

SAP HANA System Replication Scale-Out High Availability in Amazon Web Services

SUSE Linux Enterprise Server for SAP Applications 12 SP4 or newer
Amazon Web Services

Fabian Herschel, Distinguished Architect SAP (SUSE)

Bernd Schubert, SAP Solution Architect (SUSE)

Lars Pinne, System Engineer (SUSE)

Guilherme Felix, SAP Specialist (AWS)

Sander Bleijenbergh, SAP Specialist (AWS)

Somckit Khemmanivanh, SAP Specialist (AWS)

SAP HANA System Replication Scale-Out High Availability in Amazon Web Services

Date: 2024-11-12

SUSE® Linux Enterprise Server for SAP Applications is optimized in various ways for SAP* applications. This guide provides detailed information about how to install and customize SUSE Linux Enterprise Server for SAP Applications for SAP HANA Scale-Out system replication with automated failover in the Amazon Web Services Cloud.

This document is based on SUSE Linux Enterprise Server for SAP Applications 12 SP4. The concept however can also be used with newer versions.

Disclaimer: Documents published as part of the SUSE Best Practices series have been contributed voluntarily by SUSE employees and third parties. They are meant to serve as examples of how particular actions can be performed. They have been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. SUSE cannot verify that actions described in these documents do what is claimed or whether actions described have unintended consequences. SUSE LLC, its affiliates, the authors, and the translators may not be held liable for possible errors or the consequences thereof.

Contents

- 1 About this Guide 4
- 2 Scope of This Documentation 5
- 3 Planning the Installation 6
- 4 Using AWS with SUSE Linux Enterprise High Availability Extension Clusters 14
- 5 Installing the SAP HANA Databases on both sites 30
- 6 Set up the SAP HANA System Replication 33
- 7 Integration of SAP HANA with the Cluster 37
- 8 Configuration of the Cluster and SAP HANA Resources 40
- 9 Testing the Cluster 53
- 10 Administration 56
- 11 Useful Links, Manuals, and SAP Notes 61
- 12 Legal notice 65
- 13 GNU Free Documentation License 66

1 About this Guide

1.1 Introduction

SUSE® Linux Enterprise Server for SAP Applications is optimized in various ways for SAP* applications. This guide provides detailed information about installing and customizing *SUSE Linux Enterprise Server for SAP Applications* for SAP HANA Scale-Out system replication with automated failover in the Amazon Web Services (AWS) Cloud.

High availability is an important aspect of running your mission-critical SAP HANA servers.

The SAP HANA Scale-Out system replication is a replication of all data in one SAP HANA Scale-Out cluster to a secondary SAP HANA cluster in different AWS Availability Zones. SAP HANA supports asynchronous and synchronous replication modes. We describe here the synchronous replication from memory of the primary system to memory of the secondary system, because it is the only method which allows the pacemaker cluster to make decisions based on the implemented algorithms.

The recovery time objective (RTO) is minimized through the data replication at regular intervals.

1.2 Additional Documentation and Resources

Chapters in this manual contain links to additional documentation resources that are either available through Linux manual pages or on the Internet.

For the latest documentation updates, see <https://documentation.suse.com>.

You can also find numerous whitepapers, guides, and other resources at the SUSE Linux Enterprise Server for SAP Applications resource library at <https://www.suse.com/products/sles-for-sap/resource-library/>.

1.3 Feedback

Several feedback channels are available:

Bugs and Enhancement Requests

For services and support options available for your product, refer to <http://www.suse.com/support/>.

To report bugs for a product component, go to <https://scc.suse.com/support/> requests, log in, and select *Submit New SR* (Service Request).

Mail

For feedback on the documentation of this product, you can send a mail to doc-team@suse.com (<mailto:doc-team@suse.com>). Make sure to include the document title, the product version and the publication date of the documentation. To report errors or suggest enhancements, provide a concise description of the problem and refer to the respective section number and page (or URL).

2 Scope of This Documentation

This document describes how to set up an SAP HANA Scale-Out system replication cluster installed on separated AWS Availability Zones based on SUSE Linux Enterprise Server for SAP Applications 12 SP4 or SUSE Linux Enterprise Server for SAP Applications 15.

To give a better overview, the installation and setup is separated into seven steps.




