Setting up AWS Resources and the SLES Operating System for SAP HANA Installation

Amazon Web Services

March 2015

Contents

Abstract

Amazon Web Services (AWS) provides various services and tools for deploying SAP products on the AWS cloud platform. This guide discusses the steps required to set up and configure AWS resources such as Amazon Elastic Compute Cloud (Amazon EC2) instances and Amazon Elastic Block Store (Amazon EBS) volumes to install a new SAP HANA database instance or to scale out your existing SAP HANA database instance. The guide assumes familiarity with AWS services and resources such as Amazon Virtual Private Cloud (Amazon VPC), Amazon EC2, Amazon EBS, and security groups.

Introduction

AWS provides a Quick Start reference deployment to fast-track your SAP HANA deployment on the AWS cloud. The Quick Start leverages the AWS programmable infrastructure to quickly provision resources needed to deploy SAP HANA, and usually takes less than an hour to complete with minimal manual intervention.

Enterprises that are familiar with deploying SAP HANA on AWS may want to expand their SAP HANA footprint on AWS by either deploying additional SAP HANA instances or scaling out existing SAP HANA instances in their AWS cloud infrastructure. While the Quick Start deployment provides an easy hands-off way to deploy new SAP HANA instances on AWS, this guide describes how to set up AWS resources such as Amazon EC2 and Amazon EBS by using the AWS Command Line Interface (AWS CLI) to support additional customization scenarios. This guide also explains how to configure the SUSE Linux Enterprise Server (SLES) operating system for new SAP HANA instance deployments, or to scale out your existing SAP HANA instances on the AWS cloud.

Before You Get Started

You are responsible for all costs related to your use of any AWS services while following this guide. Prices are subject to change. Please refer to the pricing pages for the AWS services that you intend to use for full details.

You must already own a license for the SAP HANA software and have access to download the SAP HANA Platform Edition software from the SAP Software Download Center (requires access to SAP Service Marketplace).

Creating AWS resources for SAP HANA installation is an advanced topic. If you are new to AWS, see the following content to get familiar with AWS technologies:

- Getting started with AWS
- Amazon EC2
- Amazon EBS
- Amazon VPC
- AWS CLI (Command Line Interface)

For best practices for implementing SAP HANA on AWS, see the following:

- SAP HANA on AWS Implementation and Operations Guide
- SAP HANA Quick Start Deployment Guide (see the section “Overview of SAP HANA on AWS”)

What We'll Cover

We will cover two scenarios in this guide.

- Scenario 1: Creating AWS resources and configuring SLES for a single-node SAP HANA instance.
- Scenario 2: Creating AWS resources and configuring SLES to scale out an existing SAP HANA instance.

Choose the scenario you'd like to implement. We will cover the detailed steps for implementing each scenario with AWS CLI and operating system commands. By the end of this document, you will have the AWS infrastructure ready to install a multi-node SAP HANA instance.

This guide does not cover the setup of Amazon VPC, subnets, and security groups, or the installation of SAP HANA.

Preparation

Before you get started, complete the following preparatory tasks:

1. Make sure that the following information pertaining to your existing AWS resources is available. You will need this information while executing AWS CLI commands to create your Amazon EC2 and Amazon EBS resources.

Information You'll Need	Description
Region ID	Region where you want to deploy your AWS resources.
Availability Zone	Availability Zone within your target region where you want to deploy your resources.
Placement group	Placement group you want use to logically group your Amazon EC2 instances within an Availability Zone (recommended for scale out SAP HANA implementations).
Amazon VPC ID	Amazon VPC where you want to deploy your Amazon EC2 instance for SAP HANA installation.
Subnet ID	Subnet where you want to deploy your Amazon EC2 instance.
AMI ID	Amazon Machine Image (AMI) that will be used to launch your Amazon EC2 instance. You can find the latest 64-bit AMI by looking for the string "suse-sles-11-sp3-v*YYYYMMDD*-hvm-ssd-x86_64" (for example,

Information You'll Need	Description
	"suse-sles-11-sp3-v20141003-hvm-ssd-x86_64") in the AMI name in the AWS Management Console, and choosing the AMI name that specifies the latest date.
Key pair	Make sure that you have generated the key pair in your target region, and that you have access to the private key.
Security group ID	Name of the security group that you want to assign to your Amazon EC2 instance. See the appendix for detailed information about the security group for SAP HANA instances.
Access key ID	Access key for your AWS account that will be used with AWS CLI tools.
Secret access key	Secret key for your AWS account that will be used with AWS CLI tools.
SAP HANA media volume ID	Name of the volume where you have staged your HANA media.

2. Install the AWS Command Line Interface (AWS CLI) on the desktop, laptop, or server from which you plan to execute the commands to deploy AWS resources. If you plan to set up AWS CLI under Linux, OS X, or Unix, we recommend that you use the bundled installer. For instructions, see the AWS CLI User Guide.

> **Tip** Configure your AWS CLI profile with the target region where you want to deploy your AWS resources for SAP HANA.

Scenario 1: Creating AWS Resources and Configuring SLES for a Single-Node HANA Instance

> **Note** The syntax shown for the CLI and Linux commands in this section is specific to the scope of this document. Each command supports many additional options. For more information, use the help option on the command line or refer to the documentation.

1. **Check the region where you want to deploy your AWS resources.**

   Display the AWS CLI configuration data:

   ```
   $ aws configure list
   ```

   In the command output, make sure that the default region configured in your AWS profile is the same as the target region where you want to deploy your AWS resources.

