HANA virtual appliance • Multi-node for BW and Business Suite use cases coming soon.
HANA Licensing	Bring-your-own-License. Customers must have a current license for the SAP HANA Database.
Available from	SAP HANA Marketplace
EC2 instance types	cr1.8xlarge
Number of nodes	1-5
HANA Memory	244 GiB / 488 GiB / 732 GiB / 976 GiB / 1.22TiB

Sizing

SAP HANA is offered on AWS in both single node and multi-node configurations with a total of 244, 488, 732, 976, and 1220 GiB RAM. Since HANA is a columnar database it requires less storage to store data compared to a traditional row based RDMS. Data is highly compressed and compression ratios can range from 3:1 to over 10:1 based on the source data and source database.

As far as sizing of the HANA appliance is concerned, main memory is the most important resource. There are various sizing methods depending on the implementation scenario but in general the following methods apply:

To obtain sizing information for a system that has not yet been implemented, use the SAP QuickSizer. Please go to http://service.sap.com/quicksizer for further details. The SAP QuickSizer will provide information on both the SAP HANA In-Memory Database and the SAP NetWeaver application server where applicable.

To migrate an existing SAP NetWeaver BW system from any database platform to HANA, SAP strongly recommends to use the new ABAP sizing report for SAP NetWeaver BW described in SAP note 1736976.

To migrate an already existing Business Suite System to HANA, it's recommended to use SAP note 1872170 to estimate the main memory requirements of the HANA virtual appliance.

> **Note:** Further sizing information is also available in the SAP HANA Administration Guide.

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

Table 1: Common SAP HANA Sizing Notes

If memory requirements for the SAP HANA solution exceed the available memory of a single AWS instance, a scale out solution consisting of multiple instances can be deployed as long as the SAP solution being deployed supports a scale-out configuration.

Solution Architecture

The SAP HANA on AWS Infrastructure Subscription can be deployed in either a single node or multi-node architecture configuration consisting of up to 5 HANA nodes.

AWS Architecture Components

Single and Multi-node deployments automatically provision and deploy and stitch together all the necessary AWS components into a customer's AWS account using AWS CloudFormation. 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.

The following components are deployed and configured as part of this offering:

- An AWS Virtual Private Cloud (VPC) configured with two subnets, one public, and the other private.
- A NAT instance deployed into the public subnet and configured with an Elastic IP Address (EIP) for outbound Internet connectivity and inbound SSH access.
- A Windows Server deployed in the public subnet with HANA Studio preloaded.
- An Identity and Access Management (IAM) instance role with fine-grained permissions for backup and failure recovery.
- An S3 Bucket where HANA Backups can be stored.
- An SAP HANA System installed with the proper EBS storage volumes configuration for HANA performance needs.
- Pre-configured security groups.
- Single node or multi-node SAP HANA virtual appliances automatically configured per SAP best practices.

Single Node Architecture

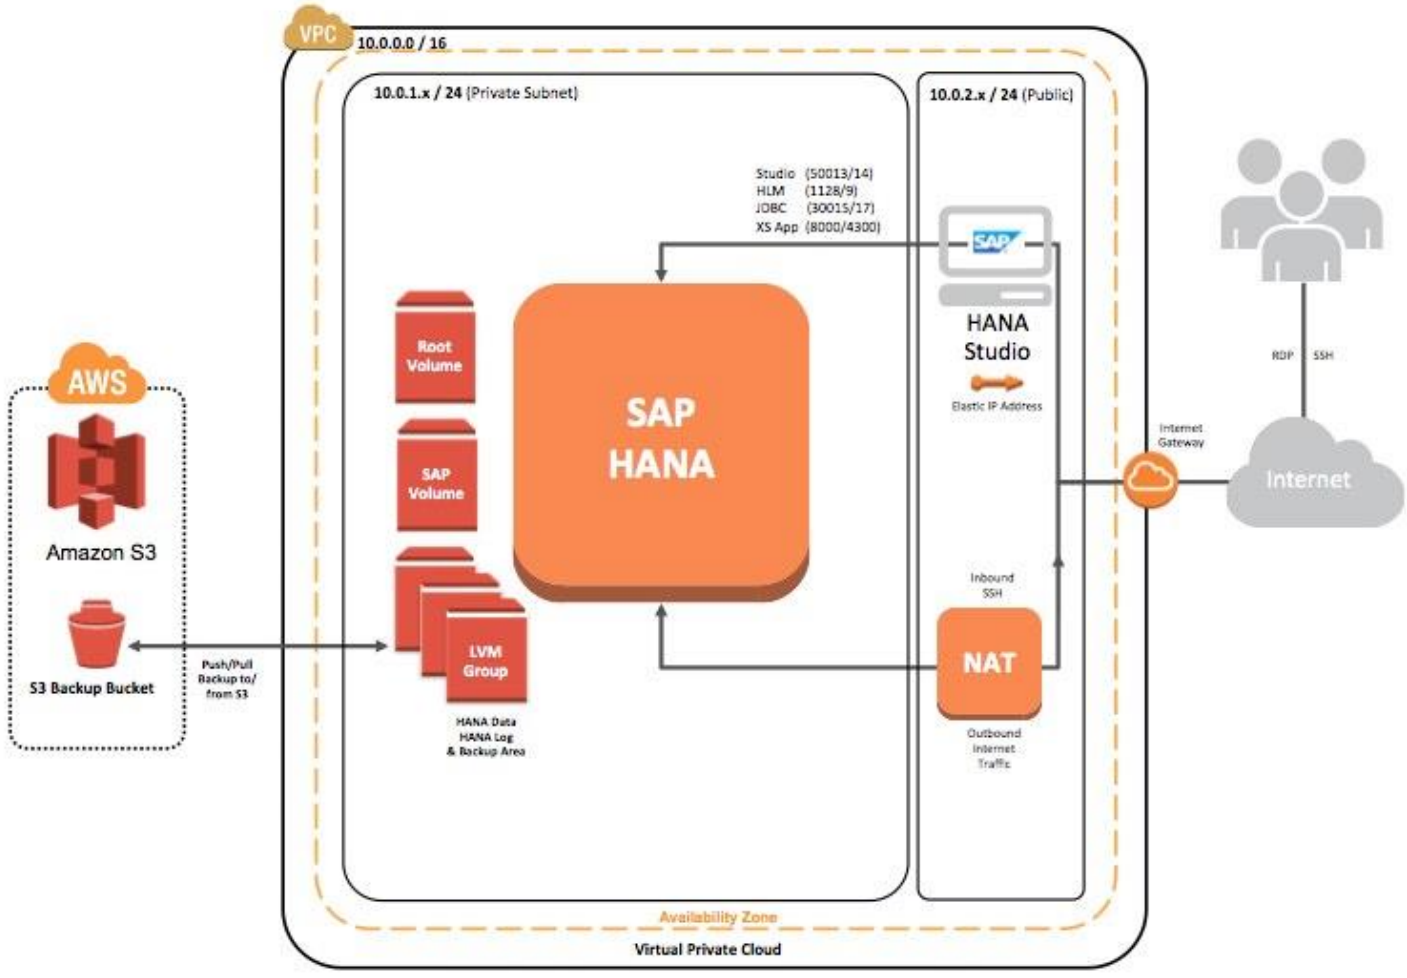

Figure 1: Single Node Architecture

Multi-Node Architecture

Multi-node deployments additionally automatically install the worker nodes based on the deployment selection for number of nodes. Worker nodes are also deployed into the same subnet as the Master node.

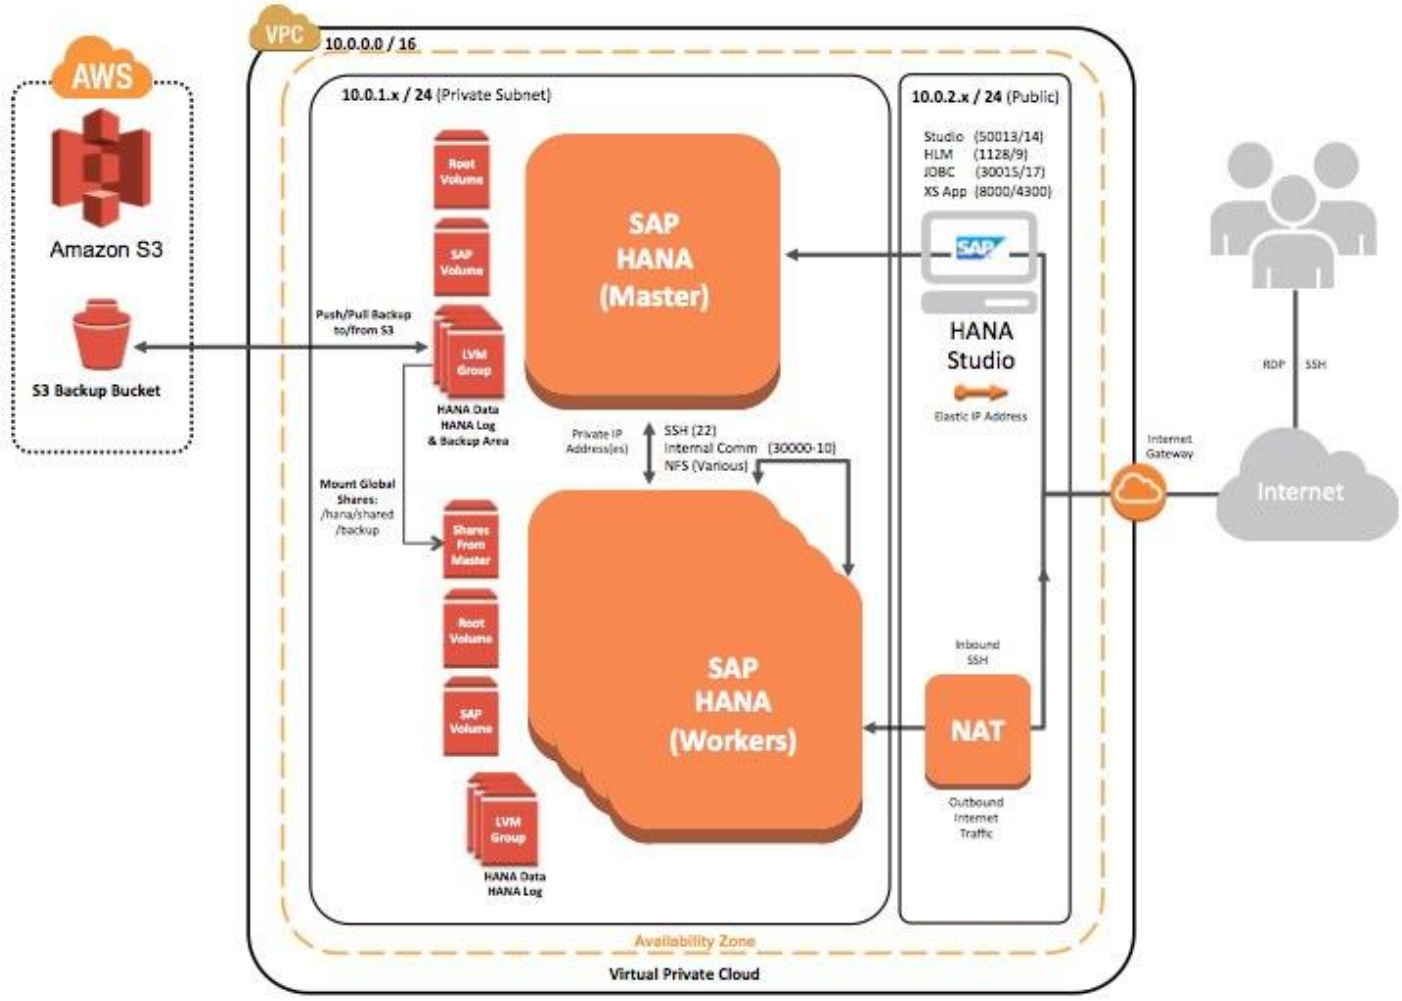

Figure 2: Multi-Node Architecture

Advanced Configurations

The provisioning process starts on the saphana.com website where the user is required to enter their AWS account. Upon submission, SAP grants access to a private Amazon Machine Image (AMI), which is used during the deployment process. After selecting the number of HANA nodes desired the users browser is redirected to an AWS CloudFormation template depending on the number of nodes selected. At this point a custom CloudFormation template can be substituted instead in order to "customize" the deployment. For example, if a customer already had an existing VPC where they wanted to deploy the solution this could be accomplished by specifying additional parameters upfront. See appendix A for sample custom CloudFormation templates.

Storage Architecture

In order to meet the HPC requirements of SAP HANA, the storage configuration used for SAP HANA on AWS is optimized for both price and performance based on KPI's provided by SAP through the SAP HANA Tailored Datacenter Integration program. As long as the deployment is done using the standard provisioning process through saphana.com and AWS CloudFormation, the storage configuration is built using an SAP supported configuration.

The storage configuration for SAP HANA on AWS is based on Elastic Block Store (EBS) Provisioned IOPS (P-IOPS) volumes. AWS Elastic Block Store (EBS) provides persistent block level storage volumes for use with EC2 instances. EBS volumes are off-instance storage that persists independently from the life of an instance.

Provisioned IOPS volumes offer storage with consistent low-latency performance, and are designed for applications with I/O-intensive workloads such as SAP HANA. Backed by Solid-State Drives (SSDs), provisioned IOPS volumes can achieve single digit millisecond latencies and are designed to deliver within 10% of the provisioned IOPS performance 99.9% of the time. Furthermore, volume striping allows for significant IOPS and throughput performance.

Each Amazon EBS volume is automatically replicated within its Availability Zone to protect from component failures, offering high availability and durability. As such, the production configuration is based on 12 x 200GB x 2000 P-IOPS volumes striped together in a Raid-0 configuration. Each SAP HANA node carries the same EBS configuration regardless of whether it is configured as Master or worker node.

The solution also uses a shared nothing storage concept for the data and log area so a single HANA node failure does not impact the availability of all the storage for the solution. However, the backup and HANA Shared file systems are owned by the HANA Master node and shared via NFS to all worker and standby nodes as per SAP best practices.

