SAP HANA on AWS Operations Guide

February 2014
(Updated: November 2015)

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.

Contents

Introduction

This guide provides best practices for operating SAP HANA systems that have been deployed on Amazon Web Services (AWS) either by using the SAP HANA Quick Start reference deployment process[1] or manually, by following the instructions in Setting up AWS Resources and the SLES Operating System for SAP HANA Installation.[2] This guide is not intended to replace any of the standard SAP documentation. SAP guides and notes can be found at the following locations:

- SAP Library (help.sap.com) - SAP HANA Administration Guide[3]
- SAP installation guides[4] (these require SAP Support Portal access)
- SAP notes[5] (these require SAP Support Portal access)

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

- Getting Started with AWS[6]
- What is Amazon EC2?[7]

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

- SAP on AWS Implementation and Operations Guide provides best practices for achieving optimal performance, availability, and reliability, and lower total cost of ownership (TCO) while running SAP solutions on AWS.[8]
- SAP on AWS High Availability Guide explains how to configure SAP systems on Amazon Elastic Compute Cloud (Amazon EC2) to protect your application from various single points of failure.[9]
- SAP on AWS Backup and Recovery Guide explains how to back up SAP systems running on AWS, in contrast to backing up SAP systems on traditional infrastructure.[10]

Administration

This section provides guidance on common administrative tasks required to operate an SAP HANA system, including information about starting, stopping, and cloning systems.

Starting and Stopping EC2 Instances Running SAP HANA Hosts

At any time, you can stop one or multiple SAP HANA hosts. Before stopping the Amazon EC2 instance of an SAP HANA host, first stop SAP HANA on that instance. When you resume the instance, it will automatically be started with the same IP address, network, and storage configuration as before.

Creating an Image of an SAP HANA System

You can create your own Amazon Machine Image (AMI) based on an existing instance by using the AWS Management Console or the command line.[11] For more information, see the AWS documentation.[12] You could use an AMI of your SAP HANA instance for the following purposes:

- **To create a full offline system backup** (of the OS, /usr/sap, HANA shared, backup, data, and log files) – AMIs are automatically saved in three different Availability Zones within the same region.
- **To move a HANA system from one region to another** – You can create an image of an existing Amazon EC2 instance and move to another region by following the instructions in the AWS documentation.[13] Once the AMI has been copied to the target region, the new instance can be launched in the target region.
- **To clone an SAP HANA system** – You can create an AMI of an existing SAP HANA system to create an exact clone of the system. See the following section for additional information.

Tip The SAP HANA system should be in a consistent state before you create an AMI. To do this, stop the SAP HANA instance before creation or by following the instructions in SAP Note 1703435 (requires SAP Support Portal access).[14]

Cloning an SAP HANA System

Single-node systems – To create a clone of a single-node SAP HANA system, you create an image in Amazon EC2, as discussed in the previous section.

Multi-node systems – Multi-node SAP HANA systems cannot be cloned by creating an image. Instead, follow these steps:

1. Use the SAP HANA Quick Start reference deployment to provision a new SAP HANA system with the same configuration as the HANA system you want to clone.
2. Perform a data backup of the original system.
3. Restore the backup of the original system into the new system.

Backup and Recovery

This section provides an overview of the AWS services used in the backup and recovery of SAP HANA systems and provides an example backup and recovery scenario. This guide does not include detailed instructions on how to execute database backups using native HANA backup/recovery features or third-party backup tools. Please refer to the standard OS, SAP, and SAP HANA documentation, or the documentation provided by backup software vendors. In addition, backup schedules, frequency, and retention periods may vary with your system type and business requirements. See the following standard SAP documentation for guidance on these topics (SAP notes require SAP Support Portal access).

> **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[15]	FAQ: SAP HANA Database Backup & Recovery
1821207[16]	Determining required recovery files

SAP Note #	Description
1869119[17]	Checking backups using `hdbbackupcheck`
1873247[18]	Checking recoverability with `hdbbackupdiag --check`
1651055[19]	Scheduling SAP HANA Database Backups in Linux

AWS Services and Components for Backup Solutions

AWS provides a number of services and options for storage and backup, including Amazon Simple Storage Service (Amazon S3), AWS Identity and Access Management, and Amazon Glacier.