2. **Create a JSON file for the Amazon EBS storage volumes.**

   Create a JSON file that contains the storage requirements for the SAP HANA master node. You will use this JSON file in the next step, when you create the Amazon EC2 instance.

   Depending on the Amazon EC2 instance type that you choose, here are the storage sizes we recommend for optimal performance:

Amazon EC2 Instance Type	RAM (GiB)	General-Purpose SSD (GiB)	Total SAP HANA Volume Size (GiB)
c3.8xlarge	60	4 x 334* 1 x 50**	1336
r3.2xlarge	61	4 x 334* 1 x 50**	1336
r3.4xlarge	122	4 x 334* 1 x 50**	1336
r3.8xlarge	244	4 x 667* 1 x 50**	2668

* These volumes will be used to store the SAP HANA data, log, backup, and shared files.

** This volume will be used to store local SAP files.

Once you determine the Amazon EC2 instance type that you want to use for your SAP HANA instance, use the vi editor to create a JSON file (e.g., ebs.json) similar to the following, and save the file in the /tmp folder. Replace the value for the *VolumeSize* parameter depending on your instance type.

Example JSON file:

```
 [
{"DeviceName":"/dev/sdf","Ebs":{"VolumeSize":667,"VolumeTyp
e":"gp2","DeleteOnTermination":true}},
{"DeviceName":"/dev/sdg","Ebs":{"VolumeSize":667,"VolumeTyp
e":"gp2","DeleteOnTermination":true}},
{"DeviceName":"/dev/sdh","Ebs":{"VolumeSize":667,"VolumeTyp
e":"gp2","DeleteOnTermination":true}},
{"DeviceName":"/dev/sdi","Ebs":{"VolumeSize":667,"VolumeTyp
e":"gp2","DeleteOnTermination":true}},
{"DeviceName":"/dev/sdj","Ebs":{"VolumeSize":50,"VolumeType
":"gp2","DeleteOnTermination":true}}]
```

Note The device name `/dev/sda` is reserved for the root volume. The recommended device names for Amazon EBS volumes are `/dev/sd[f-p]`. For more information, see *Block Device Mapping* in the AWS documentation.

3. Launch the Amazon EC2 instance.

Launch the Amazon EC2 instance for the SAP HANA master node in your target region by using the information you gathered in the preparation phase. You will also be creating the storage volumes required for the SAP HANA master node and attaching them to the Amazon EC2 instance for the SAP HANA master node, per the JSON file created in the previous step.

Make sure that you choose one of these Amazon EC2 instance types for your SAP HANA instance:

AWS Instance for SAP HANA	vCPU	RAM (GiB)	SAP Support
r3.2xlarge	8	61	Non-production use only
c3.8xlarge	16	60	Non-production and production use for SAP Business One, version for SAP HANA

AWS Instance for SAP HANA	vCPU	RAM (GiB)	SAP Support
r3.4xlarge	16	122	Non-production use only
r3.8xlarge	32	244	Non-production and production use for SAP HANA Platform Edition and SAP Business One, version for SAP HANA

Note Make sure that you have the private key of the key pair you used to launch your instance. You cannot access your instance without the private key.

Syntax:

$ **aws ec2 run-instances**
--**image-id** *AMI-ID*
--**count** *number-of-EC2-instances*
--**instance-type** *instance-type*
--**key-name**=*name-of-key-pair*
--**security-group-ids** *security-group-ID*
--**subnet-id** *subnet-ID*
--placement AvailabilityZone=*Availability-Zone-ID*,
GroupName=*placement-group-name*
--**block-device-mappings file://***path-of-the-JSON file*

Example (when using the command, please make sure to place the command and its parameters on a single line):

```
$ aws ec2 run-instances
--image-id ami-xxxxxxxx
--count 1
--instance-type r3.8xlarge
--key-name=my_key
--security-group-ids sg-xxxxxxxx
--subnet-id subnet-xxxxxxxx
--placement AvailabilityZone=us-east-1a,GroupName=My-
PlacementGroup
--block-device-mappings file:///tmp/ebs.json
```

To monitor the status of your instance, use the **describe-instance-status** command.

Syntax:

$ aws ec2 describe-instance-status --instance-id *instance-ID*

Example:

```
$ aws ec2 describe-instance-status --instance-id i-xxxxxxxx
```

Optionally, you can create tags for the instance that you just launched.

Syntax:

$ aws ec2 create-tags
--resources *resource-ID*
--tags Key=*key-name*,**Value**=*key-value*

Example:

```
$ aws ec2 create-tags --resources i-xxxxxxxx --tags
Key=Name,Value="SAP HANA Master"
```

4. **Update the host name.**

Log in to your SAP HANA worker node with SSH using the private key pair, and switch to root user to update the host name along with the domain name according to your requirements. Use the following Linux commands.

Syntax: **hostname** *name-of-your-host*

Example:

```
# hostname imdbmaster
```

This change will be valid only until the system is rebooted. Add the host name with the domain name to the /etc/HOSTNAME file to make the change permanent.

Syntax: **echo** *hostname.domain* **> /etc/HOSTNAME**

Example:

```
# echo imdbmaster.local > /etc/HOSTNAME
```

Add an entry to the /etc/hosts file with your new host name and the IP address of your instance.

Syntax: **echo "***IP-address hostname.domain hostname***" >> /etc/hosts**

Example:

```
# cp /etc/hosts /etc/hosts.bak
# echo "10.0.1.2 imdbmaster imdbmaster.local" >> /etc/hosts
```

Alternatively, you can edit the /etc/hosts file and manually add this entry.