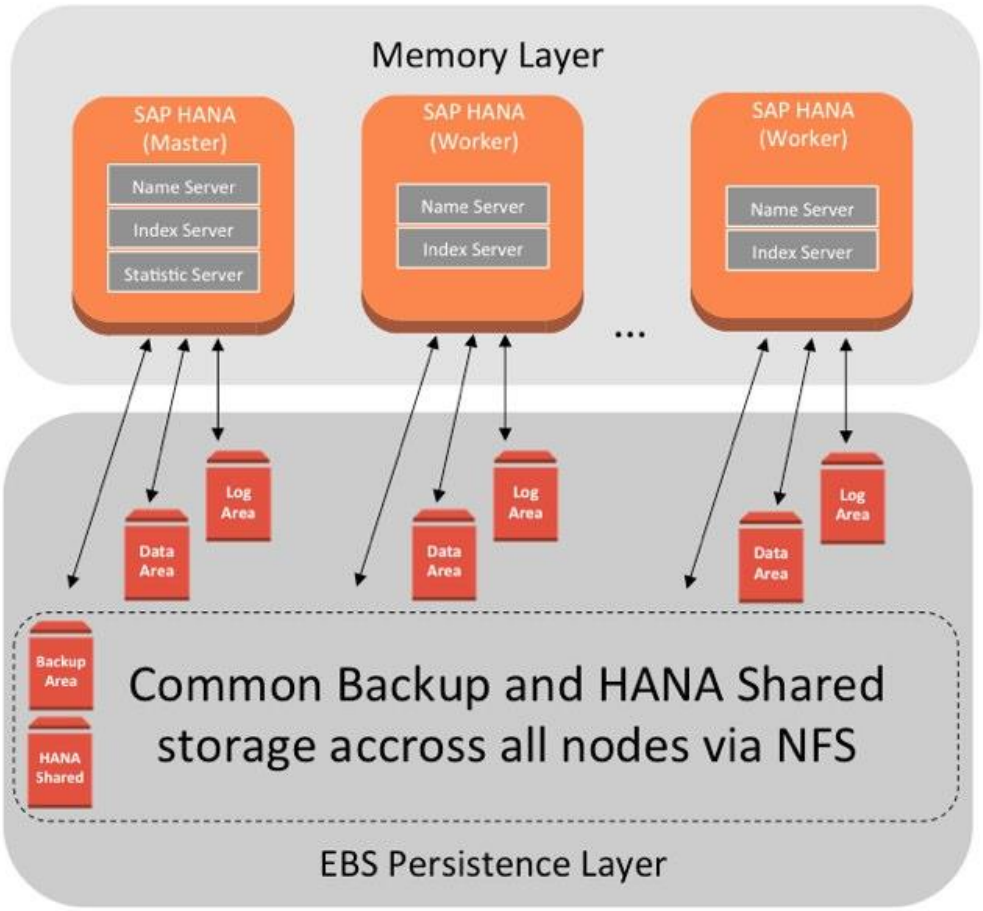

Figure 3: EBS Persistence Architecture

Tip

EBS Standard storage volumes can easily be substituted for non-production and proof of concept environments when storage performance is not critical. See Appendix A for modified CloudFormation templates and instructions.

Deployment

Note The information in this section has not yet been updated with recent changes in the deployment process. For up-to-date information, see the SAP HANA on AWS Quick Start Deployment Guide.

Preparation

1. Create an Amazon Web Services (AWS) account, if needed. – http://aws.amazon.com
2. Choose an EC2 Region to deploy the SAP HANA on AWS solution.

   Amazon EC2 locations are composed of Regions and Availability Zones. Regions are dispersed and located in separate geographic areas. Currently, the BYOL version of HANA on AWS can be deployed in the following AWS regions:

   - US, Northern Virginia (us-east-1)
   - US, Oregon (us-west-2)
   - Ireland, EU (eu-west-1)
   - Tokyo, Japan (ap-northeast-1)

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

3. Choose an Availability Zone within the region of your choice.

   Availability Zones are distinct locations within a Region that are engineered to be isolated from failures in other Availability Zones and provide inexpensive, low latency network connectivity to other Availability Zones in the same Region.

   Note:
   In most cases, choosing the first availability zone in your region should be sufficient. For example, in us-east-1, the first availability zone would be us-east-1a. For us-west-2, this would be us-west-2a and so forth.

   To find the availability zones available in your particular region:
   a. Sign in to the AWS Management Console and open the Amazon EC2 console at https://console.aws.amazon.com/ec2/
   b. From the navigation bar, view the options in the region selector.

Figure 4: Region Selection

c. After you select a region, you can view your Availability Zones within that region directly on the main EC2 page.

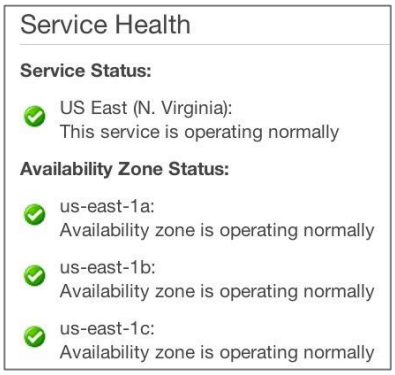

Figure 5: Availability Zones

Tip

In the case of us-east-1 (Virginia) and ap-northeast-1 (Tokyo) there are Availability Zones that do not support VPC. If you receive the message *"Value for parameter availability zone is invalid. Subnets can currently only be created in the following availability zones,"* you will need to choose a different availability zone for your deployment.

4. Create a key-pair in your preferred region.

   Amazon EC2 uses public–key cryptography to encrypt and decrypt login information. To be able to log into your instances, you must create a key pair. You will use this key pair to log into the Linux instance where HANA is installed using SSH. With Windows instances, use the key pair to obtain the administrator password via the EC2 console and then log in using RDP.
   Step-by-step instructions

5. If applicable, request EC2 and/or EBS limit increases.

   a. If you plan on deploying more than a **single** node, please request a limit increase for Elastic Block Store (EBS) Provisioned IOPS volumes here. Request 24,000 x the number of nodes you plan on deploying. For example a 5-node deployment you would add to your current Provisioned IOPS limit 5*24000, so that you would request from AWS 10,000 + 120,000 Provisioned IOPS. 10,000 is the default provisioned IOPS limit. There is no charge associated with extending the limit for Provisioned IOPS.

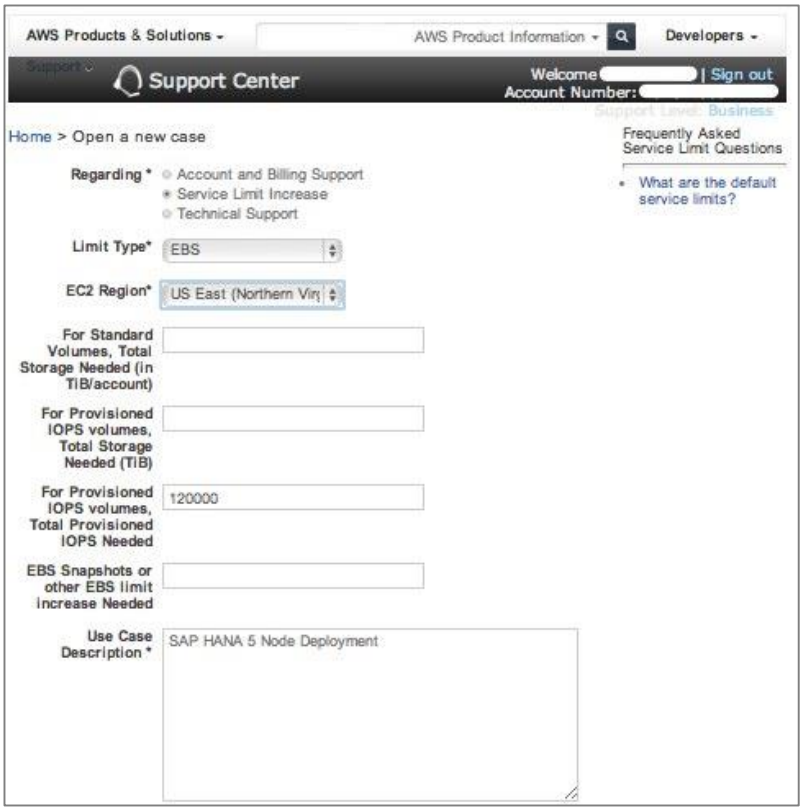

Figure 6: Sample EBS Limit Increase Request

b. If you plan on deploying more than two SAP HANA nodes, please request a limit increase for the CR1 instance type here. ***By default, each AWS account starts with a limit of 2***.

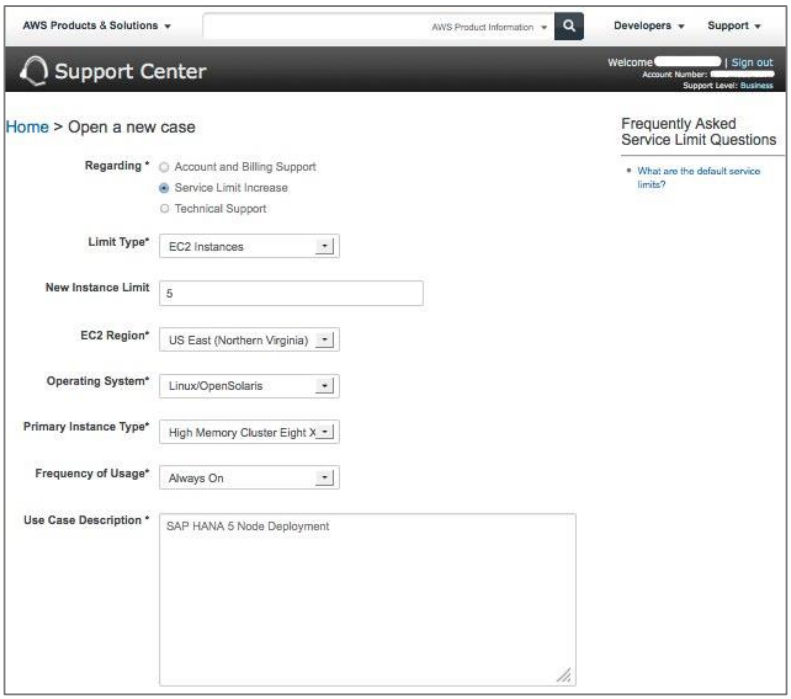

Figure 7: Sample EC2 Limit Increase Request

Receive SAP HANA Images

1. Navigate to the SAP HANA Marketplace offering.
2. Click Deploy Now.
3. Enter your Amazon Web Services account number and, if known, your SAP customer id.
4. Click OK.
5. Your screen will now look like the following:

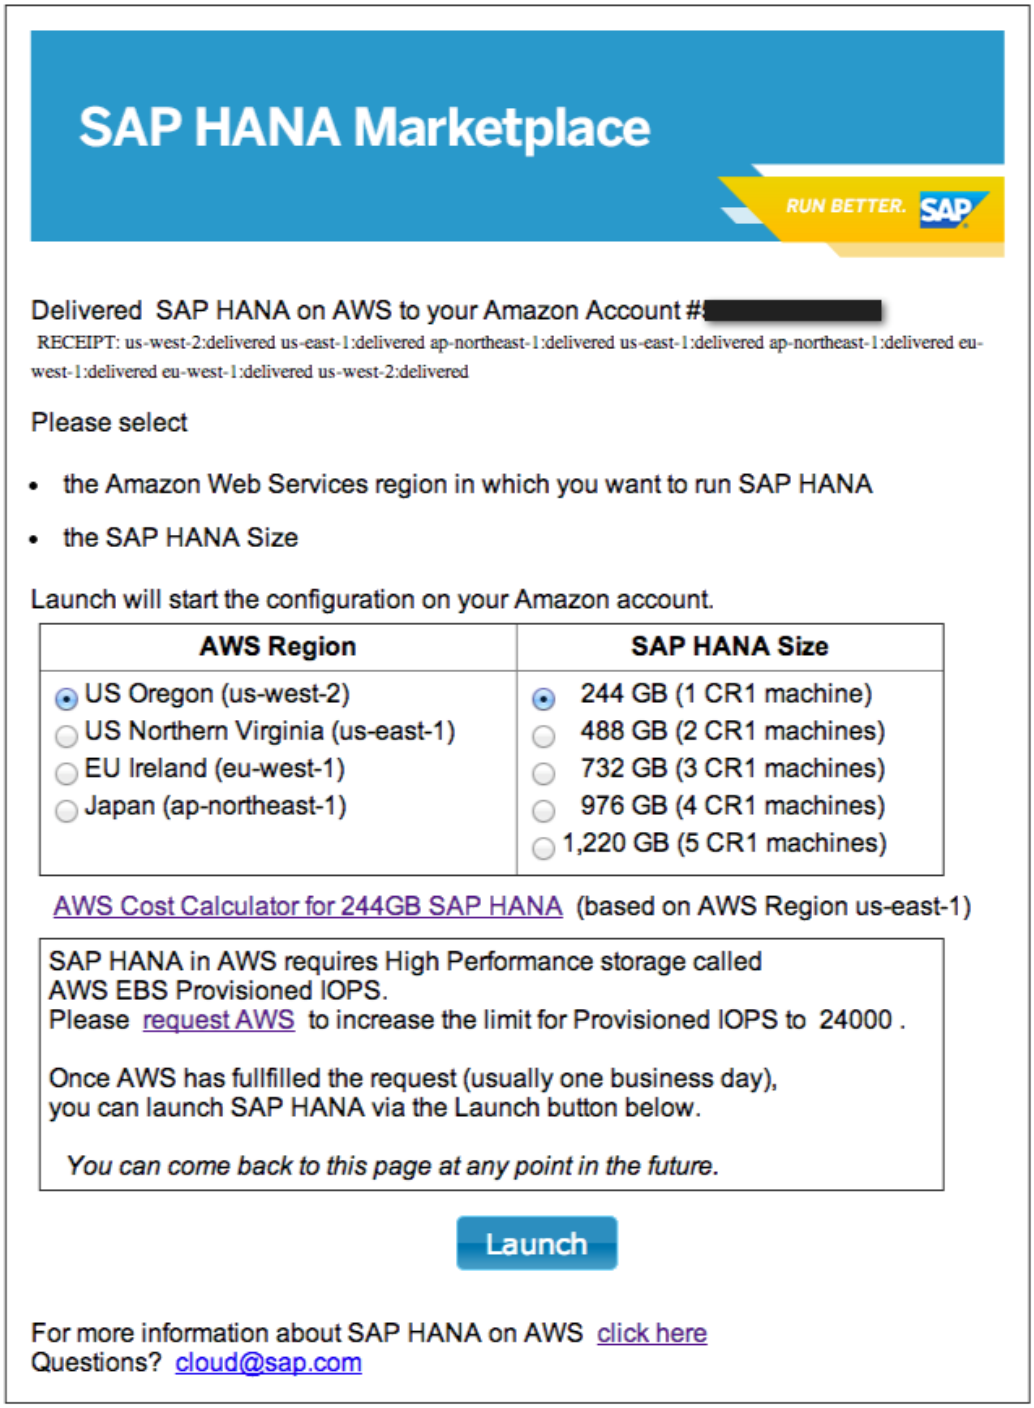

Figure 8: Deployment Selection

Deploy the SAP HANA Solution

1. Select the AWS Region and SAP HANA Size.
2. For the SAP HANA Size 244 GiB you can proceed to Launch and skip the following steps.
3. If not already done previously, click on the link “request AWS” to increase the limit for Provisioned IOPS to the specified number. Also request a limit increase for the “High Memory Cluster Eight XL” instance type if deploying more than 2 nodes. Wait until Amazon completes the request (usually within one business day).
4. Click Launch and log into your AWS account if needed.
5. Specify a name of the Stack

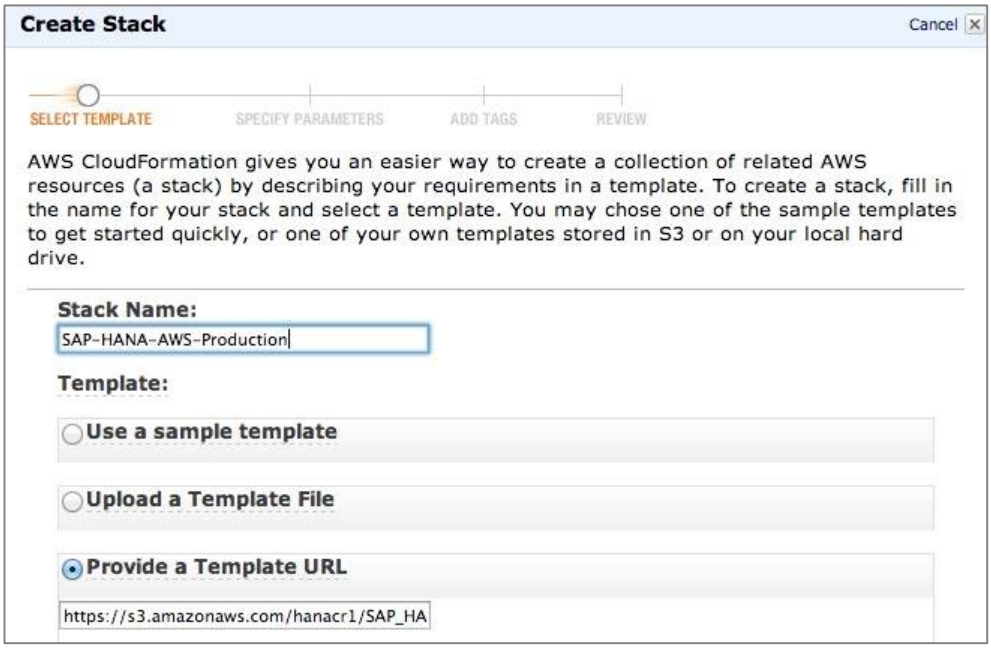

Figure 9: Cloud Formation – Choose Stack Name

6. On the next page:
   a. Specify a SID for the HANA System.
   b. Specify a CIDR range that will have SSH access to the NAT instance (TCP/22) and RDP access (TCP/3389) to the HANA Studio instance in the public subnet. A default of 0.0.0.0/0 will allow access from any IP address.