Amazon S3

Amazon S3 is the center of any SAP backup and recovery solution on AWS.[20] It provides a highly durable storage infrastructure designed for mission-critical and primary data storage, and is designed to provide 99.999999999% durability and 99.99% availability over a given year. See the Amazon S3 documentation for detailed instructions on how to create and configure an Amazon S3 bucket to store your SAP HANA backup files.[21]

AWS Identity and Access Management

With AWS Identity and Access Management (IAM), you can securely control access to AWS services and resources for your users.[22] 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.

During the deployment process, AWS CloudFormation creates an IAM role that allows access to get and/or put objects to and from Amazon S3. That role is subsequently assigned to each AWS instance that is hosting SAP HANA master and worker nodes at launch time as they are deployed.

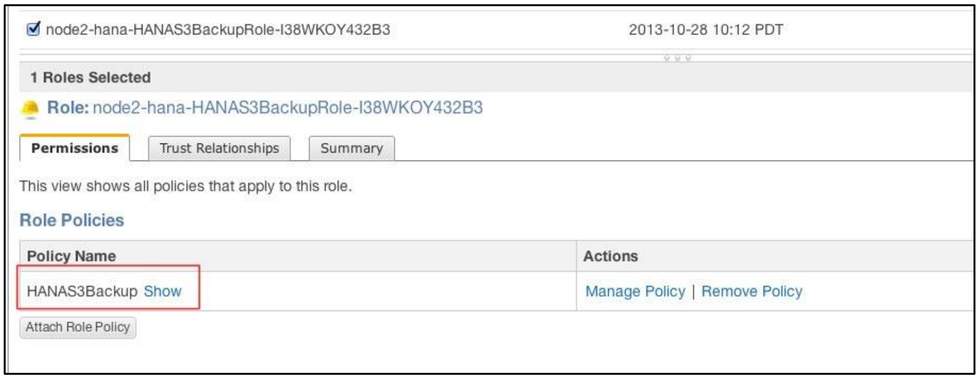

☑ node2-hana-HANAS3BackupRole-I38WKOY432B3 2013-10-28 10:12 PDT

1 Roles Selected

Role: node2-hana-HANAS3BackupRole-I38WKOY432B3

This view shows all policies that apply to this role.

Role Policies

Policy Name	Actions
HANAS3Backup Show	Manage Policy | Remove Policy

Figure 1: IAM Role Example

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

```
{"Statement":[
 {"Resource":"arn:aws:s3::: <your-s3-bucket-name>/*",

"Action":["s3:GetObject","s3:PutObject","s3:DeleteObject",
"s3:ListBucket","s3:Get*","s3:List*"],
    "Effect":"Allow"},

{"Resource":"*","Action":["s3:List*","ec2:Describe*","ec2:A
ttachNetworkInterface",

"ec2:AttachVolume","ec2:CreateTags","ec2:CreateVolume","ec2
:RunInstances",
  "ec2:StartInstances"],"Effect":"Allow"}]}
```

To add functions later, you can use the AWS Management Console to modify the IAM role.

Amazon Glacier

Amazon Glacier is an extremely low-cost service that provides secure and durable storage for data archiving and backup.[23] 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. You can use lifecycle policies, as explained in the *Amazon S3 Developer Guide*, to push SAP HANA backups to Amazon Glacier for long-term archiving.[24]

Backup Destination

The primary difference between backing up SAP systems on AWS compared with traditional on-premises infrastructure is the backup destination. The typical backup destination used with on-premises infrastructure is tape. On AWS, backups are stored in Amazon S3 instead of on tape. Amazon S3 has many benefits over tape, including the ability to automatically store backups "offsite" from the source system, since data in Amazon S3 is replicated across multiple facilities within the AWS region.

SAP HANA systems provisioned using the SAP HANA Quick Start reference deployment are configured with a set of Amazon Elastic Block Store (Amazon EBS) volumes to be used as an initial local backup destination. HANA backups are first stored on these local EBS volumes, and then copied to Amazon S3 for long-term storage.

You can use SAP HANA Studio, SQL commands, or the DBA Cockpit to start or schedule SAP HANA data backups. Log backups are written automatically unless disabled. The /backup file system is configured as part of the deployment process.