Tip For SAP systems, the maximum length of the host name should not exceed 13 characters. The name should comply with SAP standards. See SAP OSS Note 611361 for details (requires access to SAP Service Marketplace).

In the **/etc/cloud/cloud.cfg** file, change the *preserve_hostname* value to **true**. This will ensure that the Amazon EC2 instance's host name is preserved during a restart.

Finally, update the network settings using the YaST tool to disable host name reset via DHCP:

a. Start YaST.

b. Navigate to Network Devices, Network Settings (press Enter).

c. Select Hostname/DNS in the menu.

d. Change the values of the Hostname and Domain Name fields if required.

e. Clear the Change Hostname via DHCP check box.

f. Press F10 to save.

g. Press F9 to exit.

5. **Install prerequisite packages**.

   As root user, use the following **zypper** commands to install the SLES packages required for SAP HANA installation, and to remove the **ulimit** package.

   Syntax:

   To install a package: **zypper –n install** *package-name*

   To remove a package: **zypper remove** *package-name*

> **Note** Your Amazon EC2 instance should have access to the Internet to read and download required packages from the SUSE repository.

```
zypper -n install gtk2
zypper -n install java-1_6_0-ibm
zypper -n install libicu
zypper -n install mozilla-xulrunner*
zypper -n install ntp
zypper -n install sudo
zypper -n install syslog-ng
zypper -n install tcsh libssh2-1
zypper -n install autoyast2-installation
zypper -n install yast2-ncurses
zypper -n install libgcc_s1 libstdc++6
zypper -n install ntfs-3g
zypper -n install nfs-kernel-server

zypper remove ulimit
```

6. **Update the NFS server ports.**

   As root user, update the NFS server configuration file to use specific ports. Your SAP HANA master node will be the NFS server that shares the /hana/shared and /backup file systems when you decide to scale out your SAP HANA instance.

   Use an editor to update the following ports in the /etc/sysconfig/nfs file:

```
STATD_PORT="4000"
LOCKD_TCPPORT="4001"
LOCKD_UDPPORT="4001"
```

```
MOUNTD_PORT="4002"
```

7. Configure the NTP, AUTOFS, NFS, and LVM services.

As root user, add the following entries to the /etc/ntp.conf file to configure the server pool for the NTP service. Before you make the changes, we recommend that you back up the existing ntp.conf file.

```
# cp /etc/ntp.conf /etc/ntp.conf.bak
# echo "server 0.pool.ntp.org" >> /etc/ntp.conf
# echo "server 1.pool.ntp.org" >> /etc/ntp.conf
# echo "server 2.pool.ntp.org" >> /etc/ntp.conf
# echo "server 3.pool.ntp.org" >> /etc/ntp.conf
```

Now, configure the ntp, autofs, boot.lvm, and nfsserver services to start automatically during an instance reboot.

Syntax:

To enable autostart for a service: **chkconfig** *service-name* **on**
To restart a service: **service** *service-name* **restart**

```
# chkconfig ntp on
# chkconfig autofs on
# chkconfig boot.lvm on
# chkconfig nfsserver on
# service ntp restart
# service autofs restart
# service nfsserver restart
```

8. Disable HugePages.

As root user, disable HugePages for your SAP HANA instance, and add an entry to the /etc/init.d/boot.local file to keep HugePages disabled during instance reboot. We recommend that you back up the existing /etc/init.d/boot.local file before you make this change.

```
# cp /etc/init.d/boot.local /etc/init.d/boot.local.bak
# echo never > /sys/kernel/mm/transparent_hugepage/enabled
# echo "echo never >
/sys/kernel/mm/transparent_hugepage/enabled" >>
/etc/init.d/boot.local
```

9. Add and update the kernel parameter.

As root user, add a new kernel parameter and load the kernel settings to activate the change dynamically. We recommend that you back up the /etc/sysctl.conf file before you make this change.

```
# cp /etc/sysctl.conf /etc/sysctl.conf.bak
# echo "net.ipv4.tcp_slow_start_after_idle=0" >>
/etc/sysctl.conf
# sysctl -p /etc/sysctl.conf
```

10. Create physical volume(s).

As root user, create physical volume(s). You need to create physical devices only for volumes that will be used to store SAP HANA data, log, backup, and shared files. (Typically, these volumes are 334 GiB or 667 GiB in size.)

To view the list of volumes attached to your instance and their device names, run the **lsblk** command as root user. The command displays the list of devices that are attached to your instance.

Syntax: **lsblk**

Example:

```
# lsblk
NAME   MAJ:MIN RM   SIZE RO MOUNTPOINT
hda      3:0    0    10G  0
└─hda1   3:1    0    10G  0 /
xvdj   202:144  0    50G  0
xvdi   202:128  0   667G  0
xvdh   202:112  0   667G  0
xvdg   202:96   0   667G  0
xvdf   202:80   0   667G  0
```

```
$
```

Note The device might be attached with a different name than what you specify. For example, if you specify a device name of /dev/sdf, the kernel might rename your device /dev/xvdf, but in most cases, the trailing letter will remain the same.

Create physical devices by running the **pvcreate** command.

Syntax: **pvcreate** *device-name ... device-name*

Example:

```
# pvcreate /dev/xvdf /dev/xvdg /dev/xvdh /dev/xvdi
```

11. Set the I/O scheduling option to NOOP.

As root user, set the I/O scheduling option to "noop" mode to optimize the I/O performance of your SAP HANA instance on the AWS platform.