FIGURE 1: INSTALLATION STEPS

- Planning (section *Section 3, "Planning the Installation"*)
- Operating System (OS) setup (section *Section 4, "Using AWS with SUSE Linux Enterprise High Availability Extension Clusters"*)
- SAP HANA installation (section *Section 5, "Installing the SAP HANA Databases on both sites"*)
- SAP HANA system replication configuration (section *Section 6, "Set up the SAP HANA System Replication"*)
- SAP HANA cluster integration (section *Section 7, "Integration of SAP HANA with the Cluster"*)
- SLES for SAP cluster configuration (section *Section 8, "Configuration of the Cluster and SAP HANA Resources"*)
- Testing (section *Section 9, "Testing the Cluster"*)

As the result of the setup process you will have a SUSE Linux Enterprise Server for SAP Applications cluster controlling two groups of SAP HANA Scale-Out nodes in system replication configuration. This architecture was named the 'performance optimized scenario' because failovers should only take a few minutes.

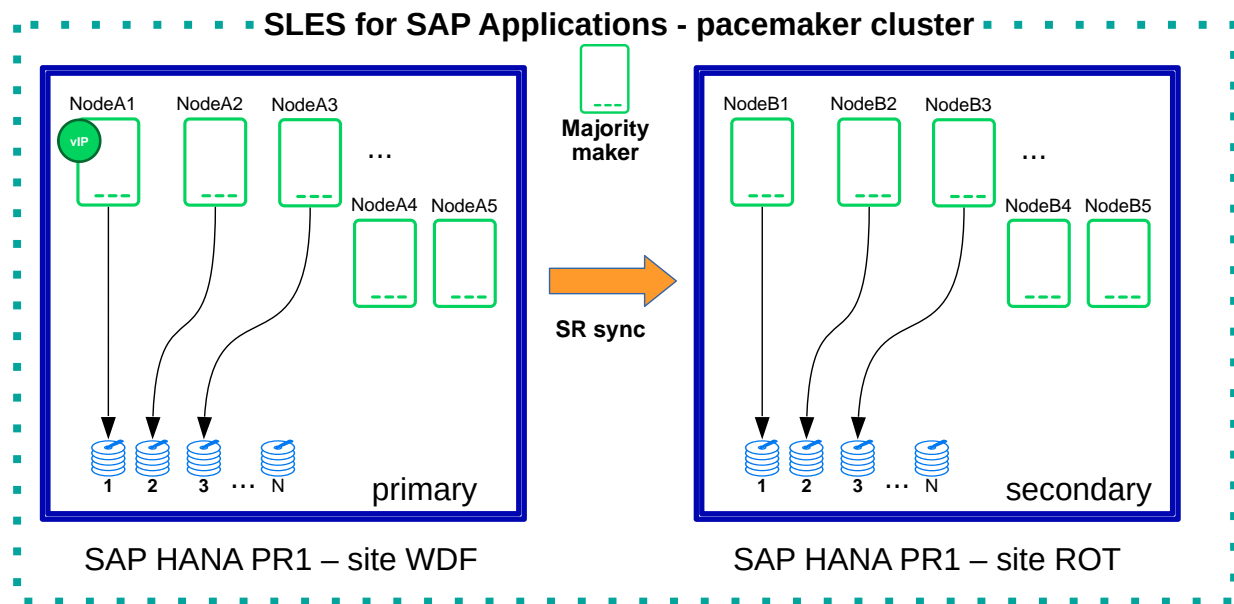


FIGURE 2: CLUSTER WITH SAP HANA SR - PERFORMANCE OPTIMIZED

3 Planning the Installation

Planning the installation is essential for a successful SAP HANA cluster setup.




FIGURE 3: PLANNING SECTION 4, "USING AWS WITH SUSE LINUX ENTERPRISE HIGH AVAILABILITY EXTENSION CLUSTERS" SECTION 5, "INSTALLING THE SAP HANA DATABASES ON BOTH SITES" SECTION 6, "SET UP THE SAP HANA SYSTEM REPLICATION" SECTION 7, "INTEGRATION OF SAP HANA WITH THE CLUSTER" SECTION 8, "CONFIGURATION OF THE CLUSTER AND SAP HANA RESOURCES" SECTION 9, "TESTING THE CLUSTER"

What you need before you start:

- EC2 instances created using "SUSE Linux Enterprise Server for SAP Applications 12 SP4" , "SUSE Linux Enterprise Server for SAP Applications 15" or later created from an Amazon Machine Image (AMI). If using a Bring Your Own Subscription (BYOS) AMI, a valid SUSE product subscription is required.
- SAP HANA installation media
- Two Amazon Elastic File System (EFS) - one per Availability Zone
- Filled parameter sheet (see below)

3.1 Environment Requirements

This section defines the minimum requirements to install SAP HANA Scale-Out and create a cluster in AWS.