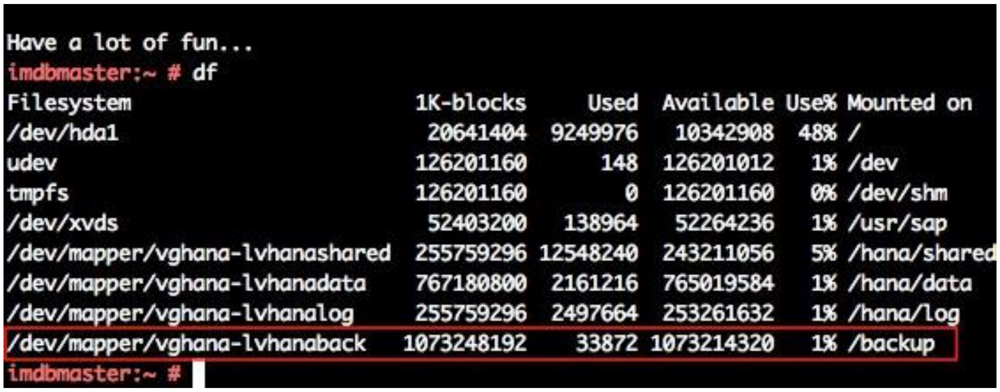

```
Have a lot of fun...
imdbmaster:~ # df
Filesystem                       1K-blocks     Used  Available Use% Mounted on
/dev/hda1                         20641404  9249976   10342908  48% /
udev                             126201160      148  126201012   1% /dev
tmpfs                            126201160        0  126201160   0% /dev/shm
/dev/xvds                         52403200   138964   52264236   1% /usr/sap
/dev/mapper/vghana-lvhanashared  255759296 12548240  243211056   5% /hana/shared
/dev/mapper/vghana-lvhanadata    767180800  2161216  765019584   1% /hana/data
/dev/mapper/vghana-lvhanalog     255759296  2497664  253261632   1% /hana/log
/dev/mapper/vghana-lvhanaback   1073248192    33872 1073214320   1% /backup
imdbmaster:~ #
```

Figure 2: SAP HANA File System Layout

The SAP HANA global.ini configuration file has been customized by the SAP HANA Quick Start reference deployment process 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, is installed as part of the base image.[25] Using various commands, you can control multiple AWS services from the command line directly, and automate them through scripts. Access to the Amazon S3 bucket is available through the IAM role assigned to the instance (discussed earlier). Using the AWS CLI commands for Amazon S3, you can list the contents of the previously created bucket, back up files, and restore files, as explained in the AWS CLI documentation.[26]

```
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

Here are the steps you might take for a typical backup task:

1. In the SAP HANA Backup Editor, choose **Open Backup Wizard**. You can also open the Backup Wizard by right-clicking the system that you want to back up and choosing **Back Up**.

   a. Select destination type **File**. This will back up the database to files in the specified file system.

   b. Specify the backup destination (/backup/data/<SID>) and the backup prefix.

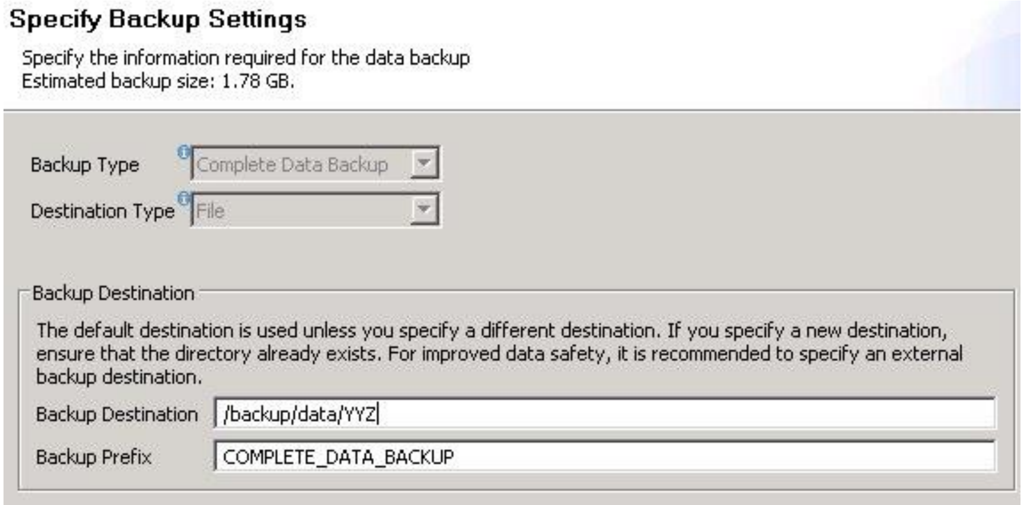

Specify Backup Settings

Specify the information required for the data backup
Estimated backup size: 1.78 GB.

Backup Type: Complete Data Backup

Destination Type: File

Backup Destination

The default destination is used unless you specify a different destination. If you specify a new destination, ensure that the directory already exists. For improved data safety, it is recommended to specify an external backup destination.

Backup Destination: /backup/data/YYZ

Backup Prefix: COMPLETE_DATA_BACKUP

Figure 3: SAP HANA Backup Example

   c. Click **Next** and then **Finish**. A confirmation message will appear when the backup is complete.

   d. Verify that 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
```