Syntax: **echo "noop" > /sys/block/***device-name***/queue/scheduler**

Example:

```
# echo "noop" > /sys/block/xvdf/queue/scheduler
# echo "noop" > /sys/block/xvdg/queue/scheduler
# echo "noop" > /sys/block/xvdh/queue/scheduler
# echo "noop" > /sys/block/xvdi/queue/scheduler
```

Important Make sure that you set the I/O scheduling option to "noop" for **all** the physical devices that you have created.

12. Create a volume group.

As root user, create a volume group called **vghana** to stripe the multiple physical volumes you created in step 10 for optimal I/O performance.

Syntax: **vgcreate** *vg-name device-name(s)*

Example:

```
# vgcreate vghana /dev/xvdf /dev/xvdg /dev/xvdh /dev/xvdi
```

13. Create logical volume(s).

As root user, create logical volumes for SAP HANA log, data, shared, and backup files. The size of the logical volumes will differ depending on the Amazon EC2 instance type you choose, as shown in the following table.

Amazon EC2 Instance Type	Log Size (GiB)	Data Size (GiB)	Shared Data Size (GiB)	Backup Size (GiB)
c3.8xlarge	60	180	60	300
r3.2xlarge	60	180	60	300
r3.4xlarge	122	366	122	610
r3.8xlarge	244	732	244	1200

Syntax: **lvcreate –n** *lv-name* **-i** *#-of-volumes-to-stripe-across* **-I 256 –L** *lv-size vg name*

Example:

```
# lvcreate -n lvhanalog -i 4 -I 256 -L 244G vghana
# lvcreate -n lvhanadata -i 4 -I 256 -L 732G vghana
# lvcreate -n lvhanashared -i 4 -I 256 -L 244G vghana
# lvcreate -n lvhanaback -i 4 -I 256 -L 1200G vghana
```

14. Format logical volume(s).

As root user, format all logical volumes you created in the previous step. You can find the list of logical volumes in the directory /dev/mapper.

Syntax: **mkfs.xfs** *full-path-of-logical-volume*

Example:

```
# mkfs.xfs /dev/mapper/vghana-lvhanalog
# mkfs.xfs /dev/mapper/vghana-lvhanadata
# mkfs.xfs /dev/mapper/vghana-lvhanashared
# mkfs.xfs /dev/mapper/vghana-lvhanaback
```

15. Format block device for /usr/sap.

As root user, format the Amazon EBS volumes attached to your instance to store local SAP files (device name: /dev/xvdj).

Syntax: **mkfs.xfs -f** *device-name*

Example:

```
# mkfs.xfs -f /dev/xvdj
```

16. Create directories to mount the file systems.

As root user, create the directories to mount the file systems required for SAP HANA master node installation. You will need the following mount points:

- /hana/log
- /hana/data
- /hana/shared
- /usr/sap
- /backup

Syntax: **mkdir** *directory-path*

Example:

```
# mkdir /hana /hana/log /hana/data /hana/shared /backup
/usr/sap
```

17. Create entries for mount points in the /etc/fstab file.

As root user, add entries to the /etc/fstab file and mount the file systems. Adding entries to /etc/fstab ensures that your file systems are mounted automatically when your Amazon EC2 instance is restarted.

Add the entries for the SAP HANA data, log, shared, backup, and local SAP file systems to the /etc/fstab file by using the following commands:

Syntax: **echo** *"logical-volume-path mount-directory file-system-type mount-options" >> /etc/fstab*

Example:

```
# echo "/dev/xvdj                    /usr/sap       xfs
nobarrier,noatime,nodiratime,logbsize=256k,delaylog 0 0" >>
/etc/fstab
# echo "/dev/mapper/vghana-lvhanadata     /hana/data
xfs nobarrier,noatime,nodiratime,logbsize=256k,delaylog 0
0" >> /etc/fstab
# echo "/dev/mapper/vghana-lvhanalog      /hana/log
xfs nobarrier,noatime,nodiratime,logbsize=256k,delaylog 0
0" >> /etc/fstab
# echo "/dev/mapper/vghana-lvhanashared      /hana/shared
xfs nobarrier,noatime,nodiratime,logbsize=256k,delaylog 0
0" >> /etc/fstab
# echo "/dev/mapper/vghana-lvhanaback      /backup      xfs
nobarrier,noatime,nodiratime,logbsize=256k,delaylog 0 0" >>
/etc/fstab
```

Alternatively, you can edit the /etc/fstab file and add these entries.

When you have added the entries to the /etc/fstab file, mount the file systems, and then run the **df** command to verify that they were mounted appropriately.

Syntax: **mount –a**
df -h

Example:

```
# mount -a
# df -h
Filesystem                    Size  Used Avail Use%
Mounted on
rootfs                        9.9G  1.8G  7.6G  20% /
udev                          121G  116K  121G   1% /dev
tmpfs                         121G     0  121G   0%
/dev/shm
/dev/hda1                     9.9G  1.8G  7.6G  20% /
/dev/xvdj                      50G   33M   50G   1%
/usr/sap
```

```
/dev/mapper/vghana-lvhanashared  244G  8.1G  236G   4%
/hana/shared
/dev/mapper/vghana-lvhanadata    732G  2.6G  730G   1%
/hana/data
/dev/mapper/vghana-lvhanalog     244G  4.2G  240G   2%
/hana/log
/dev/mapper/vghana-lvhanaback    1.2T   34M  1.2T   1%
/backup
#
```

18. Attach and mount the HANA media volume for installation.

Use AWS CLI to attach the HANA media volume to your SAP HANA instance, and mount it to a file system for installation.