Note

The minimum requirements mentioned here do not include SAP sizing information. For sizing information use the official SAP sizing tools and services.



Note

The preferred method to deploy SAP HANA scale-out clusters in AWS is to follow the AWS QuickStarts using the "single Availability Zone (AZ) and multi-node architecture" deployment option. If you are installing SAP HANA Scale-Out manually, refer to the [AWS SAP HANA Guides \(https://docs.aws.amazon.com/sap/latest/sap-hana/welcome.html\)](https://docs.aws.amazon.com/sap/latest/sap-hana/welcome.html)  documentation for detailed installation instructions, including recommended storage configuration and file systems.

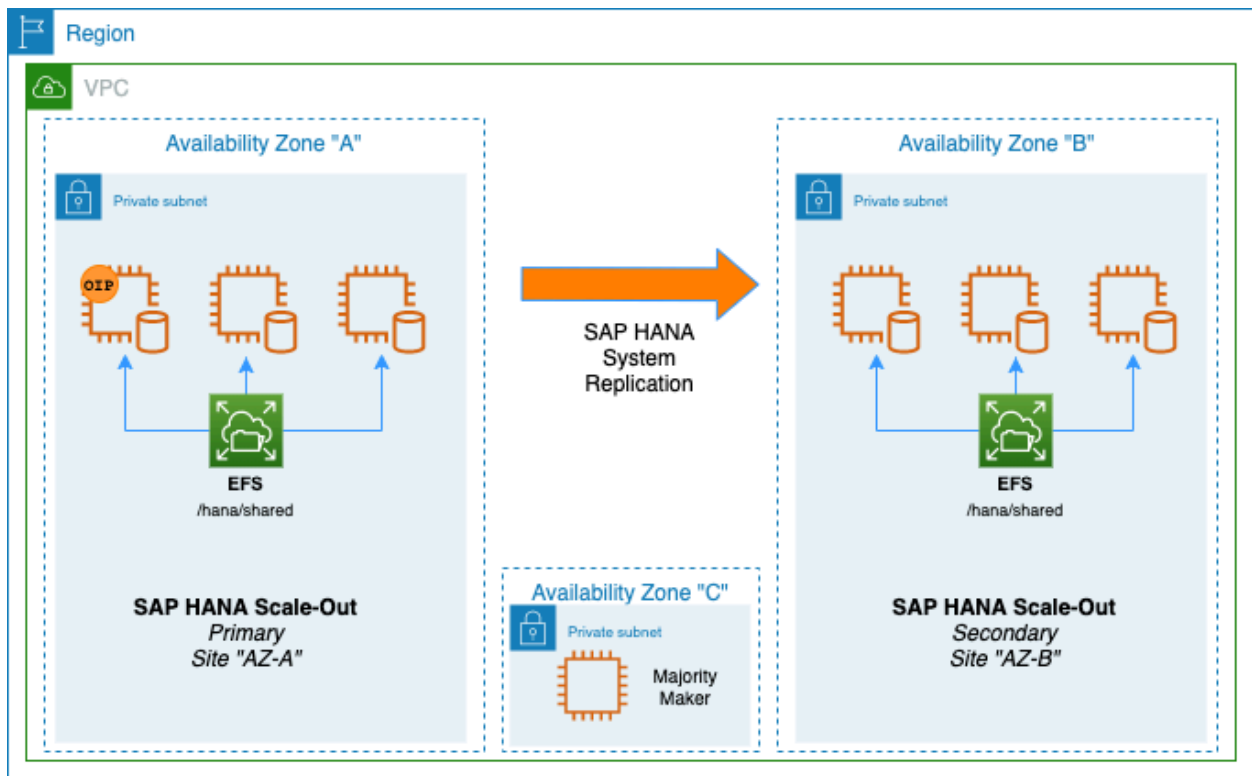


FIGURE 4: SIMPLIFIED CLUSTER ARCHITECTURE FOR SAP HANA SCALE-OUT SYSTEM REPLICATION ACROSS TWO AVAILABILITY ZONES

As an example, this guide will detail the implementation of a SUSE cluster composed of two 3-node SAP HANA Scale-Out clusters (one per Availability Zone), plus a Majority Maker (mm) node. But apart from the number of EC2 instances, all other requirements should be the same for different numbers of SAP HANA nodes:

- 3 HANA certified EC2 instances (bare metal or virtualized) in AZ-a
- 3 HANA certified EC2 instances (bare metal or virtualized) in AZ-b
- 1 EC2 instance, to be used as cluster Majority Maker, with at least 2 vCPUs, 2 GB RAM and 50 GB disk space in AZ-c
- 2 Amazon Elastic Filesystem (EFS) for /hana/shared
- 1 Overlay IP address for the primary (active) HANA System cluster

3.2 Parameter Sheet

Planning the cluster implementation can be very complex. Thus we recommend the installation to be planned properly. It is recommended to have all required parameters already in place before starting the deployment. It is a good practice to first fill out the parameter sheet and then begin with the installation.

TABLE 1: PARAMETER SHEET TO PREPARE THE NFS BASED SETUP

| Parameter | Value |
|-------------------------------------|-------|
| Path to SAP HANA media | |
| Node names AZ-a | |
| Node names AZ-b | |
| Node name majority maker | |
| IP addresses of all cluster nodes | |
| SAP HANA SID | |
| SAP Instance number | |
| Overlay IP address | |
| AWS Route Table | |
| SAP HANA site name site 1 | |
| SAP HANA site name site 2 | |
| EFS file system AZ-a (/hana/shared) | |
| EFS file system AZ-b (/hana/shared) | |
| Watchdog driver | |
| Placement Group Name | |

3.3 SAP Scale-Out Scenario in AWS with SUSE Linux Enterprise High Availability Extension

An SAP HANA Scale-Out cluster in AWS requires the use of a majority maker node in a 3rd Availability Zone. A majority maker, also known as tie-breaker node, will ensure the cluster quorum is maintained in case of the loss of one Availability Zone. In AWS, to maintain functionality of the Scale-Out cluster, at least all nodes in one Availability Zone plus the majority maker node need to be running. Otherwise the cluster quorum will be lost and any remaining SAP HANA node will be automatically rebooted.

It is also recommended that each set of cluster nodes has its own [EC2 placement group \(https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html\)](https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html) using "cluster" mode. This is needed to ensure that nodes can achieve the low-latency and high-throughput network performance needed for node-to-node communication required by SAP HANA in an Scale-Out deployment.

To automate the failover, the SUSE Linux Enterprise High Availability Extension (HAE) built into *SUSE Linux Enterprise Server for SAP Applications* is used. It includes two resource agents which have been created to support SAP HANA Scale-Out High Availability.

The first resource agent (RA) is **SAPHanaController**, which checks and manages the SAP HANA database instances. This RA is configured as a master/slave resource.

The master assumes responsibility for the active master name server of the SAP HANA database running in primary mode. All other instances are represented by the slave mode.

The second resource agent is **SAPHanaTopology**. This RA has been created to make configuring the cluster as simple as possible. It runs on all SAP HANA nodes (except the majority maker) of a SUSE Linux Enterprise High Availability Extension 12 cluster. It gathers information about the status and configuration of the SAP HANA system replication. It is designed as a normal (stateless) clone resource.