      Important:
      As a security best practice, we recommend you restrict this to your own specific CIDR Range or IP address.

   c. Enter a Master Password. This password will be used to set the password for OS users <sid>adm, sapadm, and the HANA SYSTEM DB user.
   d. Enter the Availability Zone of your choice from step 2.3c above.
   e. Specify the Key-Pair name created in Step 2.4 above.

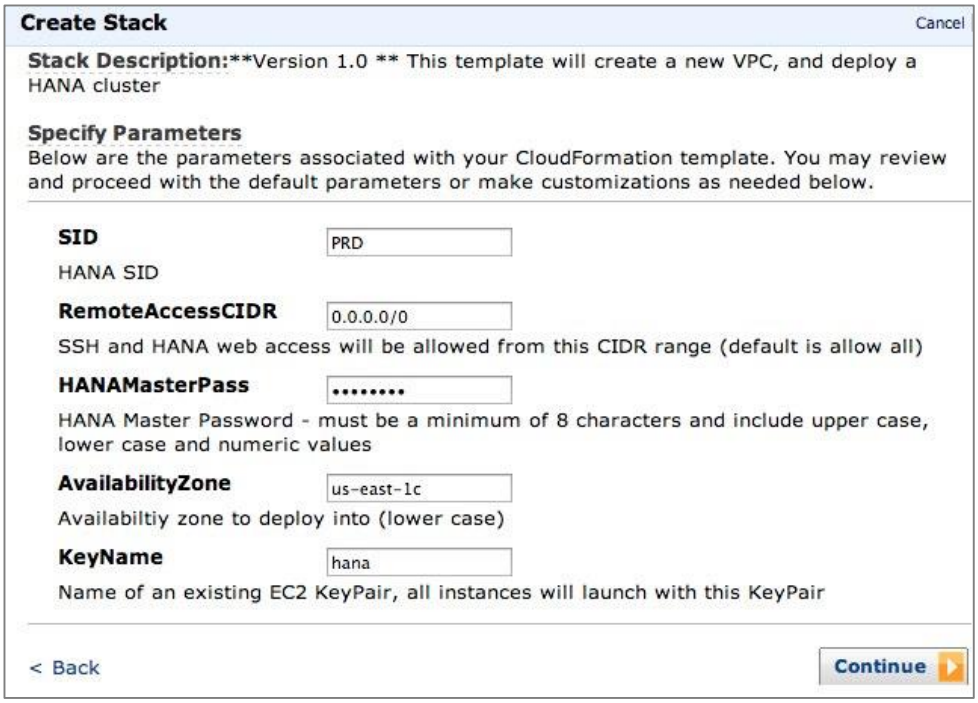

Figure 10: Cloud Formation – Deployment Parameters

f. The screen "Add Tags" is optional. Information on how tagging works in AWS can be found here.
g. Review the information and press Continue to initiate the provisioning. If you receive any warnings about the parameters you entered, use the back selection to go back and fix them.
h. Wait until the Stack is marked as CREATE_COMPLETED.

7. Monitor the provisioning process.

a. You will immediately be able to track the status of the deployment process in the description tab of the CloudFormation stack.

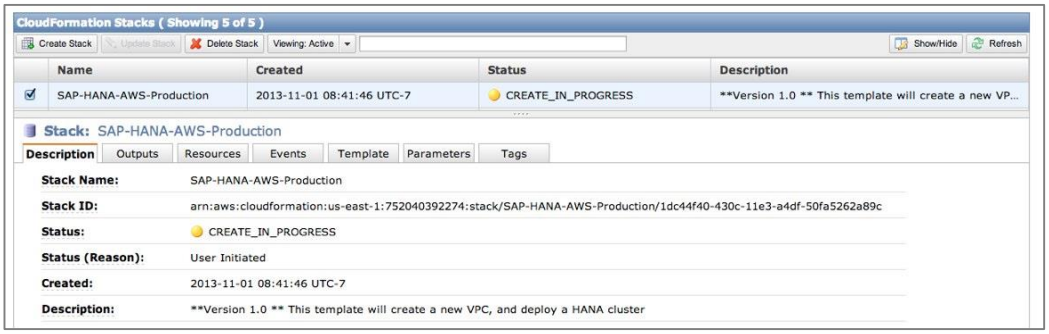

Figure 11: Cloud Formation – Overall Deployment Status

b. To see progress of the individual component and system deployments, navigate to the Events tab. Here you can monitor the progress of the entire CloudFormation stack deployment process.

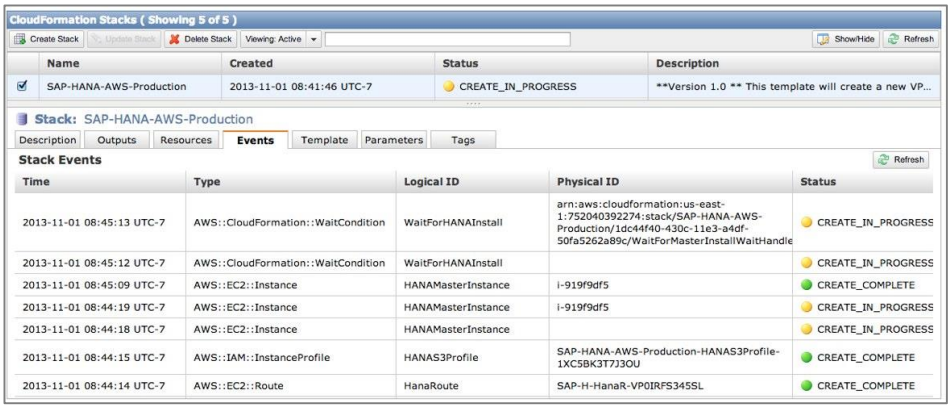

Figure 12: Cloud Formation – Overall Deployment Status

Note:
Single node SAP HANA deployments can take anywhere from 10-15 minutes.
Multi-node SAP HANA deployments will take 10-15 minutes for the master node and an additional 10-15 minutes for all worker nodes as all worker nodes are deployed in parallel.

c. Once the create process is complete you will see the stack marked as CREATE_COMPLETED.

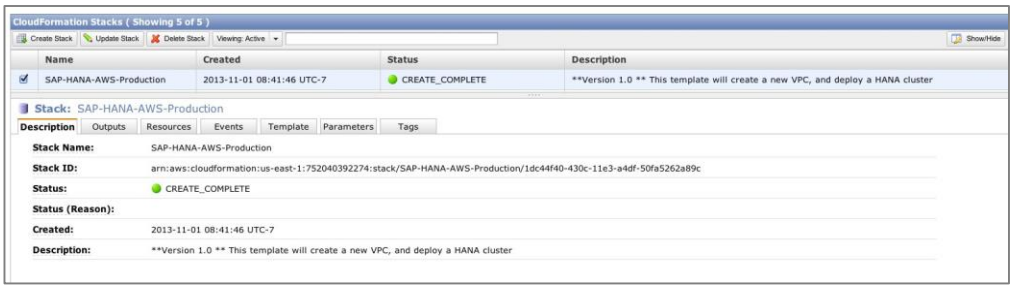

Figure 13: Cloud Formation – Create Complete

d. If you encounter the status message ROLLBACK_IN_PROGRESS or ROLL_BACK_COMPLETE, please see the next section for troubleshooting.

Troubleshooting

Most provisioning errors can be attributed to problems with account limits. If you see a ROLLBACK_IN_PROGRESS or ROLLBACK_COMPLETE status message check the events tab of the failed CloudFormation stack to determine which resource first attributed to the ROLLBACK event.

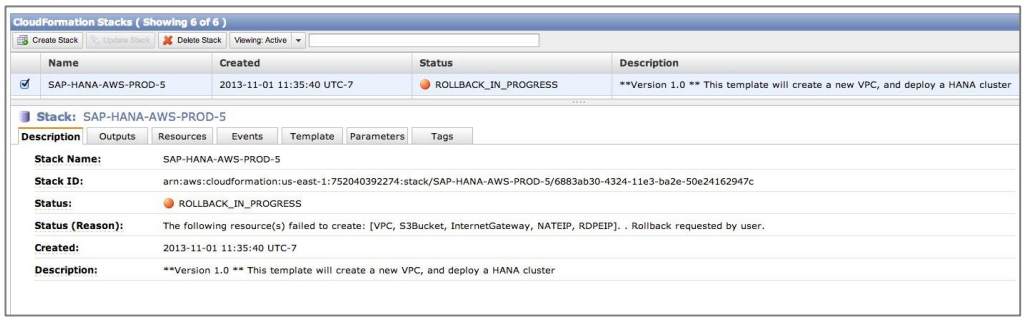

Figure 14: CloudFormation – Rollback Example

Start from the bottom and scroll up until you see the first CREATE_FAILED event. You may need to scroll to the right to see the actual error message.

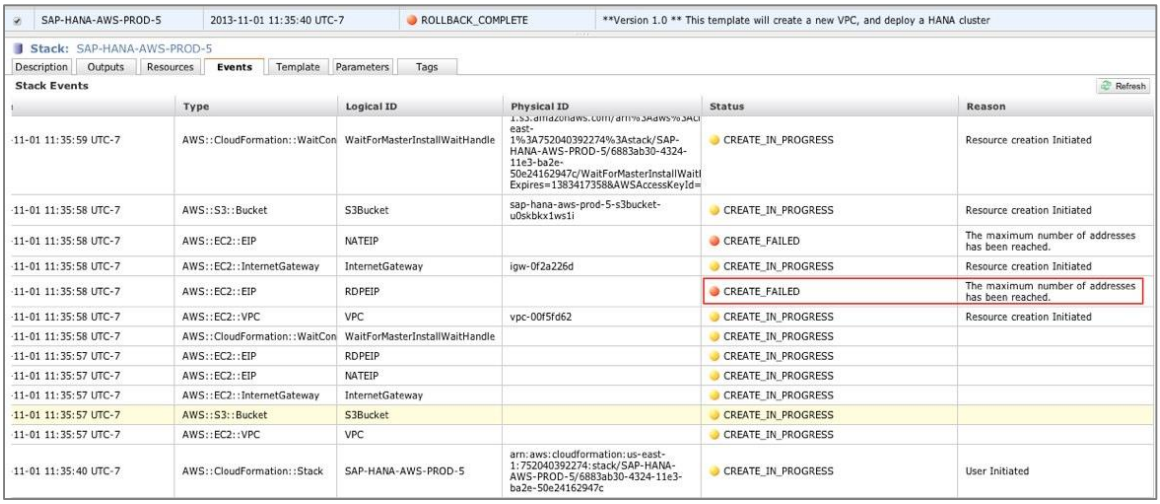

SAP-HANA-AWS-PROD-5	2013-11-01 11:35:40 UTC-7	ROLLBACK_COMPLETE	**Version 1.0 ** This template will create a new VPC, and deploy a HANA cluster	

Stack: SAP-HANA-AWS-PROD-5

Description | Outputs | Resources | **Events** | Template | Parameters | Tags

Stack Events (Refresh)

	Type	Logical ID	Physical ID	Status	Reason
11-01 11:35:59 UTC-7	AWS::CloudFormation::WaitCon	WaitForMasterInstallWaitHandle	1%3A752040392274%3Astack/SAP-HANA-AWS-PROD-5/6883ab30-4324-11e3-ba2e-50e24162947c/WaitForMasterInstallWait Expires=1383417358&AWSAccessKeyId=	CREATE_IN_PROGRESS	Resource creation Initiated
11-01 11:35:58 UTC-7	AWS::S3::Bucket	S3Bucket	sap-hana-aws-prod-5-s3bucket-u0skbkx1ws1i	CREATE_IN_PROGRESS	Resource creation Initiated
11-01 11:35:58 UTC-7	AWS::EC2::EIP	NATEIP		CREATE_FAILED	The maximum number of addresses has been reached.
11-01 11:35:58 UTC-7	AWS::EC2::InternetGateway	InternetGateway	igw-0f2a226d	CREATE_IN_PROGRESS	Resource creation Initiated
11-01 11:35:58 UTC-7	AWS::EC2::EIP	RDPEIP		CREATE_FAILED	The maximum number of addresses has been reached.
11-01 11:35:58 UTC-7	AWS::EC2::VPC	VPC	vpc-00f5fd62	CREATE_IN_PROGRESS	Resource creation Initiated
11-01 11:35:58 UTC-7	AWS::CloudFormation::WaitCon	WaitForMasterInstallWaitHandle		CREATE_IN_PROGRESS	
11-01 11:35:57 UTC-7	AWS::EC2::EIP	RDPEIP		CREATE_IN_PROGRESS	
11-01 11:35:57 UTC-7	AWS::EC2::EIP	NATEIP		CREATE_IN_PROGRESS	
11-01 11:35:57 UTC-7	AWS::EC2::InternetGateway	InternetGateway		CREATE_IN_PROGRESS	
11-01 11:35:57 UTC-7	AWS::S3::Bucket	S3Bucket		CREATE_IN_PROGRESS	
11-01 11:35:57 UTC-7	AWS::EC2::VPC	VPC		CREATE_IN_PROGRESS	
11-01 11:35:40 UTC-7	AWS::CloudFormation::Stack	SAP-HANA-AWS-PROD-5	arn:aws:cloudformation:us-east-1:752040392274:stack/SAP-HANA-AWS-PROD-5/6883ab30-4324-11e3-ba2e-50e24162947c	CREATE_IN_PROGRESS	User Initiated

Figure 15: CloudFormation – Create Failed

If you get an error that the "instance did not stabilize" (as below) this means you have exceeded your IOPS for the region and need to request an increase.