Syntax: **aws ec2 attach-volume --volume-id** *volume-name* **--instance** *instance-name* **--device** *device-name*

Example:

```
$ aws ec2 attach-volume --volume-id vol-xxxxxxxx --instance
i-yyyyyyyy --device /dev/sdk
```

Note The device might be attached with a different name than what you specify. For example, if you specify a device name of /dev/sdf, the kernel might rename your device /dev/xvdf, but in most cases, the trailing letter will remain the same.

Now mount your HANA media volume.

Syntax: **mount** *device-name mount-directory*

```
# mkdir /media
# mount /dev/xvdk1 /media
```

Tip To have /media mounted during every system reboot, add an entry to the /etc/fstab file. We recommend that you back up your

existing /etc/fstab file before you make this change. See Step 17 for instructions.

19. **Install SAP HANA.**

    Follow the standard SAP installation procedure to install your SAP HANA database instance.

Scenario 2: Creating AWS Resources and Configuring SLES to Scale out an Existing HANA Instance

Before you start:

- Make sure that you are creating your SAP HANA worker node instance(s) in the same region, Availability Zone, subnet, and placement group as your master node. Also, the Amazon EC2 instance type that you choose for your worker node should be same as the instance type for the master node.
- Make sure that SAP supports the SAP HANA scale-out scenario for the SAP solution you are implementing, such as SAP NetWeaver Business Warehouse (BW), or SAP Business Suite.

Note The syntax shown for the CLI and Linux commands in this section is specific to the scope of this document. Each command supports many additional options. For more information, use the help option on the command line or refer to the documentation.

1. **Check the region where you want to deploy your AWS resources.**

   Display the AWS CLI configuration data:

   Syntax: **$ aws configure list**

   In the command output, make sure that the default region configured in your AWS profile is the same as the target region where you want to deploy your AWS resources.

2. **Create a JSON file for the Amazon EBS storage volumes.**

   Create a JSON file that contains the storage requirements for the SAP HANA worker node. You will use this JSON file in the next step when you create the Amazon EC2 instance.

Depending on the Amazon EC2 instance type that you choose, here are the storage sizes we recommend for optimal performance:

Amazon EC2 Instance Type	RAM (GiB)	General-Purpose SSD (GiB)	Total SAP HANA Volume Size (GiB)
c3.8xlarge	60	4 x 334* 1 x 50**	1336
r3.2xlarge	61	4 x 334* 1 x 50**	1336
r3.4xlarge	122	4 x 334* 1 x 50**	1336
r3.8xlarge	244	4 x 667* 1 x 50**	2668

* These volumes will be used to store the SAP HANA data and log files.

** This volume will be used to store local SAP files.

Once you determine the Amazon EC2 instance type that you want to use for your SAP HANA instance, use the vi editor to create a JSON file (e.g., ebs.json) similar to the following, and save the file in the /tmp folder. Replace the value for the *VolumeSize* parameter depending on your instance type.

Example JSON file:

```
 [
{"DeviceName":"/dev/sdf","Ebs":{"VolumeSize":667,"VolumeTyp
e":"gp2","DeleteOnTermination":true}},
{"DeviceName":"/dev/sdg","Ebs":{"VolumeSize":667,"VolumeTyp
e":"gp2","DeleteOnTermination":true}},
{"DeviceName":"/dev/sdh","Ebs":{"VolumeSize":667,"VolumeTyp
e":"gp2","DeleteOnTermination":true}},
{"DeviceName":"/dev/sdi","Ebs":{"VolumeSize":667,"VolumeTyp
e":"gp2","DeleteOnTermination":true}},
{"DeviceName":"/dev/sdj","Ebs":{"VolumeSize":50,"VolumeType
":"gp2","DeleteOnTermination":true}}
]
```

Note The device name /dev/sda is reserved for the root volume. The recommended device names for Amazon EBS volumes are /dev/sd[f-p].

For more information, see *Block Device Mapping* in the AWS documentation.

3. Launch the Amazon EC2 instance.

Launch the Amazon EC2 instance(s) for the SAP HANA worker node in your target region by using the information you gathered in the preparation phase. You will be also be creating the storage volumes required for SAP HANA worker node(s) and attaching them to the Amazon EC2 instance for the SAP HANA worker node, per the JSON file created in the previous step.

Make sure that you choose one of these Amazon EC2 instance types for your SAP HANA instance:

AWS Instance for SAP HANA	vCPU	RAM (GiB)	SAP Support
r3.2xlarge	8	61	Non-production use only
c3.8xlarge	16	60	Non-production and production use for SAP Business One, version for SAP HANA
r3.4xlarge	16	122	Non-production use only
r3.8xlarge	32	244	Non-production and production use for SAP HANA Platform Edition and SAP Business One, version for SAP HANA

Note Make sure that you have the private key of the key pair you used to launch your instance. You cannot access your instance without the private key.

Syntax:

$ **aws ec2 run-instances**
--**image-id** *AMI-ID*
--**count** *number-of-EC2-instances*
--**instance-type** *instance-type*
--**key-name**=*name-of-key-pair*
--**security-group-ids** *security-group-ID*
--**subnet-id** *subnet-ID*
--placement AvailabilityZone=*Availability-Zone-ID*,
GroupName=*placement-group-name*
--**block-device-mappings file://***path-of-the-JSON file*

Example (when using the command, please make sure to place the command and its parameters on a single line):

```
aws ec2 run-instances
--image-id ami-xxxxxxxx
--count 1
--instance-type r3.8xlarge
--key-name=my_key
--security-group-ids sg-xxxxxxxx
--subnet-id subnet-xxxxxxxx
--placement AvailabilityZone=us-east-1a,GroupName=My-
PlacementGroup
--block-device-mappings file:///tmp/ebs.json
```

Tip If you are planning to add multiple worker nodes, use the **–count** option to specify the number of nodes you need. This will help launch all your worker nodes in parallel.