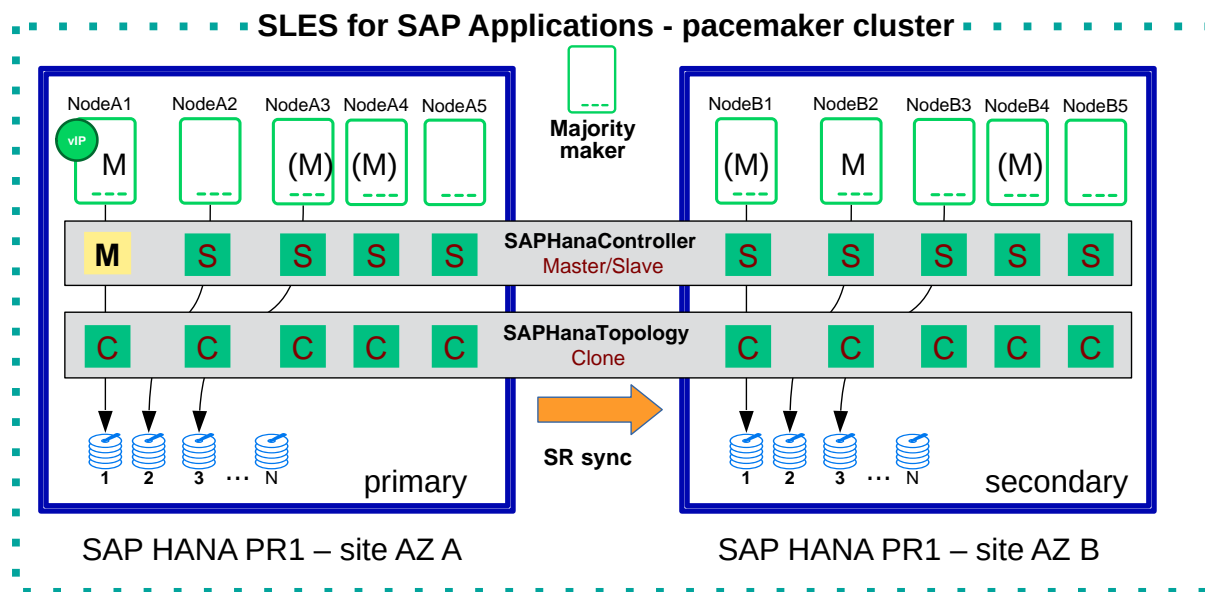


FIGURE 5: CLUSTER RESOURCE AGENTS AND MASTER/SLAVE STATUS MAPPING

SAP HANA system replication for Scale-Out is supported in the following scenarios or use cases:

Performance optimized, single container (A > B)

In the performance optimized scenario, an SAP HANA RDBMS on site "A" is synchronizing with an SAP HANA RDBMS on a second site "B". As the SAP HANA RDBMS on the second site is configured to preload the tables, the takeover time is typically very short.

Performance optimized, multi-tenancy also named MDC (%A > %B)

Multi-tenancy is supported for all above scenarios and use cases. This scenario is supported since SAP HANA 1 SPS12. The setup and configuration from a cluster point of view is the same for multi-tenancy and single containers. Thus you can use the above documents for both kinds of scenarios.

Multi-tenancy is the default installation type for SAP HANA 2.0.

3.4 The Concept of the Performance Optimized Scenario

In case of a failure of the primary SAP HANA Scale-Out cluster on AZ-a, the High Availability Extension tries to start the takeover process. This allows to use the already loaded data at the SAP HANA Scale-Out located in AZ-b. Typically the takeover is much faster than the local restart.

A site is noticed as "down" or "on error" if the **LandscapeHostConfiguration status** reflects this (return code 1). This happens when the worker nodes are going down without any local SAP HANA standby nodes left.

Without any additional intervention the resource agent will wait for the SAP internal HA cluster to repair the situation locally. An additional intervention could be a custom python hook using the SAP provider **srServiceStateChanged()** available since SAP HANA 2.0 SPS01.

To achieve an automation of this resource handling process, we can use the SAP HANA resource agents included in the SAPHanaSR-ScaleOut RPM package delivered with SUSE® Linux Enterprise Server for SAP Applications.

You can configure the level of automation by setting the parameter **AUTOMATED_REGISTER**. If automated registration is activated, the cluster will also automatically register a former failed primary to become the new secondary.

3.5 Important Prerequisites

The *SAPHanaSR-ScaleOut* resource agent software package supports Scale-Out (multiple-node to multiple-node) system replication with the following configurations and parameters:

- The cluster must include a valid STONITH method; in AWS the STONITH mechanism used is diskless SBD with watchdog.
- Since HANA primary and secondary reside in different [Availability Zones \(AZs\)](https://aws.amazon.com/about-aws/global-infrastructure/regions_az/) (https://aws.amazon.com/about-aws/global-infrastructure/regions_az/), an [Overlay IP](https://docs.aws.amazon.com/es_es/sap/latest/sap-hana/sap-ha-overlay-ip.html) (https://docs.aws.amazon.com/es_es/sap/latest/sap-hana/sap-ha-overlay-ip.html) address is required.
- Linux users and groups, such as `<sid>adm`, are defined **locally** in the Linux operating system.
- Time synchronization of all nodes relies on Amazon's Time Sync Service by default.
- SAP HANA Scale-Out groups in different Availability Zones must have the same SAP Identifier (SID) and instance number.
- EC2 instances must have different host names.
- The SAP HANA Scale-Out system must only have **one** active master name server per site. It should have up to three master name server candidates (SAP HANA nodes with a configured role 'master <N >').