2. The next step is to push or synchronize the backup files from the `/backup` file system to S3 by using the `aws s3 sync` command.[27]

```
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
```

3. Use the AWS Management Console to verify that the files have been pushed to Amazon S3. You can also use the `aws s3 ls` command shown previously in the AWS Command Line Interface section.

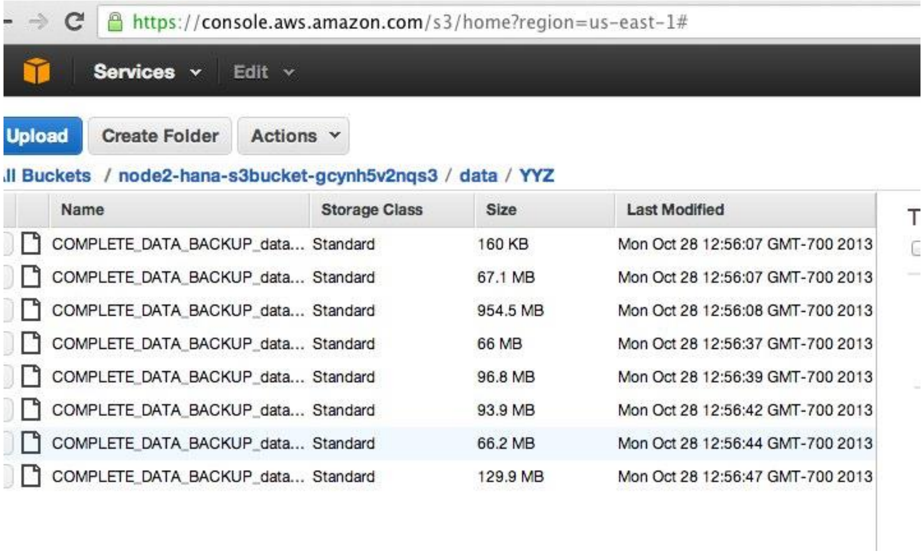
https://console.aws.amazon.com/s3/home?region=us-east-1#
Services
Edit
Upload
Create Folder
Actions
ll Buckets / node2-hana-s3bucket-gcynh5v2nqs3 / data / YYZ

Name	Storage Class	Size	Last Modified
COMPLETE_DATA_BACKUP_data...	Standard	160 KB	Mon Oct 28 12:56:07 GMT-700 2013
COMPLETE_DATA_BACKUP_data...	Standard	67.1 MB	Mon Oct 28 12:56:07 GMT-700 2013
COMPLETE_DATA_BACKUP_data...	Standard	954.5 MB	Mon Oct 28 12:56:08 GMT-700 2013
COMPLETE_DATA_BACKUP_data...	Standard	66 MB	Mon Oct 28 12:56:37 GMT-700 2013
COMPLETE_DATA_BACKUP_data...	Standard	96.8 MB	Mon Oct 28 12:56:39 GMT-700 2013
COMPLETE_DATA_BACKUP_data...	Standard	93.9 MB	Mon Oct 28 12:56:42 GMT-700 2013
COMPLETE_DATA_BACKUP_data...	Standard	66.2 MB	Mon Oct 28 12:56:44 GMT-700 2013
COMPLETE_DATA_BACKUP_data...	Standard	129.9 MB	Mon Oct 28 12:56:47 GMT-700 2013

Figure 4: Amazon S3 Bucket Contents After Backup

Tip The `aws s3 sync` command will only upload new files that don't exist in Amazon S3. Use a periodically scheduled cron job to sync, and then delete files that have been uploaded. See SAP note 1651055 for scheduling periodic backup jobs in Linux, and extend the supplied scripts with `aws s3 sync` commands.

Restore Example

When you need to restore your SAP HANA database from a backup, you can take the following steps.