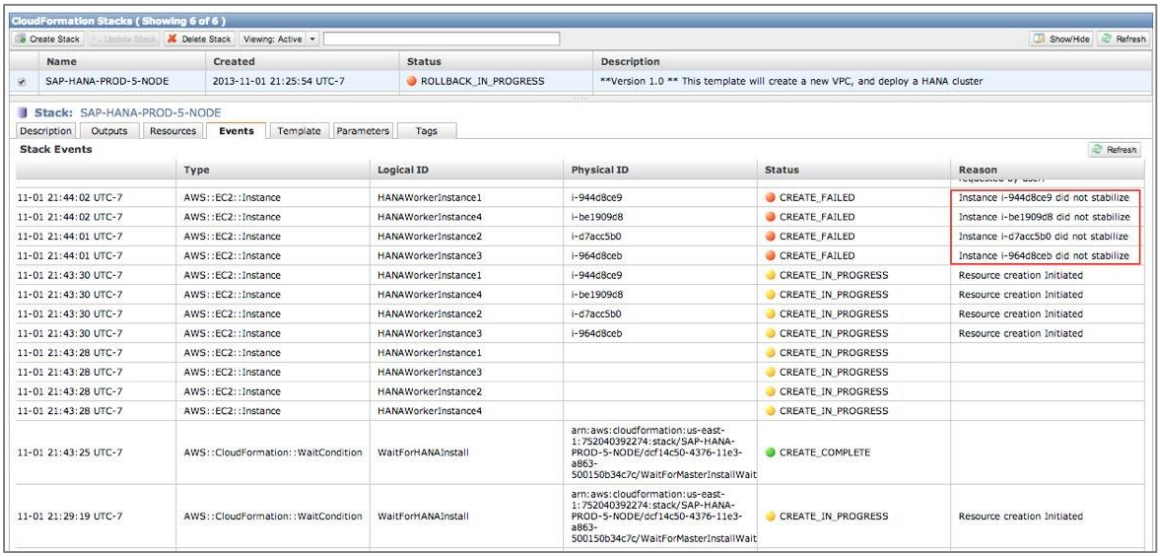

CloudFormation Stacks (Showing 6 of 6)

Create Stack | Update Stack | Delete Stack | Viewing: Active | Show/Hide | Refresh

	Name	Created	Status	Description
	SAP-HANA-PROD-5-NODE	2013-11-01 21:25:54 UTC-7	ROLLBACK_IN_PROGRESS	**Version 1.0 ** This template will create a new VPC, and deploy a HANA cluster

Stack: SAP-HANA-PROD-5-NODE

Description | Outputs | Resources | **Events** | Template | Parameters | Tags

Stack Events (Refresh)

	Type	Logical ID	Physical ID	Status	Reason
11-01 21:44:02 UTC-7	AWS::EC2::Instance	HANAWorkerInstance1	i-944d8ce9	CREATE_FAILED	Instance i-944d8ce9 did not stabilize
11-01 21:44:02 UTC-7	AWS::EC2::Instance	HANAWorkerInstance4	i-be1909d8	CREATE_FAILED	Instance i-be1909d8 did not stabilize
11-01 21:44:01 UTC-7	AWS::EC2::Instance	HANAWorkerInstance2	i-d7acc5b0	CREATE_FAILED	Instance i-d7acc5b0 did not stabilize
11-01 21:44:01 UTC-7	AWS::EC2::Instance	HANAWorkerInstance3	i-964d8ceb	CREATE_FAILED	Instance i-964d8ceb did not stabilize
11-01 21:43:30 UTC-7	AWS::EC2::Instance	HANAWorkerInstance1	i-944d8ce9	CREATE_IN_PROGRESS	Resource creation Initiated
11-01 21:43:30 UTC-7	AWS::EC2::Instance	HANAWorkerInstance4	i-be1909d8	CREATE_IN_PROGRESS	Resource creation Initiated
11-01 21:43:30 UTC-7	AWS::EC2::Instance	HANAWorkerInstance2	i-d7acc5b0	CREATE_IN_PROGRESS	Resource creation Initiated
11-01 21:43:30 UTC-7	AWS::EC2::Instance	HANAWorkerInstance3	i-964d8ceb	CREATE_IN_PROGRESS	Resource creation Initiated
11-01 21:43:28 UTC-7	AWS::EC2::Instance	HANAWorkerInstance1		CREATE_IN_PROGRESS	
11-01 21:43:28 UTC-7	AWS::EC2::Instance	HANAWorkerInstance3		CREATE_IN_PROGRESS	
11-01 21:43:28 UTC-7	AWS::EC2::Instance	HANAWorkerInstance2		CREATE_IN_PROGRESS	
11-01 21:43:28 UTC-7	AWS::EC2::Instance	HANAWorkerInstance4		CREATE_IN_PROGRESS	
11-01 21:43:25 UTC-7	AWS::CloudFormation::WaitCondition	WaitForHANAInstall	arn:aws:cloudformation:us-east-1:752040392274:stack/SAP-HANA-PROD-5-NODE/dcf14c50-4376-11e3-a863-500150b34c7c/WaitForMasterInstallWait	CREATE_COMPLETE	
11-01 21:29:19 UTC-7	AWS::CloudFormation::WaitCondition	WaitForHANAInstall	arn:aws:cloudformation:us-east-1:752040392274:stack/SAP-HANA-PROD-5-NODE/dcf14c50-4376-11e3-a863-500150b34c7c/WaitForMasterInstallWait	CREATE_IN_PROGRESS	Resource creation Initiated

Figure 16: CloudFormation – Instance Did Not Stabilize

If you get an error “Value for parameter availabilityZone is invalid. Subnets can currently only be created in the following availability zones,“ you will need to choose a different availability zone for your deployment.

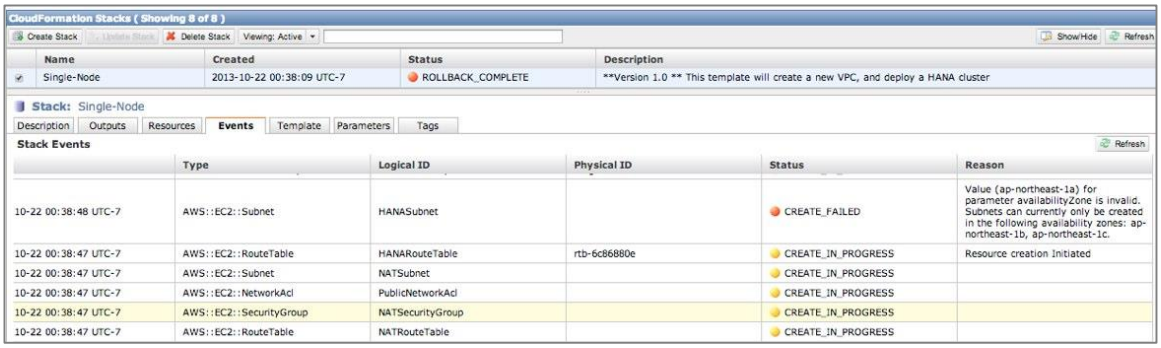

Figure 17: CloudFormation – Choose another Availability Zone

Getting Access to SAP HANA

The default network security setup of this solution follows security best practices of AWS. The provisioning logic creates the solution architecture described in the solution 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 only be accessed through instances placed in the public subnet or DMZ layer.

Through this DMZ layer, two methods of access are available.

- **HANA Studio Access**
  Connect to the Windows Instance using a Remote Desktop Client where HANA Studio has been preloaded.
- **OS Level Access**
  SSH to the NAT instance and then to the SAP HANA instance(s) using a SSH client of your choice.

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 datacenter and your VPC. See http://aws.amazon.com/vpc/ for more details.

HANA Studio Access using the RDP Instance

1. In the output window please note down the Elastic IP address (EIP) of the RDP Instance.

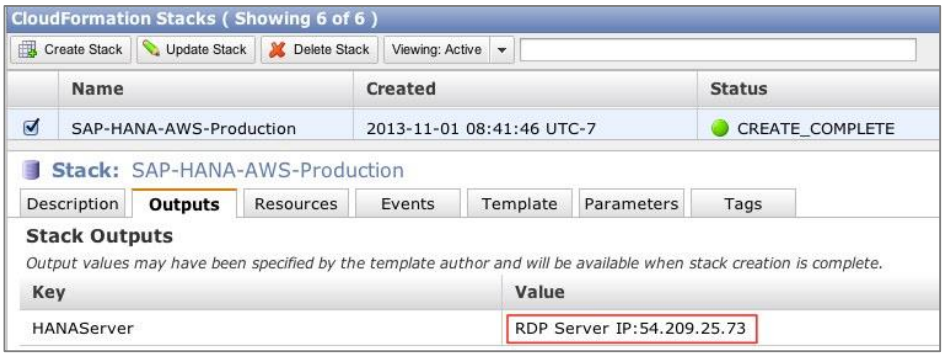

Figure 18: Cloud Formation – RDP Server IP info

2. Get the Windows Administrator Password from the EC2 console.
   - Go to Services -> EC2 -> Instances -> Select your RDP Instance
   - Choose Connect -> Get password
   - Choose or paste in the contents of your private key in the space provided.
   - The password will be decrypted and shown to you.

3. Choose Download Remote Desktop File or connect via an RDP client of your choice.

4. Start HANA Studio and add a System
   - IP Address or hostname of Master Node (imdbmaster)
   - Instance Number: 00
   - User: SYSTEM
   - Password: < your password from 5.3.c >

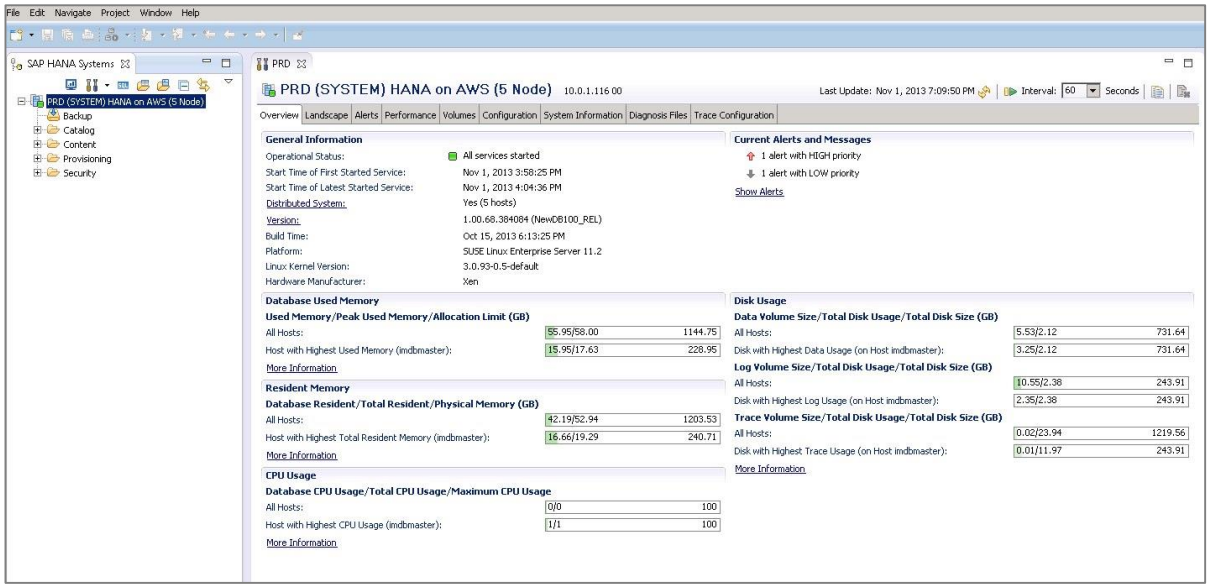

Figure 19: HANA Studio – SAP HANA Overview

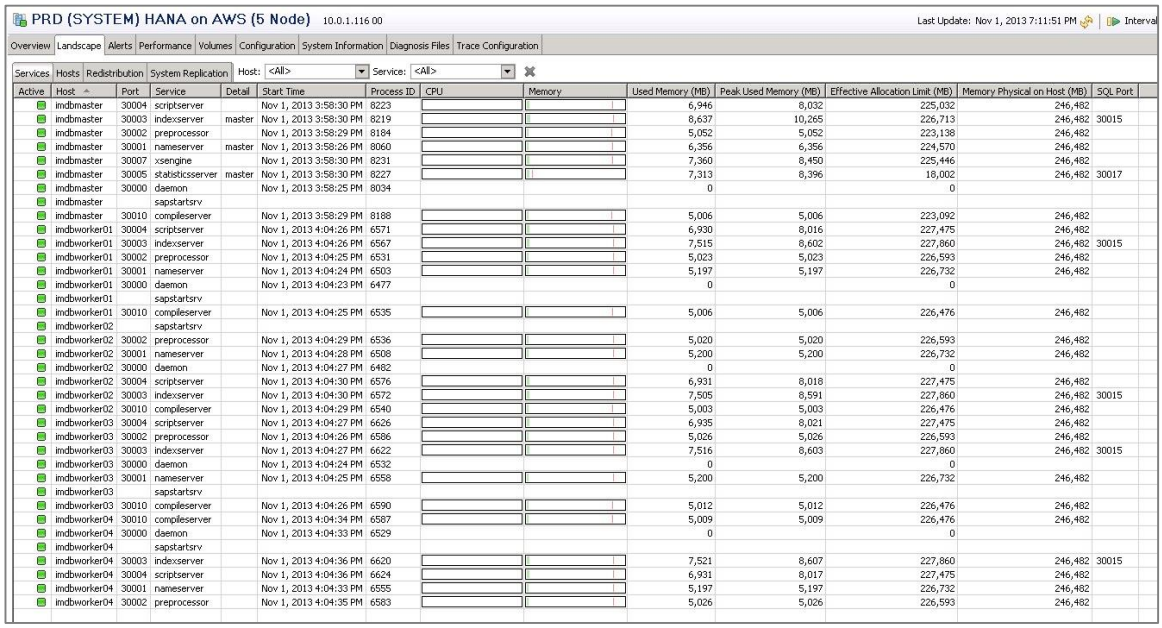

Figure 20: HANA Studio – Landscape Status

Note

We recommend you take a backup at this point. This can be done via HANA Studio for HANA. You can also take complete system image (Amazon Machine Image) through the EC2 console for recovery later.

SSH Access

1. Navigate to Services -> EC2 -> Instances and find your NAT instance and note the public Elastic IP Address.

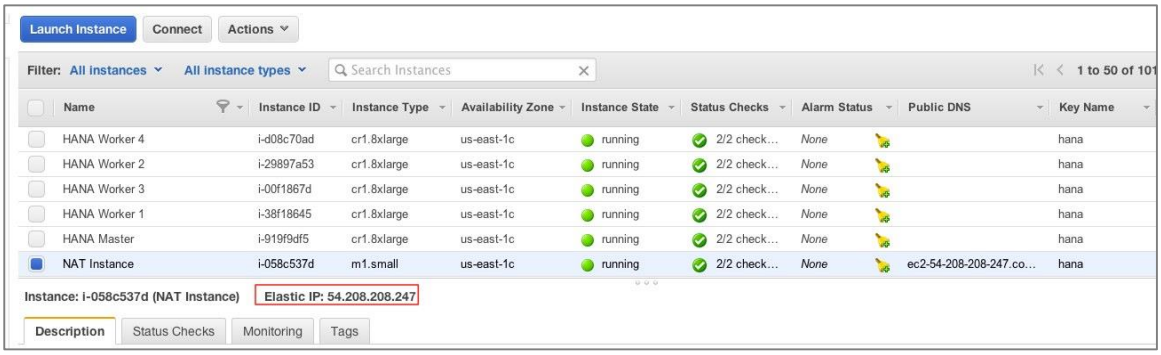

Figure 21: NAT – Elastic IP Address

2. Using an ssh client of your choice (i.e. Putty or ITerm), ssh into the NAT instance using the key-pair specified during the deployment process.