To monitor the status of your instance, use the **describe-instance-status** command.

Syntax:

$ aws ec2 describe-instance-status --instance-id *instance-ID*

Example:

```
$ aws ec2 describe-instance-status --instance-id i-xxxxxxxx
```

Optionally, you can create tags for the instance that you just launched.

Syntax:

$ aws ec2 create-tags
--resources *resource-ID*
--tags Key=*key-name*,**Value**=*key-value*

Example:

```
$ aws ec2 create-tags --resources i-xxxxxxxx --tags
Key=Name,Value="SAP HANA Master"
```

Important You need to repeat steps 4 through the end of the procedure for every worker node.

4. **Update the host name.**

   Log in to your SAP HANA worker node with SSH using the private key pair, and switch to root user to update the host name along with the domain name according to your requirements. Use the following Linux commands.

   Syntax: **hostname** *name-of-your-host*

   Example:

```
# hostname imdbworker01
```

   This change will be valid only until the system is rebooted. Add the hostname with the domain name to the /etc/HOSTNAME file to make the change permanent.

   Syntax: **echo** *hostname.domain* **> /etc/HOSTNAME**

   Example:

```
# echo imdbworker01.local > /etc/HOSTNAME
```

   Add an entry to the /etc/hosts file with your new host name and the IP address of your instance.

   Syntax: **echo "***IP-address hostname.domain hostname***" >> /etc/hosts**

   Example:

```
# cp /etc/hosts /etc/hosts.bak
```

```
# echo "10.0.1.3 imdbworker01 imdbworker01.local" >>
/etc/hosts
```

Alternatively, you can edit the /etc/hosts file and manually add this entry.

> **Tip** For SAP systems, the maximum length of the host name should not exceed 13 characters. The name should comply with SAP standards. See SAP OSS Note 611361 for details (requires access to SAP Service Marketplace).

In the **/etc/cloud/cloud.cfg** file, change the *preserve_hostname* value to **true**. This will ensure that the Amazon EC2 instance's host name is preserved during a restart.

Finally, update the network settings using the YaST tool to disable host name reset via DHCP.

a. Start YaST.

b. Navigate to Network Devices, Network Settings (press Enter).

c. Select Hostname/DNS in the menu.

d. Change the values of the Hostname and Domain Name fields if required.

e. Clear the Change Hostname via DHCP check box.

f. Press F10 to save.

g. Press F9 to exit.

5. Install prerequisite packages.

As root user, use the following **zypper** commands to install the SLES packages required for SAP HANA installation and to remove the **ulimit** package.

Syntax:

To install a package: **zypper –n install** *package-name*

To remove a package: **zypper remove** *package-name*

> **Note** Your Amazon EC2 instance should have access to the Internet to read and download required packages from the SUSE repository.

```
zypper -n install gtk2
zypper -n install java-1_6_0-ibm
zypper -n install libicu
zypper -n install mozilla-xulrunner*
zypper -n install ntp
zypper -n install sudo
zypper -n install syslog-ng
zypper -n install tcsh libssh2-1
zypper -n install autoyast2-installation
zypper -n install yast2-ncurses
zypper -n install libgcc_s1 libstdc++6
zypper -n install ntfs-3g
zypper -n install nfs-kernel-server

zypper remove ulimit
```

6. **Update NFS server ports.**

   As root user, update the NFS server configuration file to use specific ports.

   Use an editor to update the following ports in the /etc/sysconfig/nfs file:

```
STATD_PORT="4000"
LOCKD_TCPPORT="4001"
LOCKD_UDPPORT="4001"
MOUNTD_PORT="4002"
```

7. **Configure the NTP, AUTOFS, NFS and LVM services.**

   As root user, add the following entries to the /etc/ntp.conf file to configure the server pool for the NTP service. Before you make the changes, we recommend that you back up the existing ntp.conf file.

```
# cp /etc/ntp.conf /etc/ntp.conf.bak
# echo "server 0.pool.ntp.org" >> /etc/ntp.conf
# echo "server 1.pool.ntp.org" >> /etc/ntp.conf
# echo "server 2.pool.ntp.org" >> /etc/ntp.conf
# echo "server 3.pool.ntp.org" >> /etc/ntp.conf
```

Now, configure the ntp, autofs, boot.lvm, and nfsserver services to start automatically during an instance reboot.

Syntax:

To enable autostart for a service: **chkconfig** *service-name* **on**

To restart a service: **service** *service-name* **restart**

```
# chkconfig ntp on
# chkconfig autofs on
# chkconfig boot.lvm on
# chkconfig nfsserver on
# service ntp restart
# service autofs restart
# service nfsserver restart
```

8. Disable HugePages.

As root user, disable HugePages for your SAP HANA instance, and add an entry to the /etc/init.d/boot.local file to keep HugePages disabled during instance reboot. We recommend that you back up the existing /etc/init.d/boot.local file before you make this change.

```
# cp /etc/init.d/boot.local /etc/init.d/boot.local.bak
# echo never > /sys/kernel/mm/transparent_hugepage/enabled
# echo "echo never >
/sys/kernel/mm/transparent_hugepage/enabled" >>
/etc/init.d/boot.local
```

9. Add and update the kernel parameter.

As root user, add a new kernel parameter and load the kernel settings to activate the change dynamically. We recommend that you back up the /etc/sysctl.conf file before you make this change.

```
# cp /etc/sysctl.conf /etc/sysctl.conf.bak
# echo "net.ipv4.tcp_slow_start_after_idle=0" >>
/etc/sysctl.conf
```

```
# sysctl -p /etc/sysctl.conf
```

10. **Create physical volume(s).**

As root user, create physical volume(s). You need to create physical devices only for volumes that will be used to store SAP HANA data and log files. (Typically, these volumes are 334 GiB or 667 GiB in size.)