1. If the backup files are not readily available already in the /backup file system but are in Amazon S3, restore the files from Amazon S3 by using the `aws s3 cp` command,[28] which has the following syntax:

```
aws --region <region> cp <s3-bucket/path> --recursive
<backup-prefix>*.
```

For 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 by using the Recovery Wizard, as outlined in the SAP HANA Administration Guide. Take care to specify **File** as the **Destination Type** and enter the correct backup prefix.

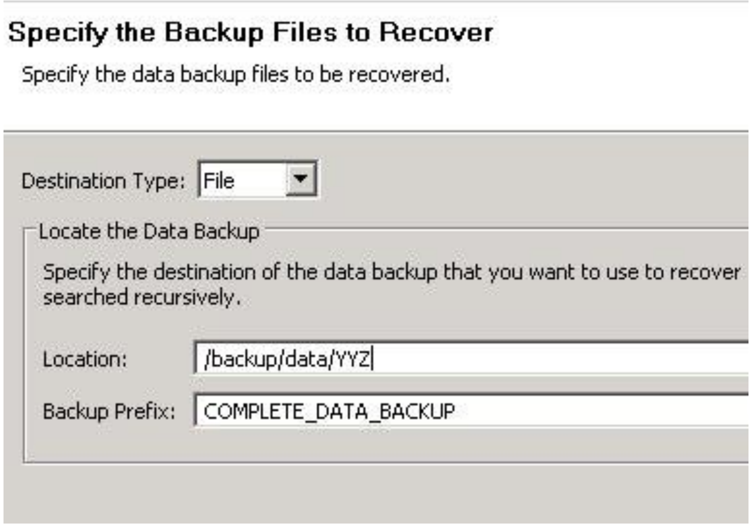

Specify the Backup Files to Recover

Specify the data backup files to be recovered.

Destination Type: File

Locate the Data Backup

Specify the destination of the data backup that you want to use to recover searched recursively.

Location: /backup/data/YYZ

Backup Prefix: COMPLETE_DATA_BACKUP

Figure 5: Restore Example

3. When the recovery is complete, you can resume normal operations and clean up backup files from the `/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. The following information serves only as a supplement to the information contained in the “Getting Support” section of the SAP HANA Administration Guide.

A few steps are required to configure proper connectivity to SAP. These steps differ depending on whether you want to use an existing remote network connection to SAP, or 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 direct support connection to SAP from AWS, consider the following steps:

1. Create and configure a specific SAProuter security group, which only allows the required inbound and outbound access to the SAP support network, for the SAProuter instance. This should be limited to a specific IP address that SAP gives you to connect to, along with TCP port 3299. See the Amazon EC2 security group documentation for additional details about creating and configuring security groups.[29]
2. Launch the instance that the SAProuter software will be installed on into a public subnet of the Amazon Virtual Private Network (Amazon VPC) and assign it an Elastic IP address (EIP).
3. Install the SAProuter software and create a `saprouttab` file that allows access from SAP to your SAP HANA systems on AWS.
4. Set up the connection with SAP. For your Internet connection, use **Secure Network Communication (SNC)**. For more information, see the SAP Remote Support – Help page.[30]
5. Modify the existing SAP HANA security groups to trust the new SAProuter security group you have created.