Note
If your connection times out, you may need to adjust the security group rules for the NAT instance to allow access from your computers IP address or proxy server.

ITerm Example:

- Add private key to authentication agent (ssh-add)
- ssh to NAT instance with –A option to forward the key.
- Note that entries for the servers hosting SAP HANA have already been maintained in /etc/hosts.
- ssh to the SAP HANA server

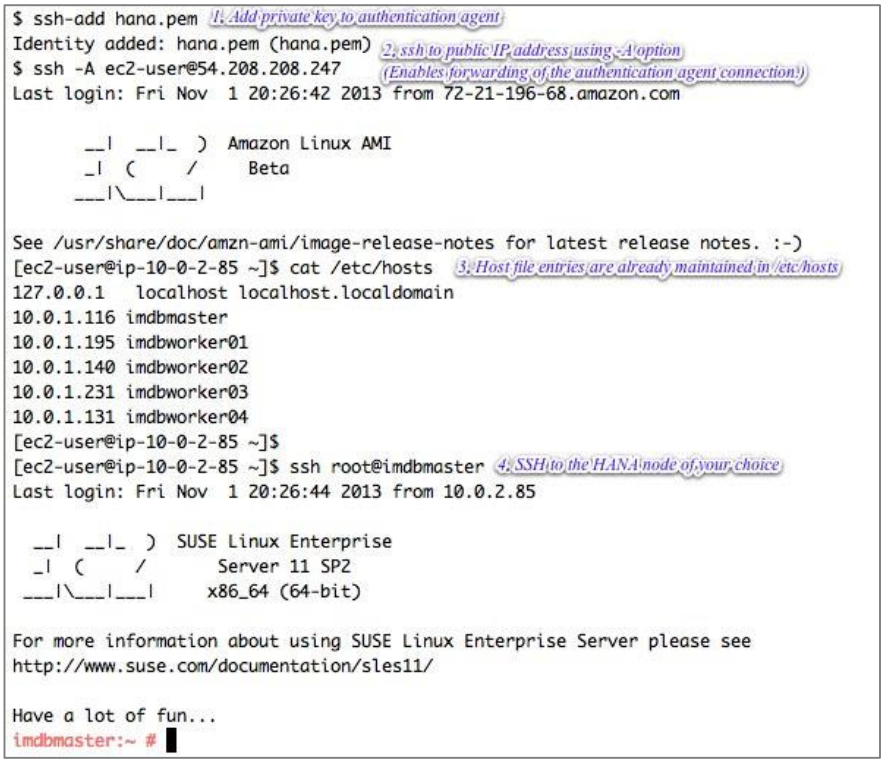

Figure 22: SSH – ITERM Example

Putty Example:

a. Download Putty.exe, Puttygen.exe, and pageant.exe
b. Load your private key into Puttygen and save as ppk file that putty can use.
c. Execute Pageant.exe, and add your new ppk key. The Pageant process must be running in order for agent forwarding to work.
d. Configure putty with the private key and choose allow agent forwarding.

Figure 23: SSH – Putty Example

e. Save configuration
f. SSH to NAT instance, then to HANA node

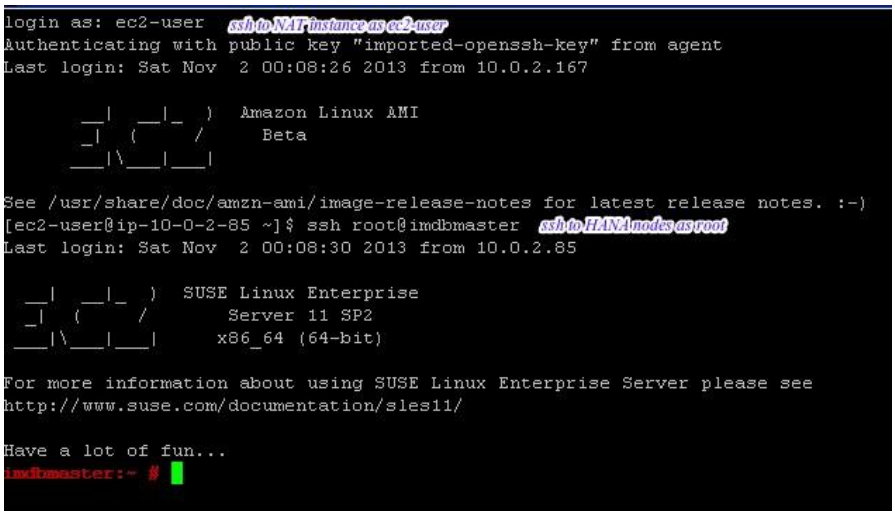

Figure 24: SSH – Putty Example Continued

Administration

Start / Stop of EC2 instances running SAP HANA Hosts

At any time one or multiple SAP HANA Hosts can be stopped. Before stopping the EC2 instance of an SAP HANA host, it is recommended to first stop SAP HANA on that instance. When resuming the EC2 Instance, the instance will automatically be started with the same IP address, network, and storage configuration as before.

Creating an Image of a SAP HANA System

There are multiple reasons for creating an image of a SAP HANA System. These include:

- Create a full system backup (OS, /usr/sap, HANA Shared, Backup, Data, Log) via Amazon Machine Image (AMI). Amazon Machine Images are automatically saved in 3 different availability zones within the same region.
- Change Storage Performance
  During instantiation of an Image it is possible to specify EBS performance ranging from EBS Standard to EBS Provisioned IOPS with 4000 IOPS per Volume. The default storage performance for SAP HANA is 2000 IOPS per Volume. Storage performance has a significant impact on AWS infrastructure cost.
- Relocating a HANA system from one region to another.
  This can be done by leveraging Image Copy and specify the new target region. The SAP HANA System can be resumed in the new region

Tip
The SAP HANA system should be in a consistent state before creating an Amazon Machine Image (AMI). This can be accomplished by stopping the SAP HANA Instance before creation or by following instructions in SAP Note 1703435.

Cloning a SAP HANA System

Cloning a SAP HANA System via imaging and re-instantiating is currently only supported for a HANA system with a single Host.

In order to clone a multi-Host SAP HANA deployment:

1. Provision a new SAP HANA system with the same configuration.
2. Perform a data backup of the original system
3. And restore the backup on the new system.

Backup/Recovery

Apart from some examples, this guide does not include detailed instructions how to execute database backups using either native HANA backup/recovery features or 3rd party backup tools. Please refer the standard OS, SAP and SAP HANA documentation or the documentation provided by the backup software vendor. In addition, backup schedules, frequency, and retention periods, are primarily based on your system type and business requirements. Please refer to the standard SAP documentation for guidance on these topics.

Note
Both general and advanced backup and recovery concepts for SAP Systems on AWS can be found in detail in the SAP on AWS Backup and Recovery Guide.

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

AWS Services and Components for Backup Solutions

Simple Storage Service (S3) – http://aws.amazon.com/s3
Amazon S3 provides a highly durable storage infrastructure designed for mission-critical and primary data storage. Amazon S3 is designed to provide 99.999999999% durability and 99.99% availability over a given year. Amazon S3 is the center of any SAP backup and recovery solution on AWS.

The deployment process automatically creates a private S3 bucket where SAP HANA backups can be stored off instance to provide more protection and durability. Only the AWS account that is used to create the bucket has access to this bucket. The S3 Bucket follows the naming convention ***<template-name-randomly_chosen_characters>*** (for example: ***node2-hana-s3bucket-gcynh5v2nqs3***).

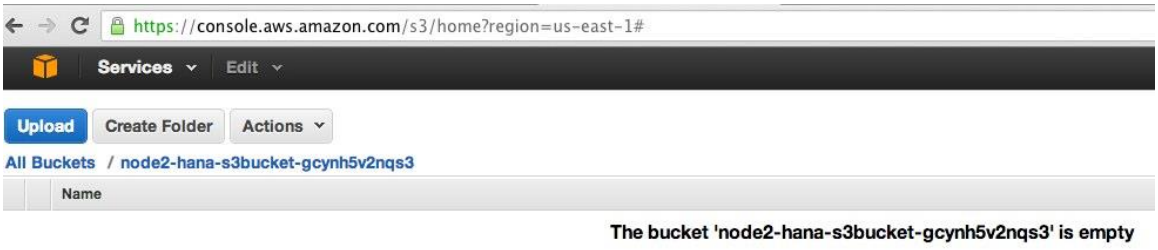

Figure 25: SSH – S3 Bucket Example

Note
Additional S3 buckets can be created if needed through the AWS console or using the AWS command line interface.

AWS Identity and Access Management
AWS Identity and Access Management (IAM) enables you to securely control access to AWS services and resources for your users. Using IAM, you can create and manage AWS users and groups and use permissions to allow and deny their access to AWS resources. You can create roles in IAM, and manage permissions to control which operations can be performed by the entity, or AWS service, that assumes the role. You can also define which entity is allowed to assume the role.

An IAM role allowing access to get/put objects to and from S3 created during the CloudFormation deployment process and is subsequently assigned to each AWS instance hosting SAP HANA master and worker nodes at launch time as they are deployed.

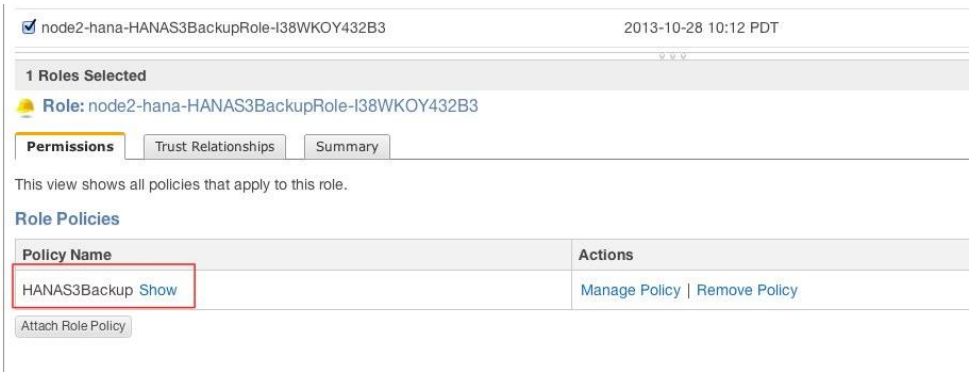

Figure 26: SSH – IAM Role Example

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

```
{"Statement":[
 {"Resource":"arn:aws:s3::: node2-hana-s3bucket-gcynh5v2nqs3/*",
   "Action":["s3:GetObject","s3:PutObject","s3:DeleteObject", "s3:ListBucket","s3:Get*","s3:List*"],
   "Effect":"Allow"},
 {"Resource":"*","Action":["s3:List*","ec2:Describe*","ec2:AttachNetworkInterface",
  "ec2:AttachVolume","ec2:CreateTags","ec2:CreateVolume","ec2:RunInstances",
  "ec2:StartInstances"],"Effect":"Allow"}]}
```

If additional functions are later desired, the IAM Role can be modified using the AWS Console.

Amazon Glacier – http://aws.amazon.com/glacier
Amazon Glacier is an extremely low-cost storage service that provides secure and durable storage for data archiving and backup. In order to keep costs low, Amazon Glacier is optimized for data that is infrequently accessed and for which retrieval times of several hours are suitable. With Amazon Glacier, customers can reliably store large or small amounts of data for as little as $0.01 per gigabyte per month, a significant savings compared to on-premises solutions. SAP HANA backups can be pushed to Glacier for long-term archival using lifecycle policies.

SAP HANA Backup Destination

The primary difference between backing up SAP systems on Amazon Web Services compared to traditional on-premises infrastructure is the backup destination. The typical backup destination used with on-premises infrastructure is tape. On AWS, instead of storing backups on tape, backups are stored in Amazon S3. There are many benefits to storing backups in Amazon S3 vs. tape. Backups stored in Amazon S3 are automatically stored “offsite” from the source system since data in Amazon S3 is replicated across multiple facilities within the AWS region.

SAP HANA Data backups can be triggered and/or scheduled using SAP HANA studio, SQL commands, or the DBA Cockpit. While log backups are written automatically (unless disabled). The /backup file system has been configured as part of the deployment process.

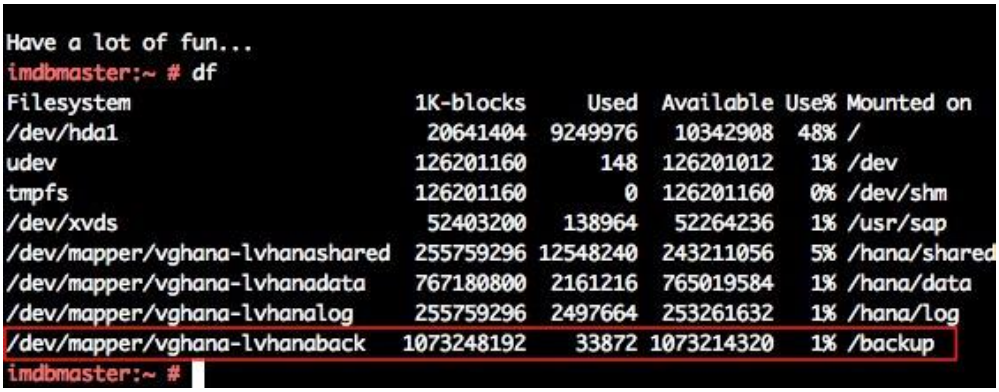

Figure 27: SSH – File system Layout

The SAP HANA global.ini configuration file has been customized as follows. Database backups go directly to /backup/data/<SID> while automatic log archival files go to /backup/log/<SID>.

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

AWS Command Line interface

The AWS Command Line Interface (CLI), which is a unified tool to manage AWS services, has already been installed as part of the base image. Using various commands you are able to control multiple AWS services from the command line directly and automate them through scripts. Access to the S3 bucket is obtained through the aforementioned IAM role assigned to the instance. Using the AWS S3 commands, we can list the contents of the previously created bucket, backup files, and restore files.

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

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

Backup Example

1. In the SAP HANA Backup editor, choose “Open Backup Wizard.” Right-clicking the system that you want to back up and choose “Back Up” can also open the backup wizard.
2. Select destination type “File.” This will back up the database to files in file system specified.
3. Specify the backup destination (/backup/data/<SID>) and the backup prefix.