- The SAP HANA Scale-Out system must only have **one** failover group.
- The cluster described in this document does not manage any service IP address for a read-enabled secondary site.
- There is only one SAP HANA system replication setup - from AZ-a to AZ-b.
- The setup implements the performance optimized scenario but not the cost optimized scenario.
- The saphostagent must be running on all SAP HANA nodes, as it is needed to translate between the system node names and SAP host names used during the installation of SAP HANA.
- The replication mode should be either 'sync' or 'syncmem'.
- All SAP HANA instances controlled by the cluster must not be activated via *sapinit* autostart.



Warning

Automated registration of a failed primary after takeover is possible. But as a good starting configuration, it is recommended to **switch off** the automated registration of a failed primary. Therefore `AUTOMATED_REGISTER = "false"` is set by **default**.

In this case, you need to manually register a failed primary after a takeover. Use SAP tools like **hanastudio** or **hdbnsutil**.

- Automated start of SAP HANA instances during system boot must be switched **off** in any case.
- You need at least SAPHanaSR-ScaleOut version 0.161, "*SUSE Linux Enterprise Server for SAP Applications 12 SP4*" or "*SUSE Linux Enterprise Server for SAP Applications 15*" and SAP HANA 1.0 SPS12 (122) or SAP HANA 2.0 SPS03 for all mentioned setups. Refer to SAP Note 2235581.



Important

You must implement a valid STONITH method. Without a valid STONITH method, the complete cluster is unsupported and will not work properly. For the document at hand, diskless SBD is used as STONITH.

This setup guide focuses on the performance optimized setup as it is the only supported scenario at the time of writing.

If you need to implement a different scenario, or customize your cluster configuration, it is strongly recommended to define a Proof-of-Concept (PoC) with SUSE and AWS. This PoC will focus on testing the existing solution in your scenario.

4 Using AWS with SUSE Linux Enterprise High Availability Extension Clusters

The SUSE Linux Enterprise High Availability Extension cluster will be installed in an AWS Region. An AWS Region consists of multiple Availability Zones (AZs). An Availability Zone (AZ) is one or more discrete data centers with redundant power, networking, and connectivity in an AWS Region. AZs give customers the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center. All AZs in an AWS Region are interconnected with high-bandwidth, low-latency networking, over fully redundant, dedicated metro fiber providing high-throughput, low-latency networking between AZs. All traffic between AZs is encrypted. The network performance is sufficient to accomplish synchronous replication between AZs. AZs are physically separated by a meaningful distance, many kilometers, from any other AZ, although all are within 100 km (60 miles) of each other. AWS recommends [architectural patterns \(https://d0.awsstatic.com/enterprise-marketing/SAP/sap-hana-on-aws-high-availability-disaster-recovery-guide.pdf\)](https://d0.awsstatic.com/enterprise-marketing/SAP/sap-hana-on-aws-high-availability-disaster-recovery-guide.pdf) ⁷ where redundant cluster nodes are being spread across different Availability Zones to overcome individual Availability Zones failures.

An AWS Virtual Private Network (VPC) is spanning all Availability Zones. We assume that a customer will have:

- Identified 3 Availability Zones (AZs) to be used
- Created subnets in the 3 AZs used to host the cluster nodes
- A routing table attached to the subnets

The virtual IP address for the HANA services will be an [Overlay IP address \(https://docs.aws.amazon.com/es_es/sap/latest/sap-hana/sap-ha-overlay-ip.html\)](https://docs.aws.amazon.com/es_es/sap/latest/sap-hana/sap-ha-overlay-ip.html) ⁷. This is a specific routing entry which can send network traffic to an instance, no matter which Availability Zones (and subnet) the instance is located in.