To view the list of volumes attached to your instance and their device name, run the **lsblk** command as root user. The command displays the list of devices that are attached to your instance.

Syntax: **lsblk**

Example:

```
# lsblk
NAME   MAJ:MIN RM   SIZE RO MOUNTPOINT
hda      3:0    0    10G  0
└─hda1   3:1    0    10G  0 /
xvdj   202:144  0    50G  0
xvdi   202:128  0   667G  0
xvdh   202:112  0   667G  0
xvdg   202:96   0   667G  0
xvdf   202:80   0   667G  0
$
```

Note The device might be attached with a different name than what you specify. For example, if you specify a device name of /dev/sdf, the kernel might rename your device /dev/xvdf, but in most cases, the trailing letter will remain the same.

Create physical devices by running the **pvcreate** command.

Syntax: **pvcreate** *device-name ... device-name*

Example:

```
# pvcreate /dev/xvdf /dev/xvdg /dev/xvdh /dev/xvdi
```

11. Set the I/O scheduling option to NOOP.

As root user, set the I/O scheduling option to "noop" mode to optimize the I/O performance of your SAP HANA instance on the AWS platform.

Syntax: **echo "noop" > /sys/block/***device-name***/queue/scheduler**

Example:

```
# echo "noop" > /sys/block/xvdf/queue/scheduler
# echo "noop" > /sys/block/xvdg/queue/scheduler
# echo "noop" > /sys/block/xvdh/queue/scheduler
# echo "noop" > /sys/block/xvdi/queue/scheduler
```

Important Make sure that you set the I/O scheduling option to "noop" for **all** the physical devices that you have created.

12. Create a volume group.

As root user, create a volume group called **vghana** to stripe the multiple physical volumes you created in step 10 for optimal I/O performance.

Syntax: **vgcreate** *vg-name device-name(s)*

Example:

```
# vgcreate vghana /dev/xvdf /dev/xvdg /dev/xvdh /dev/xvdi
```

13. Create logical volume(s)

As root user, create logical volumes for SAP HANA log and data files. The size of the logical volumes will differ depending on the Amazon EC2 instance type you choose, as shown in the following table.

Amazon EC2 Instance Type	Log Size (GiB)	Data Size (GiB)
c3.8xlarge	60	180
r3.2xlarge	60	180
r3.4xlarge	122	366
r3.8xlarge	244	732

Syntax: **lvcreate –n** *lv-name* **-i** *#-of-volumes-to-stripe-across* **-I 256 –L** *lv-size vg name*

Example:

```
# lvcreate -n lvhanalog -i 4 -I 256 -L 244G vghana
# lvcreate -n lvhanadata -i 4 -I 256 -L 732G vghana
```

14. Format logical volume(s).

As root user, format all logical volumes you created in the previous step. You can find the list of logical volumes in the directory /dev/mapper.

Syntax: **mkfs.xfs** *full-path-of-logical-volume*

Example:

```
# mkfs.xfs /dev/mapper/vghana-lvhanalog
# mkfs.xfs /dev/mapper/vghana-lvhanadata
```

15. Format block device for /usr/sap.

As root user, format the Amazon EBS volumes attached to your instance to store local SAP files (device name: /dev/xvdj).

Syntax: **mkfs.xfs -f** *device-name*

Example:

```
# mkfs.xfs -f /dev/xvdj
```

16. Create directories to mount file system(s).

As root user, create directories to mount the file systems required for SAP HANA worker node installation. You will need the following mount points:

— /hana/log

— /hana/data

— /usr/sap

— /hana/shared and /backup (these will be NFS-mounted from your master node)

Syntax: **mkdir** *directory-path*

Example:

```
# mkdir /hana /hana/log /hana/data /hana/shared /backup
/usr/sap
```

17. Create entries for mount points in the /etc/fstab file.

As root user, add entries to the /etc/fstab file and mount the file systems. Adding entries to /etc/fstab will ensure that your file systems are mounted automatically when your Amazon EC2 instance is restarted.

Add the entries for the SAP HANA data, log, and local SAP file systems to the /etc/fstab file by using the following commands.

Syntax: **echo** *"logical-volume-path mount-directory file-system-type mount-options" >> /etc/fstab*

Example:

```
# echo "/dev/xvdj                    /usr/sap       xfs
nobarrier,noatime,nodiratime,logbsize=256k,delaylog 0 0" >>
/etc/fstab
# echo "/dev/mapper/vghana-lvhanadata     /hana/data
xfs nobarrier,noatime,nodiratime,logbsize=256k,delaylog 0
0" >> /etc/fstab
# echo "/dev/mapper/vghana-lvhanalog      /hana/log
xfs nobarrier,noatime,nodiratime,logbsize=256k,delaylog 0
0" >> /etc/fstab
```

Alternatively, you can edit the /etc/fstab file and add these entries.

When you have added the entries to the /etc/fstab file, mount the file systems, and then run the **df** command to verify that they were mounted appropriately.