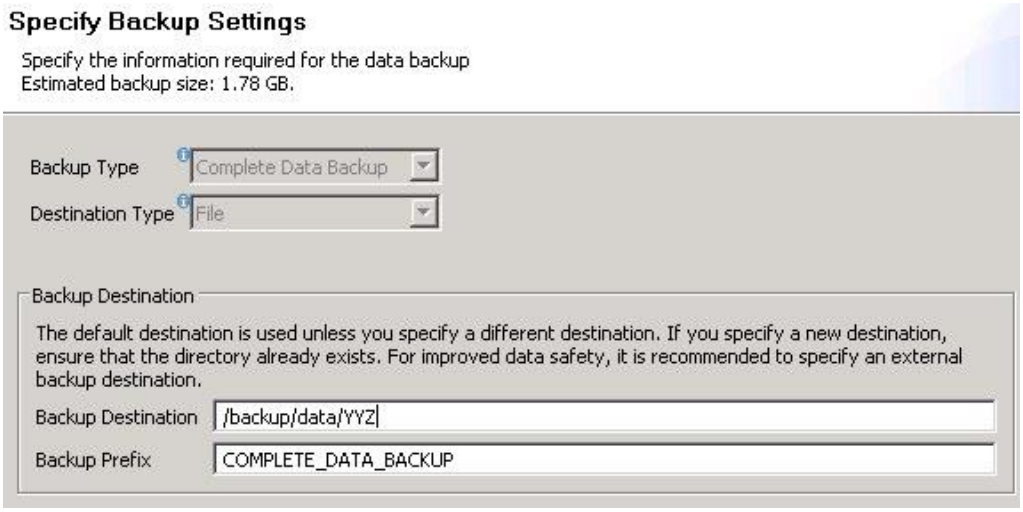

Figure 28: SSH – Backup Example

4. Chose next, and Finish
5. When the backup is complete a confirmation message will be displayed.
6. Verify the backup files are available at the operating system level.

imdbmaster:/backup # ll */*

```
data/YYZ:
total 1588080
-rw-r--r-- 1 yyzadm sapsys    163840 Oct 28 18:44 COMPLETE_DATA_BACKUP_databackup_0_1
-rw-r--r-- 1 yyzadm sapsys   70443008 Oct 28 18:44 COMPLETE_DATA_BACKUP_databackup_1_1
-rw-r--r-- 1 yyzadm sapsys 1000955904 Oct 28 18:44 COMPLETE_DATA_BACKUP_databackup_2_1
-rw-r--r-- 1 yyzadm sapsys   69292032 Oct 28 18:44 COMPLETE_DATA_BACKUP_databackup_3_1
-rw-r--r-- 1 yyzadm sapsys  101605376 Oct 28 18:44 COMPLETE_DATA_BACKUP_databackup_4_1
-rw-r--r-- 1 yyzadm sapsys   98521088 Oct 28 18:44 COMPLETE_DATA_BACKUP_databackup_5_1
-rw-r--r-- 1 yyzadm sapsys   69488640 Oct 28 18:44 COMPLETE_DATA_BACKUP_databackup_6_1
-rw-r--r-- 1 yyzadm sapsys  136269824 Oct 28 18:44 COMPLETE_DATA_BACKUP_databackup_7_1

log/YYZ:
total 34928
-rw-r--r-- 1 yyzadm sapsys    12288 Oct 28 18:44 log_backup_0_0_0_0.1382985855848
-rw-r--r-- 1 yyzadm sapsys    12288 Oct 28 18:44 log_backup_0_0_0_0.1382985856054
-rw-r--r-- 1 yyzadm sapsys    12288 Oct 28 18:44 log_backup_0_0_0_0.1382985856098
-rw-r--r-- 1 yyzadm sapsys    12288 Oct 28 18:44 log_backup_0_0_0_0.1382985856110
-rw-r--r-- 1 yyzadm sapsys    12288 Oct 28 18:44 log_backup_0_0_0_0.1382985860695
-rw-r--r-- 1 yyzadm sapsys    12288 Oct 28 18:44 log_backup_0_0_0_0.1382985864944
-rw-r--r-- 1 yyzadm sapsys    16384 Oct 28 18:44 log_backup_0_0_0_0.1382985864955
-rw-r--r-- 1 yyzadm sapsys    16384 Oct 28 18:59 log_backup_0_0_0_0.1382986752676
```

7. The next step is to push or synchronize the backup files from the /backup file system to S3 using the AWS S3 CLI.

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

```
upload: ../backup/data/YYZ/COMPLETE_DATA_BACKUP_databackup_0_1 to s3://node2-hana-s3bucket-
gcynh5v2nqs3/data/YYZ/COMPLETE_DATA_BACKUP_databackup_0_1
upload: ../backup/data/YYZ/COMPLETE_DATA_BACKUP_databackup_1_1 to s3://node2-hana-s3bucket-
gcynh5v2nqs3/data/YYZ/COMPLETE_DATA_BACKUP_databackup_1_1
upload: ../backup/data/YYZ/COMPLETE_DATA_BACKUP_databackup_2_1 to s3://node2-hana-s3bucket-
gcynh5v2nqs3/data/YYZ/COMPLETE_DATA_BACKUP_databackup_2_1
upload: ../backup/data/YYZ/COMPLETE_DATA_BACKUP_databackup_3_1 to s3://node2-hana-s3bucket-
gcynh5v2nqs3/data/YYZ/COMPLETE_DATA_BACKUP_databackup_3_1
upload: ../backup/data/YYZ/COMPLETE_DATA_BACKUP_databackup_4_1 to s3://node2-hana-s3bucket-
gcynh5v2nqs3/data/YYZ/COMPLETE_DATA_BACKUP_databackup_4_1
upload: ../backup/data/YYZ/COMPLETE_DATA_BACKUP_databackup_5_1 to s3://node2-hana-s3bucket-
gcynh5v2nqs3/data/YYZ/COMPLETE_DATA_BACKUP_databackup_5_1
upload: ../backup/data/YYZ/COMPLETE_DATA_BACKUP_databackup_6_1 to s3://node2-hana-s3bucket-
gcynh5v2nqs3/data/YYZ/COMPLETE_DATA_BACKUP_databackup_6_1
```

upload: ../backup/log/YYZ/log_backup_0_0_0_0.1382985855848 to s3://node2-hana-s3bucket-gcynh5v2nqs3/log/YYZ/log_backup_0_0_0_0.1382985855848
upload: ../backup/log/YYZ/log_backup_0_0_0_0.1382985856054 to s3://node2-hana-s3bucket-gcynh5v2nqs3/log/YYZ/log_backup_0_0_0_0.1382985856054
upload: ../backup/log/YYZ/log_backup_0_0_0_0.1382985856098 to s3://node2-hana-s3bucket-gcynh5v2nqs3/log/YYZ/log_backup_0_0_0_0.1382985856098
upload: ../backup/log/YYZ/log_backup_0_0_0_0.1382985856110 to s3://node2-hana-s3bucket-gcynh5v2nqs3/log/YYZ/log_backup_0_0_0_0.1382985856110
upload: ../backup/log/YYZ/log_backup_0_0_0_0.1382985860695 to s3://node2-hana-s3bucket-gcynh5v2nqs3/log/YYZ/log_backup_0_0_0_0.1382985860695
upload: ../backup/log/YYZ/log_backup_0_0_0_0.1382985864944 to s3://node2-hana-s3bucket-gcynh5v2nqs3/log/YYZ/log_backup_0_0_0_0.1382985864944
upload: ../backup/log/YYZ/log_backup_0_0_0_0.1382985864955 to s3://node2-hana-s3bucket-gcynh5v2nqs3/log/YYZ/log_backup_0_0_0_0.1382985864955

8. Verify the files have been pushed to S3 through the AWS Console or with the “aws s3 ls” command shown previously.

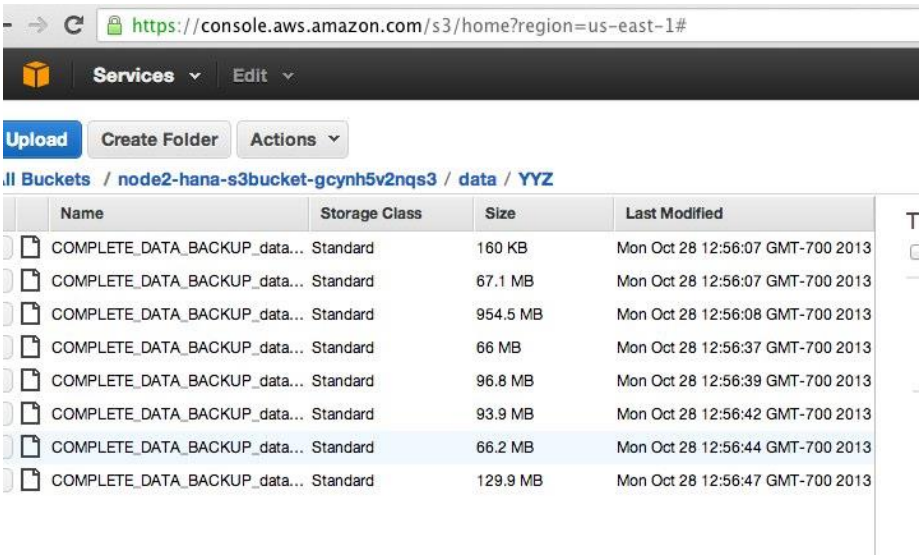

Figure 29: S3 Bucket Contents

Tip

The S3 sync command will only upload new files that don’t exist in S3. Use a periodic scheduled cron job to sync then delete files that have been uploaded. See note 1651055 for scheduling periodic backup jobs in Linux and extend the supplied scripts with the AWS S3 sync commands.

Restore Example

1. If the backup files are not readily available already in the /backup file system but are in S3, restore the files from S3 using the AWS S3 CLI command *“aws --region <region> cp <s3-bucket/path> --recursive <backup-prefix>*”*

   Example:
   imdbmaster:/backup/data/YYZ # **aws --region us-east-1 s3 cp s3://node2-hana-s3bucket-gcynh5v2nqs3/data/YYZ . --recursive --include COMPLETE***
   download: s3://node2-hana-s3bucket-gcynh5v2nqs3/data/YYZ/COMPLETE_DATA_BACKUP_databackup_0_1 to ./COMPLETE_DATA_BACKUP_databackup_0_1
   download: s3://node2-hana-s3bucket-gcynh5v2nqs3/data/YYZ/COMPLETE_DATA_BACKUP_databackup_1_1 to ./COMPLETE_DATA_BACKUP_databackup_1_1
   download: s3://node2-hana-s3bucket-gcynh5v2nqs3/data/YYZ/COMPLETE_DATA_BACKUP_databackup_2_1 to ./COMPLETE_DATA_BACKUP_databackup_2_1
   download: s3://node2-hana-s3bucket-gcynh5v2nqs3/data/YYZ/COMPLETE_DATA_BACKUP_databackup_3_1 to ./COMPLETE_DATA_BACKUP_databackup_3_1
   download: s3://node2-hana-s3bucket-gcynh5v2nqs3/data/YYZ/COMPLETE_DATA_BACKUP_databackup_4_1 to ./COMPLETE_DATA_BACKUP_databackup_4_1
   download: s3://node2-hana-s3bucket-gcynh5v2nqs3/data/YYZ/COMPLETE_DATA_BACKUP_databackup_5_1 to ./COMPLETE_DATA_BACKUP_databackup_5_1
   download: s3://node2-hana-s3bucket-gcynh5v2nqs3/data/YYZ/COMPLETE_DATA_BACKUP_databackup_6_1 to ./COMPLETE_DATA_BACKUP_databackup_6_1
   download: s3://node2-hana-s3bucket-gcynh5v2nqs3/data/YYZ/COMPLETE_DATA_BACKUP_databackup_7_1 to ./COMPLETE_DATA_BACKUP_databackup_7_1

2. Recover the SAP HANA database using the recovery wizard as outlined in the SAP HANA Administration Guide, being sure to specify file as the destination type and the correct backup prefix.

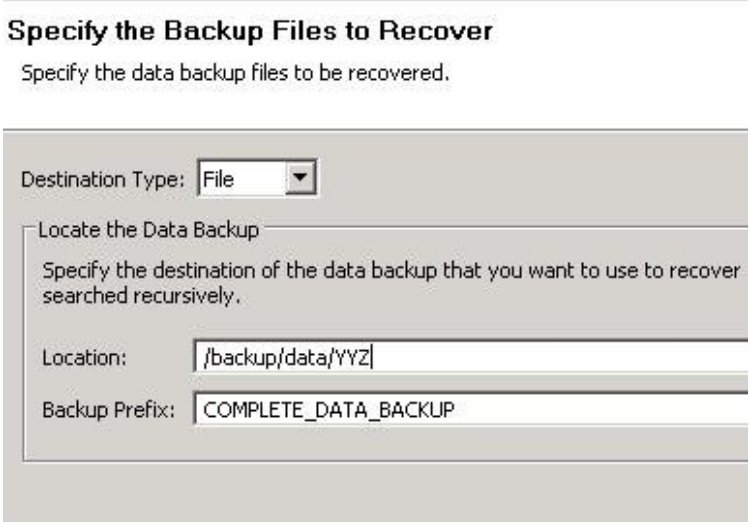

Figure 30: Restore Example

3. When the recovery is complete, resume operation and cleanup backup files from /backup/<SID>/* directories.

SAP Support Access

In some situations it may be necessary to allow an SAP support engineer to access your SAP HANA Systems on AWS. This information serves only as a supplement to the information contained in "Getting Support" section of the SAP HANA Administration guide.

There are a few steps that need to be followed in order to configure proper connectivity to SAP. These steps differ depending on whether you want to leverage an existing remote network connection to SAP or if you are setting up a new connection directly with SAP from systems on AWS.

Support Channel Setup with SAProuter on AWS

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

- Configure a specific SAProuter Security Group SAProuter instance, which only allows the required inbound and outbound access to the SAP support network. This should be limited to a specific IP address SAP gives you to connect to along with TCP port 3299.
- The instance that the SAProuter software will be installed on should be launched into a public subnet of the VPC and should be assigned an Elastic IP Address (EIP).
- Install the SAProuter software and create a saprouttab file allowing access from SAP to your SAP HANA systems on AWS.
- Setup the connection with SAP. The type of Internet connection that should be used is **Secure Network Communication (SNC)**, see https://service.sap.com/internetconnection
- Modify the existing SAP HANA security groups to trust the SAProuter Security Group.

Tip
For added security, shut down the AWS instance hosting the SAProuter service when it is not needed for support purposes.