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

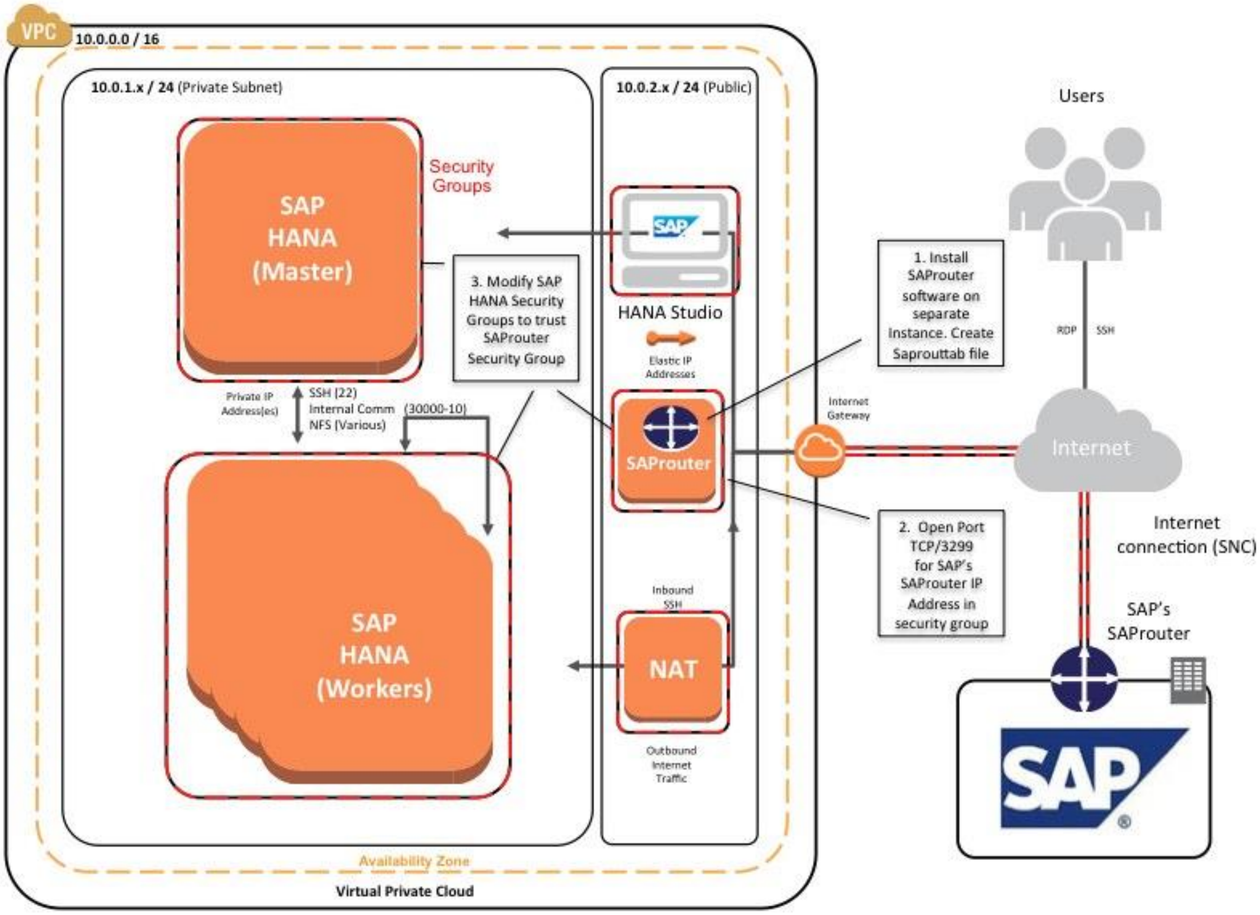

Figure 6: Support Connectivity with SAProuter on AWS

Support Channel Setup with SAProuter on Premises

In many cases, you may already have a support connection configured between your data center and SAP. This can easily be extended to support SAP systems on AWS. This scenario assumes that connectivity between your data center and AWS has already been established, either by way of a secure VPN tunnel over the Internet, or by using AWS Direct Connect.

Extending this connectivity takes only a few steps:

1. Ensure that the proper `saprouttab` entries exist to allow access from SAP to resources in the Amazon VPC.
2. Modify the SAP HANA security groups to allow access from the on-premises SAProuter IP address.