Syntax: **mount –a**
df -h

Example:

```
# mount -a
# df -h
Filesystem                      Size  Used Avail Use%
Mounted on
rootfs                          9.9G  1.8G  7.6G  20% /
udev                            121G  116K  121G   1% /dev
tmpfs                           121G     0  121G   0%
/dev/shm
/dev/hda1                       9.9G  1.8G  7.6G  20% /
/dev/xvdj                        50G   33M   50G   1%
/usr/sap
/dev/mapper/vghana-lvhanadata  732G   34M  732G   1%
/hana/data
/dev/mapper/vghana-lvhanalog   244G   33M  244G   1%
/hana/log
#
```

18. Update the /etc/hosts file with the master and worker IP addresses.

As root user, add the IP addresses and host names of your SAP HANA master and **all** the worker nodes into the /etc/hosts file of the HANA master node as well as **all** the worker nodes. This is required so that the master and worker nodes can resolve their host names for cross communication.

We recommend that you back up your existing /etc/hosts file before you make these changes.

Syntax: *IP-address hostname hostname.domain*

Example:

```
10.0.1.2 imdbmaster imdbmaster.local
10.0.1.3 imdbworker01 imdbworker01.local
```

19. Configure NFS exports in the SAP HANA master node.

Connect to your SAP HANA master node by using SSH. As root user, add the following entries to the /etc/exports file, and export the file systems.

Syntax: *filesystem-name NFS-client-host(s)* (*mount-options*)

```
# echo "/hana/shared
imdbworker01(rw,no_root_squash,no_subtree_check)" >>
/etc/exports
# echo "/backup
imdbworker01(rw,no_root_squash,no_subtree_check)" >>
/etc/exports
```

Tip If you have multiple worker nodes that have common nomenclature in their host names, you can use wildcard characters while specifying NFS client hosts. For example, you can specify `imdbworker* (rw,no_root_squash,no_subtree_check)` to export your file systems to multiple hosts (imdbworker01, imdbworker02, and so on.).

When the entries have been added to the /etc/exports file, export the NFS file systems and verify that they have been exported to the correct hosts.

Syntax:

exportfs -a
showmount -e

Example:

```
# exportfs -a
# showmount –e
```

20. **Configure the AUTOFS service in SAP HANA worker nodes.**

As root user, configure the **autofs** service in the SAP HANA worker nodes to access the NFS file systems that have been exported from the SAP HANA master node.

Edit the /etc/auto.master file, comment out the entry `+auto.master`, and append the entry `/- auto.direct`:

```
#+auto.master
```

```
/- auto.direct
```

Now, add the entries to the /etc/auto.direct file to automatically mount /hana/shared and /backup.

Syntax: *mount-directory mount-options nfsserver.domain*:*exported-mount-point*

```
# echo "/hana/shared      -
rw,rsize=32768,wsize=32768,timeo=14,intr
imdbmaster.local:/hana/shared" >> /etc/auto.direct
# echo "/backup           -
rw,rsize=32768,wsize=32768,timeo=14,intr
imdbmaster.local:/backup" >> /etc/auto.direct
```

Finally, restart the **autofs** service to activate the changes and try accessing the /hana/shared and /backup directories.

Syntax: **service** *service-name* **restart**

```
# service autofs restart
# ls /hana/shared
# ls /backup
```

21. Attach and mount the HANA media volume for installation.

Use AWS CLI to attach the HANA media volume to your SAP HANA worker node, and mount it to a file system for installation.

Syntax: **aws ec2 attach-volume --volume-id** *volume-name* **--instance** *instance-name* **--device** *device-name*

Example:

```
$ aws ec2 attach-volume --volume-id vol-xxxxxxxx --instance
i-yyyyyyyy --device /dev/sdk
```

> **Note** The device might be attached with a different name than what you specify. For example, if you specify a device name of /dev/sdf, the kernel might rename your device /dev/xvdf, but in most cases, the trailing letter will remain the same.

Now mount your HANA media volume.

Syntax: **mount** *device-name mount-directory*

```
# mkdir /media
# mount /dev/xvdk1 /media
```

> **Tip** To have /media mounted during every system reboot, add an entry to the /etc/fstab file. We recommend that you back up your existing /etc/fstab file before you make this change. See Step 17 for instructions.

22. **Install SAP HANA.**

Follow the standard SAP installation procedure to add additional worker nodes to your SAP HANA instance.

> **Note** If you are using SAP HANA Service Pack 9 (SP9), you might have to install SAP Host Agent manually first, before adding worker nodes. See OSS Note 2101350 for details (requires access to SAP Service Marketplace).

Appendix: Security Group Specifics

The following inbound and outbound protocols and ports are recommended for HANA scale-out deployment. The CIDR blocks listed in the table are examples; you will need to adjust them according to your network setup. In the table, 10.0.1.0/24 is the private subnet CIDR, and 10.0.2.0/24 is the public subnet CIDR.

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	Internal database 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 NAT instance
10.0.2.0/24	TCP	1128 - 1129	Host agent access
10.0.2.0/24	TCP	43##	Access to XSEngine (HTTPS) from 10.0.2.0 subnet
10.0.2.0/24	TCP	80##	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	3##15	Database client access
10.0.2.0/24	TCP	3##17	Database client access
10.0.2.0/24	TCP	5##13 – 5##14	Allow access for HANA Studio from RDP instance
Outbound			
0.0.0.0/0	TCP	1 - 65535	Outbound access from SAP HANA master allowed to anywhere

Send Us Your Feedback

We welcome your feedback. For comments or questions, please fill out the form at http://aws.amazon.com/sap/interest.

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.