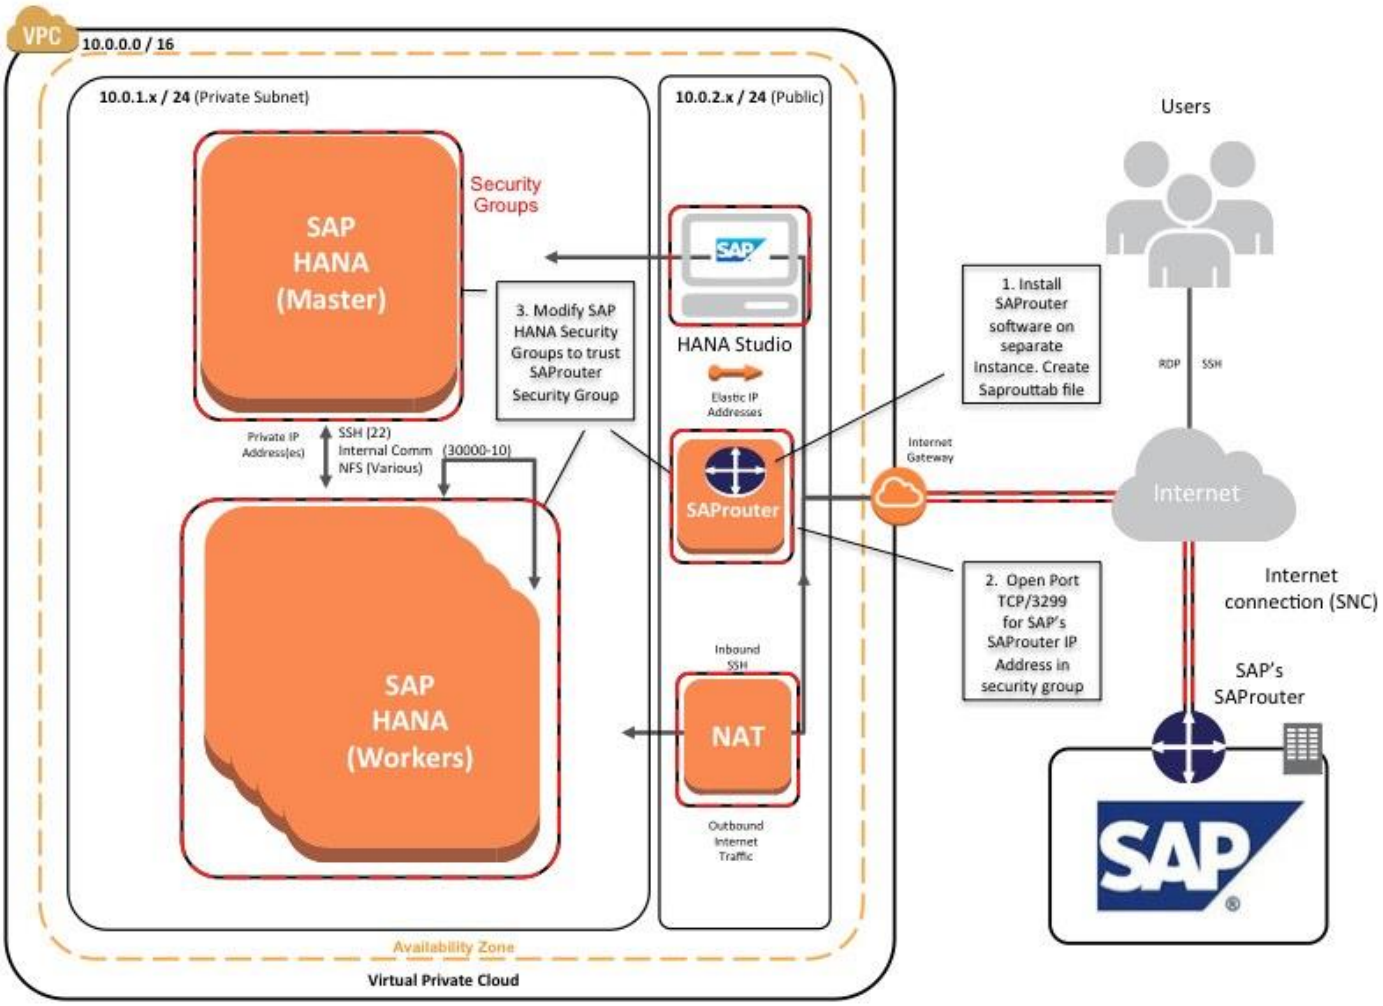

Figure 31: Support Connectivity with SAProuter on AWS

Support Channel Setup with SAProuter on-premises

In many cases a customer will already have a support connection configured between their own datacenter and SAP. This can easily be extended to allow for support of SAP systems on AWS. This scenario assumes connectivity between the customers datacenter and AWS has already been established either by way of a secure VPN tunnel over the internet or by using AWS Direct Connect.

There are a only a few steps to perform to extend this connectivity:

- Ensure the proper saprouttab entries exist to allow access from SAP to resources in the AWS VPC.
- Modify the SAP HANA Security groups to allow access from the on-premises SAProuter IP address.
- Ensure the proper firewall ports are open on the customer gateway to allow traffic to pass over TCP port 3299.

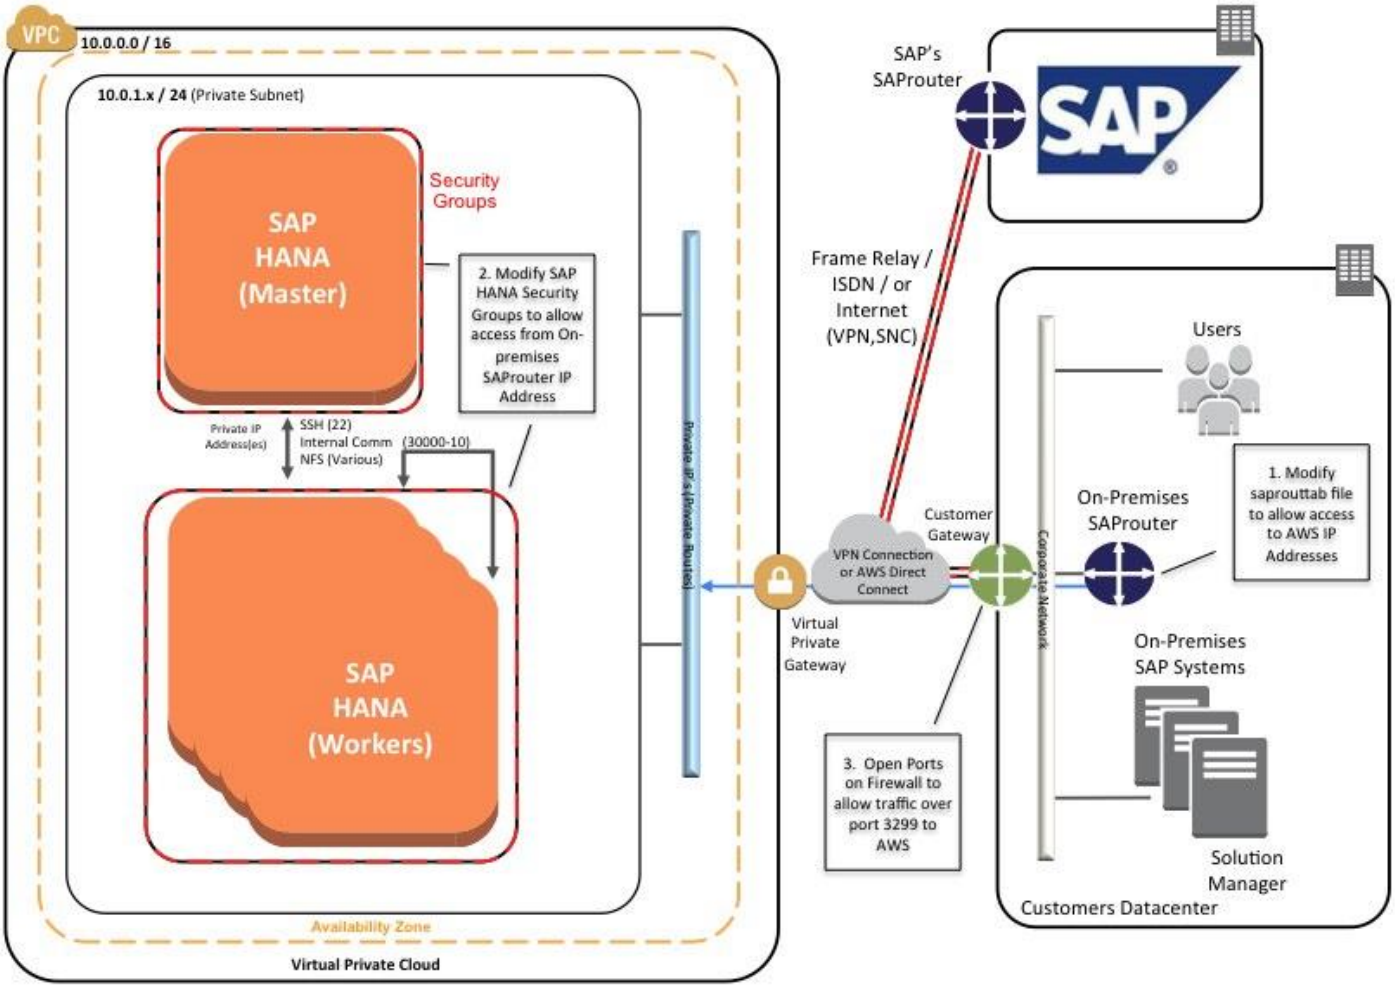

Figure 32: Support Connectivity with SAProuter On-Premises

High Availability / Disaster Recovery

This section outlines number of options for ensuring the SAP HANA system is deployed in a highly available manner. Your particular approach should only be decided after discussions with key stakeholders to understand availability requirements in terms of both recovery point and recovery time objectives (RPO/RTO).

Spare AWS Capacity

Sometimes with on-premises deployments, customers choose to purchase additional hardware to protect the SAP HANA environment in case of a hardware failure. On AWS, this may not be depending on your availability requirements. Instead of failing over to a "standby server,", you can simply start the failed virtual machine back up again and your virtual machine will be placed on a new physical host. Keep in mind, this solution is not the same as a hot standby as the SAP HANA DB will be unavailable for the time it takes to boot the virtual machine back up. However, if some downtime can be tolerated, this can save a considerable amount of money for your business.

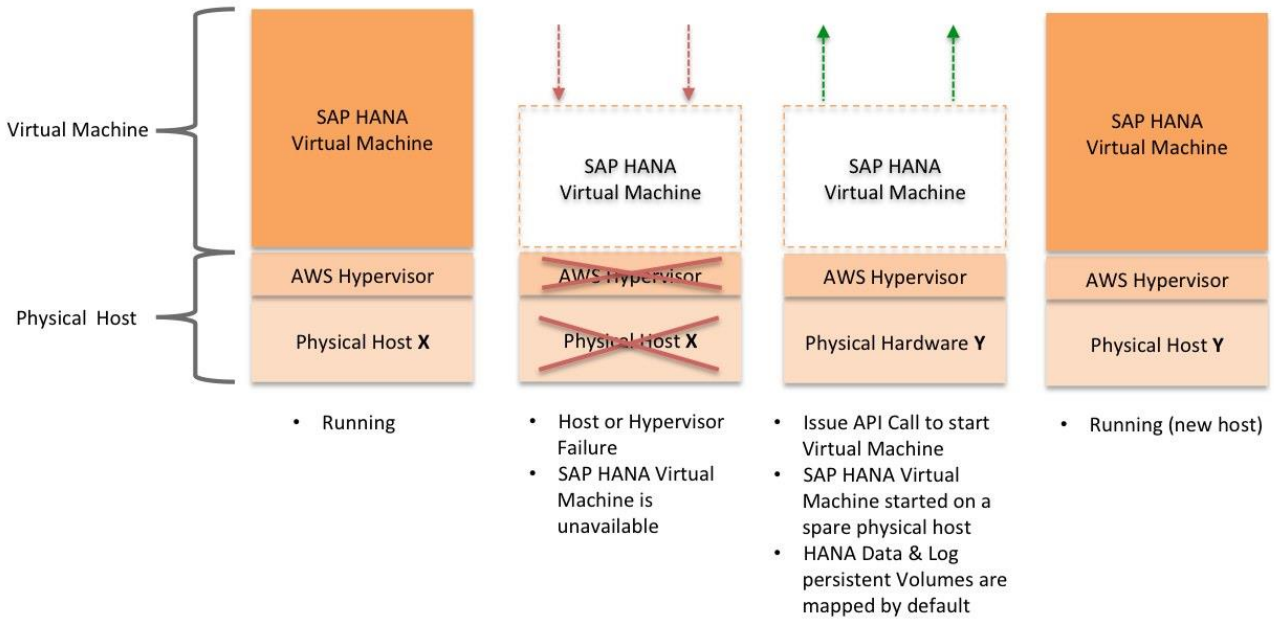

Figure 33: Leveraging Spare Capacity

Additional recommendations for this approach:

Leverage Reserved Instances
Reserved Instances can potentially provide significant cost savings depending on usage model. In addition, Reserved Instances provide a capacity reservation so that you can have confidence in your ability to launch the number of instances you have reserved when you need them.

Health Monitor
We recommend you configure a monitoring solution external to the SAP HANA System that can detect the availability of the SAP HANA Solution. Upon failure detection you can simply script appropriate actions to take based on your scenario and availability requirements.

For Example:

- Check the instance Status and availability using the AWS Command Line Interface (CLI)
  aws ec2 describe-instance-status --region <region> --instance-ids <instance-id>..<instance-id>

- If the state of the instance is stopped, just issue the start-instances command.
  aws ec2 start-instances --region us-east-1 --instance-ids <instance-id>

- If either of the status checks show as failed you may have an impaired host and should restart your instance.
  aws ec2 stop-instances --region us-east-1 --instance-ids <instance-id>

Tip
Generally it's best to allow an instance to gracefully shutdown. However, if you have issued a stop command and the instance appears to be stuck in this state you can issue the stop-instances command with the –force flag. This means the instance does not have an opportunity to flush file system caches or file system metadata. If you use this option, you must perform file system check and repair procedures.

SAP HANA High Availability using System Replication – Single Region

SAP HANA now supports system replication, which provides for a continuous update of a secondary set of HANA systems by the primary system. System replication is documented in detail in the SAP HANA Administration guide but in general system replication is configured such that the secondary systems are configured as copies of the primary systems. The number of active hosts in each system must be identical. Each SAP HANA service on the primary HANA instances communicates with its counterpart on the secondary system.

System replication can be configured for either asynchronous or synchronous replication. In synchronous mode, the primary system only commits a transaction after it has received acknowledgment from the secondary system that it has received the changes. This provides immediate consistency and provides the highest protection from data loss. While this works well for primary and secondary systems deployed in close proximity, care should be taken when system replication is configured across longer distances as this could introduce transaction delay in the system.

Within a single AWS region, Availability Zones are engineered to be isolated from failures in other Availability Zones and provide inexpensive, low latency network connectivity to the other Availability Zones in the same region. Furthermore a single VPC can be configured with separate subnets existing in different Availability Zones. These constructs then provide the ability to configure a SAP HANA environment that spans multiple datacenters to serve as a rapid failover solution for not only unplanned downtime but also planned downtime activities such as system upgrades or other maintenance activities.