3. Ensure that the proper firewall ports are open on your gateway to allow traffic to pass over TCP port 3299.

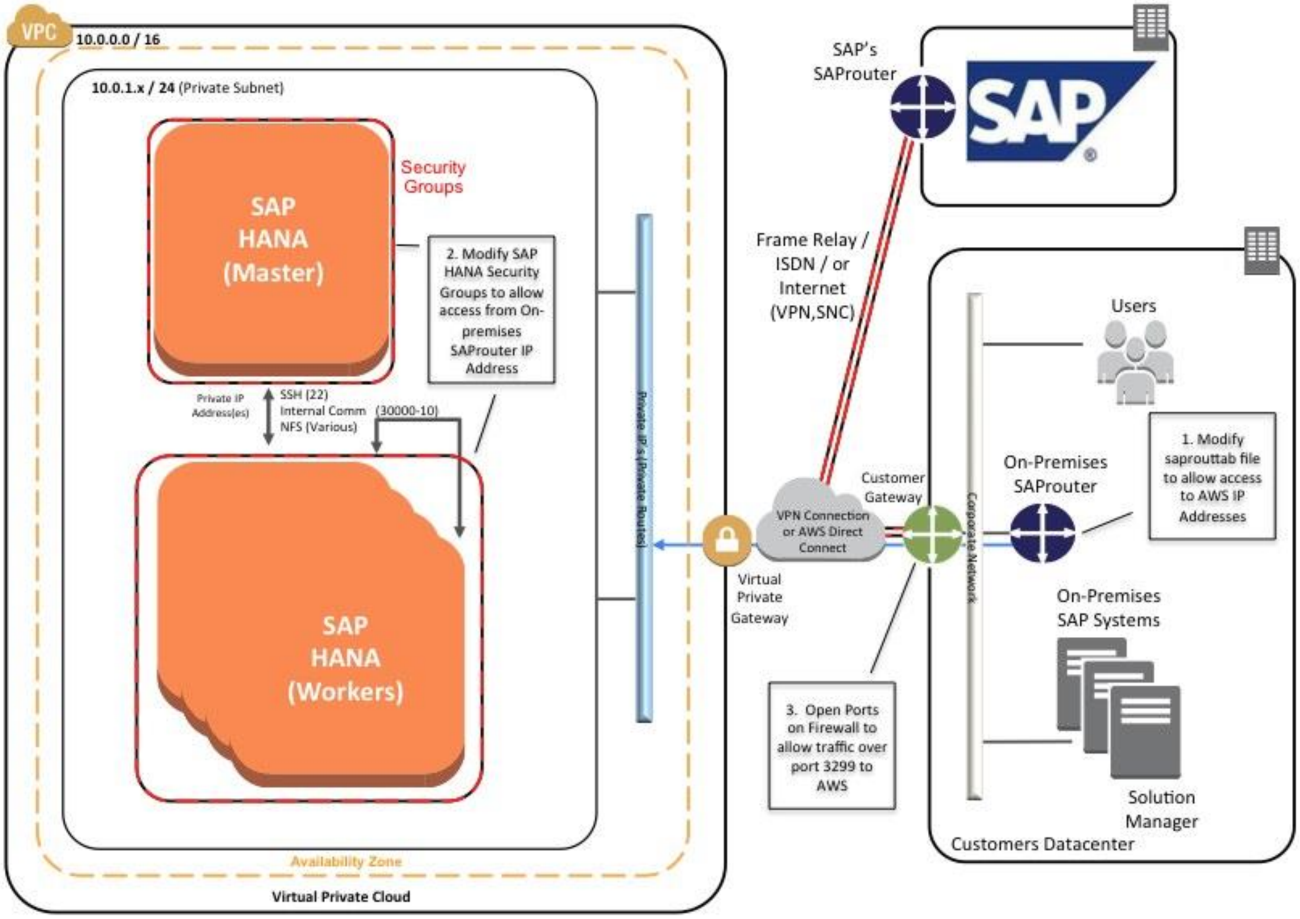
VPC
10.0.0.0 / 16
10.0.1.x / 24 (Private Subnet)
SAP HANA (Master)
Security Groups
2. Modify SAP HANA Security Groups to allow access from On-premises SAProuter IP Address
Private IP Address(es)
SSH (22)
Internal Comm (30000-10)
NFS (Various)
SAP HANA (Workers)
Private IP's (Private Routes)
Availability Zone
Virtual Private Cloud
Virtual Private Gateway
VPN Connection or AWS Direct Connect
Customer Gateway
SAP's SAProuter
SAP
Frame Relay / ISDN / or Internet (VPN,SNC)
3. Open Ports on Firewall to allow traffic over port 3299 to AWS
Corporate Network
Users
On-Premises SAProuter
1. Modify saprouttab file to allow access to AWS IP Addresses
On-Premises SAP Systems
Solution Manager
Customers Datacenter

Figure 7: Support Connectivity with SAProuter on Premises

Security

This section discusses additional security topics you may want to consider that are not covered in the SAP HANA Quick Start reference deployment guide.

Here are additional AWS security resources to help you achieve the level of security you require for your SAP HANA environment on AWS:

- AWS Cloud Security Center - http://aws.amazon.com/security/
- AWS Cloud Security Whitepaper[31]
- AWS Cloud Security Best Practices Whitepaper[32]

OS Hardening

You may want to lock down the OS configuration further—for example, to avoid providing a DB administrator with root credentials when logging into an instance.