The cluster will update this routing entry as required. All SAP system components in the VPC can reach an AWS instance with an SAP system component inside a VPC through this Overlay IP address.

Overlay IP addresses have one requirement, they need to have a CIDR range outside of the VPC. Otherwise they would be part of a subnet and a given Availability Zone.

On premises users like HANA Studio cannot reach this IP address since the AWS Virtual Private Network (VPN) gateway will not route traffic to such an IP address.

4.1 AWS Environment Configurations

Here are the prerequisites which need to be met before starting the installation in AWS:

- Have an AWS account
- Have an AWS user with administrator permissions, or with the below permissions:
 - Create security groups
 - Modify AWS routing tables
 - Create policies and attach them to IAM roles
 - Enable/Disable EC2 instances' Source/Destination Check
 - Create placement groups
- Understand your landscape:
 - Know your AWS Region and its AWS name
 - Know your VPC and its AWS VPC ID
 - Know which Availability Zones (AZs) you will use
 - Have a subnet in each of the Availability Zones:
 - Have a routing table which is implicitly or explicitly attached to the subnets
 - Have free IP addresses in the subnets for your SAP installation
 - Allow network traffic in between the subnets
 - Allow outgoing Internet access from the subnets



Note

Using AWS SAP HANA QuickStart will automatically deploy all the required AWS resources listed above: This is the quickest and safest method to ensure all applicable SAP Notes and configurations are applied to the AWS resources.

4.2 Security Groups

The following ports and protocols need to be configured to allow the cluster nodes to communicate with each other:

- Port 5405 for inbound UDP: It is used to configure the corosync communication layer. Port 5405 is being used in common examples. A different port may be used depending on the corosync configuration.
- Port 7630 for inbound TCP: It is used by the SUSE "HAWK" Web GUI.



Note

This section lists the ports which need to be available for the SUSE Linux Enterprise HAE cluster only. It does not include SAP related ports.

4.3 Placement Group

One cluster placement group per Availability Zone is required to ensure that the SAP HANA nodes will achieve the high network throughput required by SAP HANA. For more information about placement groups, refer to the AWS documentation at <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>.

4.4 AWS EC2 Instance Creation

Create all EC2 instances to configure the SUSE Linux Enterprise HAE cluster. The EC2 instances will be located in 3 different Availability Zones to make them independent of each other.

This document will cover 3 EC2 instances in AZ-a for SAP HANA (Primary site), 3 EC2 instances in AZ-b for SAP HANA (Secondary site), and 1 instance in AZ-c as cluster Majority Maker (MM).

AMI selection:

- Use a "SUSE Linux Enterprise Server for SAP" AMI. Search for "suse-sles-sap-12-sp4" or "suse-sles-sap-15" in the list of AMIs. There are several BYOS (Bring Your Own Subscription) AMIs available. Use these AMIs if you have a valid SUSE subscription. Register your system with the Subscription Management Tool (SMT) from SUSE, SUSE Manager or directly with the SUSE Customer Center!
- Use the AWS Marketplace AMI *SUSE Linux Enterprise Server for SAP Applications 15* which already includes the SUSE subscription and the HAE software components.

Launch all EC2 instances into the Availability Zones (AZ) specific subnets, and placement groups. The subnets need to be able to communicate with each other.

4.5 Host Names

By default, the EC2 instances will have automatically generated host names. But it is recommended to assign host names that comply with the SAP requirements. See SAP note 611361. You need to edit `/etc/cloud/cloud.cfg` for host names to persist:

```
preserve_hostname: true
```



Note

To learn how to change the default host name for an EC2 instance running SUSE Linux Enterprise, refer to the AWS' public documentation at <https://aws.amazon.com/premium-support/knowledge-center/linux-static-hostname-suse/>.

4.6 AWS CLI Profile Configuration

The cluster's resource agents use the AWS Command Line Interface (CLI). They will use an AWS CLI profile which needs to be created for the `root` user on all instances. The SUSE cluster resource agents require a profile which generates output in text format.

The name of the profile is arbitrary, and will be added later to the cluster configuration. The name chosen in this example is `cluster`. The region of the instance needs to be added as well. Replace the string `region-name` with your target region in the following example.