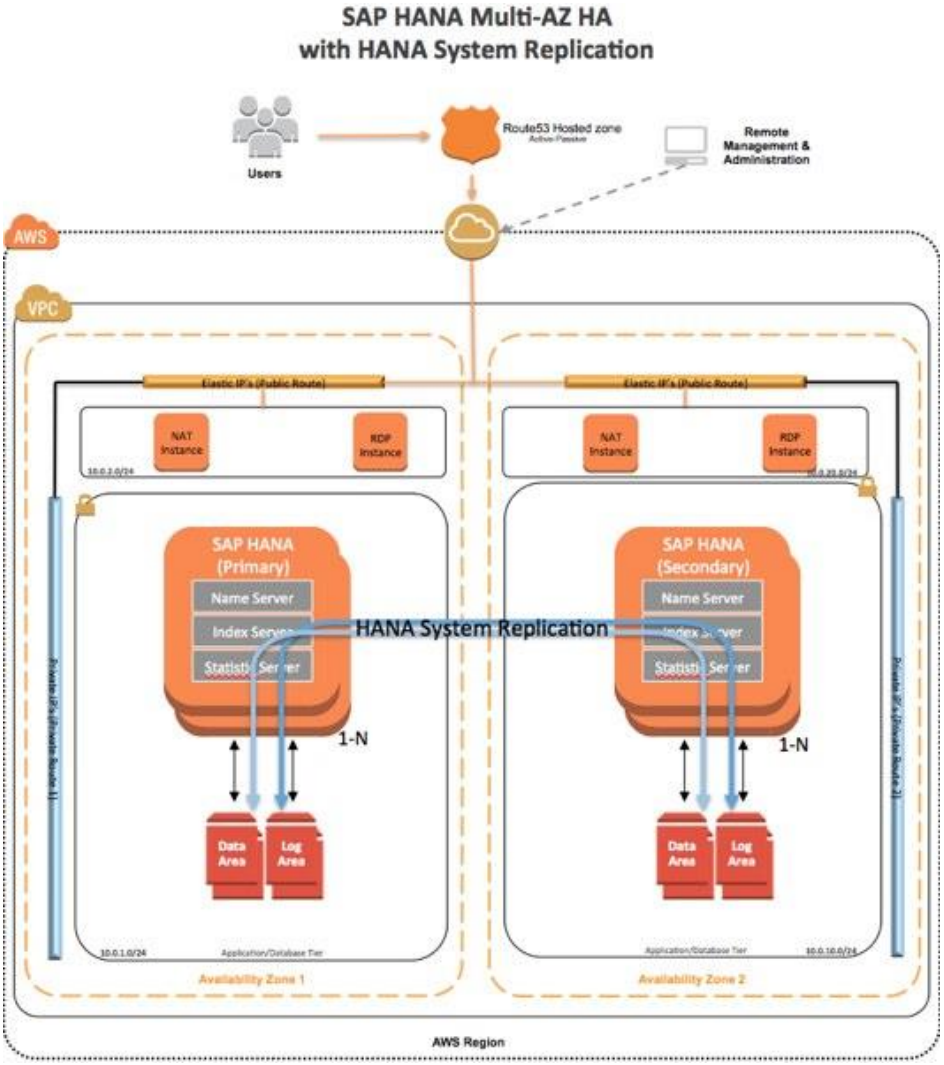

Figure 34: Multi-AZ System Replication

The process for setting up the additional systems in the secondary availability zone are as follows:

1. Create additional subnets in the VPC where SAP HANA has been deployed leveraging a second availability zone. One subnet should be private for the SAP HANA Database and or SAP Application servers and the other public if you require High Availability for the NAT and RDP instances.

   Note
   If you have connected your VPC to your own Datacenter through a secure VPN connection over the internet or via AWS Direct Connect you may not have the public subnets.

2. Associate the new subnets with the appropriate route tables in the VPC console.
3. Shutdown the primary SAP HANA Database Instance(s) and create full Amazon Machine Images (AMI's) of each instance.
4. Modify the SAP HANA Master and Worker security groups to include the new subnets to allow traffic to pass between primary and secondary HANA nodes.
5. Launch new SAP HANA systems into the new subnet leveraging the recently created AMI's.
6. Once the new systems are up and running, change the hostnames for each new HANA DB instance and update the /etc/hosts file with the proper IP Address/Hostname entry.
7. Change the hostname for the secondary SAP HANA DB Nodes using the HANA Lifecycle Manager (HLM) or command line as described in the SAP HANA Update and Configuration guide.
8. Verify that the new SAP HANA Nodes are up and running.
9. Follow the steps in section 4.1.2.1 of the SAP HANA Administration Guide to configure System Replication.
10. Test failover procedure as documented in the SAP HANA Administration Guide.

SAP HANA Disaster Recovery using System Replication – Multiple Regions

AWS also provides the ability to deploy SAP HANA environments in a multi-region deployment model. AWS Regions are dispersed and located in separate geographic areas. Currently, the BYOL version of HANA on AWS can be deployed in the following AWS regions:

- US, Northern Virginia (us-east-1)
- US, Oregon (us-west-2)
- Ireland, EU (eu-west-1)
- Tokyo, Japan (ap-northeast-1)

This method uses two separate VPC's configured in separate regions with the same number of primary and secondary SAP HANA systems. System Replication is configured using asynchronous mode. This means the primary system commits a transaction when it has been written to the log file of the primary system and sent to the secondary system through the network. It does not wait for confirmation from the secondary system. Therefore transactions are not held up on the primary system as in synchronous mode. This has potential to improve performance but also introduces the possibility of data loss upon failover if not all changes have been transferred or committed on the secondary prior to takeover.

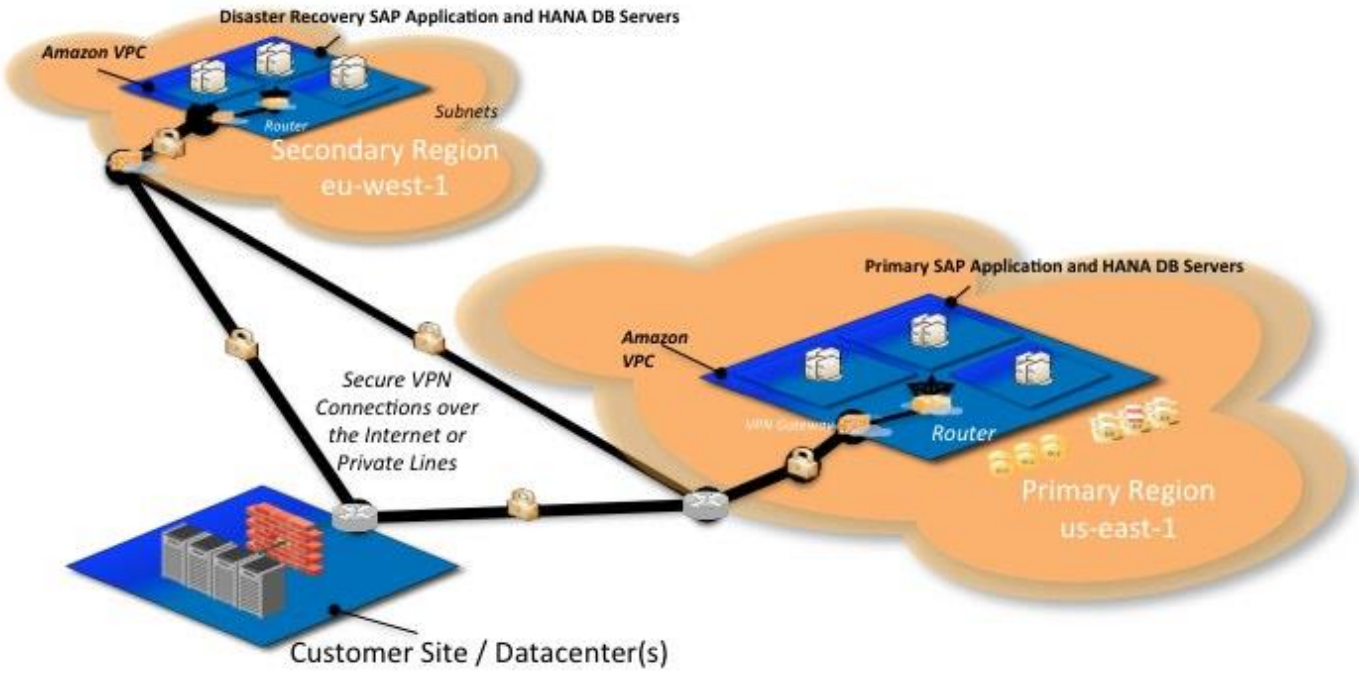

Figure 35: Multi-AZ System Replication

Note
This setup requires advanced configuration and is often influenced by custom requirements by the customer. Please contact saphana@amazon.com for additional help with this scenario.

Security

The AWS cloud infrastructure has been architected to be one of the most flexible and secure cloud computing environments available today. It provides an extremely scalable, highly reliable platform that enables customers to deploy applications and data quickly and securely.

With the AWS cloud, not only are infrastructure headaches removed, but so are many of the security issues that come with them. AWS's world-class, highly secure data centers utilize state-of-the art electronic surveillance and multi-factor access control systems. Data centers are staffed 24x7 by trained security guards, and access is authorized strictly on a least privileged basis. Environmental systems are designed to minimize the impact of disruptions to operations. And multiple geographic regions and Availability Zones allow you to remain resilient in the face of most failure modes, including natural disasters or system failures.

The AWS virtual infrastructure has been designed to provide optimum availability while ensuring complete customer privacy and segregation. For a complete list of all the security measures built into the core AWS cloud infrastructure, platforms, and services, please read our Overview of Security Processes whitepaper.

When building systems on top of the AWS infrastructure, the security responsibilities are shared between AWS and the customer. AWS secures the datacenters, infrastructure components, on up through the hypervisor layer. It is the responsibility of the customer and/or a managed service provider employed by the customer to secure the operating system, applications, and restrict access to the deployed instances from a network perspective. More information can be found at http://aws.amazon.com/security/.

Network Security

The default network security setup of this solution follows security best practices of AWS. The provisioning logic creates the solution architecture described in the solution architecture section. The provisioned SAP HANA instances can only be accessed:

1. From the CIDR block specified as "RemoteAccessCIDR" during the provisioning process.
2. By connecting to either the HANA Studio Windows Instance using Remote Desktop Client or the NAT Linux Instance using SSH.
3. Alternatively if a VPN tunnel is provisioned between the customers own data center and AWS, access can be restricted to a known CIDR block.

Identity and Access Management (IAM)

As described previously, this solution leverages an IAM role with least privileged access. It is not necessary or recommended to store SSH keys or secret keys and/or access keys on the provisioned instances.

OS Security

Access to root user on Linux or the Administrator on the Windows RDP instance can only be gained by using the SSH key specified during the deployment process. Amazon Web Services does not store these SSH keys so if you lose your SSH key you can lose access to these instances.

Operating system patches are the responsibility of the customer and should be performed on a periodic basis. The command "zypper up" will update SuSE Linux to the latest patch level available in the SuSE Linux repos on AWS.

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 created as part of this solution are restricted as much as possible while allowing access to the various functions of SAP HANA. See Appendix B for a complete list of ports and protocols configured as part of this solution.

Additional Security Options

OS Hardening

Some customers would like to lock down the OS configuration further for instance to avoid providing a DB admin with root credentials when logging into an instance.

Please also refer to SAP Notes:
1730999: Configuration changes in HANA appliance
1731000: Unrecommended configuration changes

Disabling HANA Services

HANA Services such as HANA XS are optional and should be deactivated in the case they are not needed. For instructions, see SAP Note 1697613: Remove XS Engine out of SAP HANA Database. In case of service deactivation the TCP ports should also be removed from the SAP HANA AWS Security groups for complete security.

AWS Cloud Trail

AWS Cloud Trail is a recently introduced service, which logs all AWS API calls that are made including the identity of the caller.

Notifications on Access

Notifications on SSH Login to your email address or mobile phone can be setup using AWS SNS or through 3rd party applications.

Summary

Now with AWS you don't need to wait days, weeks or even months to deploy the infrastructure needed to support your SAP HANA environment. Furthermore, AWS is completely self-service and you only pay for the resources you use. This provides a lot of flexibility for all types of SAP HANA projects and you can quickly convert these to production directly on the AWS platform.

For feedback or questions please contact us at sap-on-aws@amazon.com.

Appendix A: Custom CloudFormation Template Examples

Because the SAP HANA on AWS Infrastructure Subscription is largely based on CloudFormation, the overall solution that is deployed is largely customizable. Keep in mind that the storage configuration and instance type configurations should not be customized if "Production Support" is desired from SAP. Note you must still go through the sign-up process at saphana.com to gain access to the solution before you can use any of the custom CloudFormation templates.

Production templates:

https://s3.amazonaws.com/cf-templates-hana/SAP_HANA_AWS_1.template

https://s3.amazonaws.com/cf-templates-hana/SAP_HANA_AWS_2.template

https://s3.amazonaws.com/cf-templates-hana/SAP_HANA_AWS_3.template

https://s3.amazonaws.com/cf-templates-hana/SAP_HANA_AWS_4.template

https://s3.amazonaws.com/cf-templates-hana/SAP_HANA_AWS_5.template

Same templates with EBS Standard Volumes for non-prod and/or POCs:

https://s3.amazonaws.com/cf-templates-hana/SAP_HANA_AWS_1_EBS_Standard.template

https://s3.amazonaws.com/cf-templates-hana/SAP_HANA_AWS_2_EBS_Standard.template

https://s3.amazonaws.com/cf-templates-hana/SAP_HANA_AWS_3_EBS_Standard.template

https://s3.amazonaws.com/cf-templates-hana/SAP_HANA_AWS_4_EBS_Standard.template

https://s3.amazonaws.com/cf-templates-hana/SAP_HANA_AWS_5_EBS_Standard.template

The following CloudFormation templates provide a significant amount of customization that the default delivery including the ability to specify the following:

- Domain name – The Linux hosts are automatically configured using the domain name specified.
- Hostnames for the Linux hosts where the SAP HANA Master and Worker nodes are deployed.
- VPC-ID of existing VPC where the HANA
- Subnet-ID of existing subnet within the aforementioned VPC where SAP HANA nodes are deployed.
- Private IP Addresses of SAP HANA Virtual machines. These must be valid for the aforementioned Subnet
- Existing IAM Role to be assigned to each virtual machine (i.e. for backup functions)
- Existing Security group to be applied to each instance deployed.
- Placement group (optional)
- No RDP or NAT instance

https://s3.amazonaws.com/cf-templates-hana/SAP_HANA_AWS_1_single_subnet_existing_vpc.template

https://s3.amazonaws.com/cf-templates-hana/SAP_HANA_AWS_1_single_subnet_existing_vpc_EBS_Standard.template

https://s3.amazonaws.com/cf-templates-hana/SAP_HANA_AWS_2_single_subnet_existing_vpc.template

https://s3.amazonaws.com/cf-templates-hana/SAP_HANA_AWS_2_single_subnet_existing_vpc_EBS_Standard.template

Appendix B: Security Group Specifics

The following are the configured inbound and outbound protocols and ports allowed for the various instances deployed as part of this solution.

RDP Security Group			
Inbound			
Source	**Protocol**	**Port Range (Service)**	**Comments**
Restricted to CIDR Block specified during the deployment process	TCP	3389 (RDP)	Allow inbound RDP access to Windows instance 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

NAT Security Group			
Inbound			
Source	**Protocol**	**Port Range (Service)**	**Comments**
Restricted to CIDR Block specified during the deployment process	TCP	22 (SSH)	Allow inbound SSH access to Linux instance from your network (over the Internet gateway)
10.0.0.0/16	TCP	80 (HTTP)	Allow inbound HTTP access only from instances deployed in the VPC
10.0.0.0/16	TCP	443 (HTTPS)	Allow inbound HTTPS access from only instances deployed in the VPC
Outbound			
Destination	**Protocol**	**Port Range**	**Comments**
10.0.1.0/24	TCP	22 (SSH)	Allow SSH access from NAT instance to 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 & Worker** Security Groups			
Inbound			
Source	**Protocol**	**Port Range (Service)**	**Comments**
10.0.1.0/24	TCP	1 - 65535	Communication between instances within private subnet
10.0.1.0/24	TCP	30000 - 30010	Database Internal Communication & SAP Support Access
**10.0.1.0/24	TCP	22 (SSH)	Allow SSH access from other HANA Nodes
10.0.2.0/24	TCP	22 (SSH)	Allow SSH access from NAT instance
10.0.2.0/24	TCP	1128 - 1129	Host Agent Access
10.0.2.0/24	TCP	4300	Access to XSEngine (HTTPS) from 10.0.2.0 subnet
10.0.2.0/24	TCP	8000	Access to XSEngine (HTTP) from 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	30015	DB Client Access
10.0.2.0/24	TCP	30017	DB Client Access
10.0.2.0/24	TCP	50013 - 50014	Allow Access for HANA Studio from RDP Instance
Outbound			
0.0.0.0/0	TCP	1 - 65535	Outbound access from HANA Master allowed to anywhere