Please also refer to the following SAP notes:

- 1730999: *Configuration changes in HANA appliance*[33]
- 1731000: *Unrecommended configuration changes*[34]

Disabling HANA Services

HANA services such as HANA XS are optional and should be deactivated if they are not needed. For instructions, see SAP Note 1697613: *Remove XS Engine out of SAP HANA database*.[35] In case of service deactivation, you should also remove the TCP ports from the SAP HANA AWS security groups for complete security.

API Call Logging

AWS CloudTrail is a web service that records AWS API calls for your account and delivers log files to you.[36] The recorded information includes the identity of the API caller, the time of the API call, the source IP address of the API caller, the request parameters, and the response elements returned by the AWS service.

With CloudTrail, you can get a history of AWS API calls for your account, including API calls made via the AWS Management Console, AWS SDKs, command line tools, and higher-level AWS services (such as AWS CloudFormation). The AWS API call history produced by CloudTrail enables security analysis, resource change tracking, and compliance auditing.

Notifications on Access

You can use Amazon Simple Notification Service (Amazon SNS) or third-party applications to set up notifications on SSH login to your email address or mobile phone.[37]

High Availability and Disaster Recovery

For details and best practices for high availability and disaster recovery of SAP HANA systems running on AWS, see High Availability and Disaster Recovery Options for SAP HANA on AWS.[38]

Conclusion

This guide discussed best practices for the operation of SAP HANA systems on the AWS cloud. The best practices provided in this guide will help you efficiently manage and achieve maximum benefits from running your SAP HANA systems on the AWS cloud.

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

Contributors

The following individuals and organizations contributed to this document:

- Bill Timm, Partner Solution Architect, AWS

Notes

[1] http://docs.aws.amazon.com/quickstart/latest/sap-hana/ or https://s3.amazonaws.com/quickstart-reference/sap/hana/latest/doc/SAP+HANA+Quick+Start.pdf

[2] http://d0.awsstatic.com/enterprise-marketing/SAP/SAP-HANA-on-AWS-Manual-Setup-Guide.pdf

[3] https://help.sap.com/hana/SAP_HANA_Administration_Guide_en.pdf

[4] http://service.sap.com/instguides

[5] http://service.sap.com/notes

[6] http://docs.aws.amazon.com/gettingstarted/latest/awsgsg-intro/intro.html

[7] http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts.html

[8] http://d0.awsstatic.com/enterprise-marketing/SAP/SAP_on_AWS_Implementation_Guide.pdf

[9] http://d0.awsstatic.com/enterprise-marketing/SAP/SAP_on_AWS_High_Availability_Guide_v3.2.pdf

[10] http://d0.awsstatic.com/enterprise-marketing/SAP/sap-on-aws-backup-and-recovery-guide-v2-2.pdf

[11] http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIs.html

[12] http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/creating-an-ami-ebs.html

[13] http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/CopyingAMIs.html

[14] https://service.sap.com/notes/1703435

[15] http://service.sap.com/sap/support/notes/1642148

[16] http://service.sap.com/sap/support/notes/1821207

[17] http://service.sap.com/sap/support/notes/1869119

[18] http://service.sap.com/sap/support/notes/1873247

[19] http://service.sap.com/sap/support/notes/1651055

[20] http://aws.amazon.com/s3/

[21] http://aws.amazon.com/documentation/s3/

[22] http://aws.amazon.com/iam/

[23] http://aws.amazon.com/glacier/

[24] http://docs.aws.amazon.com/AmazonS3/latest/dev/object-archival.html

[25] http://aws.amazon.com/cli/

[26] http://docs.aws.amazon.com/cli/latest/reference/s3/

[27] http://docs.aws.amazon.com/cli/latest/reference/s3/sync.html

[28] http://docs.aws.amazon.com/cli/latest/reference/s3/cp.html

[29] http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html

[30] https://support.sap.com/remote-support/help.html

[31] http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf

[32] http://d0.awsstatic.com/whitepapers/aws-security-best-practices.pdf

[33] https://service.sap.com/sap/support/notes/1730999

[34] https://service.sap.com/sap/support/notes/1731000

[35] https://service.sap.com/sap/support/notes/1697613

[36] https://aws.amazon.com/cloudtrail/

[37] https://aws.amazon.com/sns/

[38] http://d0.awsstatic.com/enterprise-marketing/SAP/sap-hana-on-aws-high-availability-disaster-recovery-guide.pdf