One way to create such a profile is to create a file `/root/.aws/config` with the following content:

```
[default]
region = region-name
[profile cluster]
region = region-name
output = text
```

Another method is to use the `aws configure` CLI command in the following way:

```
# aws configure
AWS Access Key ID [None]:
AWS Secret Access Key [None]:
Default region name [None]: _region-name_
Default output format [None]:

# aws configure --profile cluster
AWS Access Key ID [None]:
AWS Secret Access Key [None]:
Default region name [None]: region-name
Default output format [None]: text
```

The above commands will create two profiles: a default profile and a cluster profile.



Note

AWS recommends *NOT* to store any AWS user credentials nor API signing keys in these profiles. Leave these fields blank and attach an EC2 IAM profile with the required permissions to the EC2 instance.

4.7 Configure HTTP/HTTPS Proxies (Optional)

This action is not needed if the system has direct access to the Internet.

Since the cluster resource agents will execute AWS API calls throughout the cluster lifecycle, they need HTTP/HTTPS access to AWS API endpoints. Systems which do not offer transparent Internet access may require an HTTP/HTTPS proxy. The configuration of the proxy access is described in full detail in the AWS documentation.

Add the following environment variables to the root user's `.bashrc` file:

```
export HTTP_PROXY=http://a.b.c.d:n
export HTTPS_PROXY=http://a.b.c.d:m
```

```
export NO_PROXY="169.254.169.254"
```

AWS' Data Provider for SAP will need to reach the instance meta data service directly. Add the following environment variable to the root user's `.bashrc` file:

```
export NO_PROXY="127.0.0.1,localhost,localaddress,169.254.169.254"
```

SUSE Linux Enterprise HAE also requires to add the proxy configurations to `/etc/sysconfig/pacemaker` configuration file in the following format:

```
export HTTP_PROXY=http://username:password@a.b.c.d:n
export HTTPS_PROXY=http://username:password@a.b.c.d:m
export NO_PROXY="127.0.0.1,localhost,localaddress,169.254.169.254"
```

4.7.1 Verify HTTP Proxy Settings (Optional)

Make sure that the SUSE instance can reach the EC2 instance metadata address URL <http://169.254.169.254/latest/meta-data>, as multiple system components will required to access it. Therefore it is recommended to disable any firewall rules that restrict it, and to disable proxy access to this URL.

For more information about EC2 Instance metadata server, refer to AWS' documentation at <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-metadata.html>.

4.8 Disable the Source/Destination Check for the Cluster Instances

The source/destination check needs to be disabled on all EC2 instances that are part of the cluster. This can be done through scripts using the AWS command line interface (AWS-CLI) or by using the AWS console. The following command needs to be executed *one time* only on all EC2 instances part of the cluster:

```
# aws ec2 modify-instance-attribute --instance-id EC2-instance-id --no-source-dest-check
```

Replace the variable `EC2-instance-id` with the instance ID of the AWS EC2 instances. The system on which this command gets executed needs temporarily a role with the following policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
```

```

    "Sid": "Stmt1424870324000",
    "Effect": "Allow",
    "Action": [ "ec2:ModifyInstanceAttribute" ],
    "Resource": [
      "arn:aws:ec2:region-name:account-id:instance/instance-a",
      "arn:aws:ec2:region-name:account-id:instance/instance-b"
    ]
  }
]
}

```

Replace the following individual parameter with the appropriate values:

- region-name (Example: us-east-1)
- account-id (Example: 123456789012)
- instance-a, instance-b (Example: i-1234567890abcdef)



Note

The string "instance" in the policy is a fixed string. It is not a variable which needs to be substituted!

The source/destination check can be disabled as well from the AWS console. It takes the execution of the following drop-down box in the console for both EC2 instances (see below).




FIGURE 6: DISABLE SOURCE/DESTINATION CHECK AT CONSOLE

4.9 Avoid Deletion of Overlay IP Address on the eth0 Interface

SUSE's `cloud-netconfig-ec2` package may erroneously remove any secondary IP address which is managed by the cluster agents from the `eth0` interface. This can cause service interruptions for users of the cluster service. Perform the following task on all cluster nodes:

Check whether the package `cloud-netconfig-ec2` is installed with the command:

```
# zypper info cloud-netconfig-ec2
```

If the package is installed, update the file `/etc/sysconfig/network/ifcfg-eth0` and change the following line to a "no" setting or add the line if the package is not yet installed:

```
CLOUD_NETCONFIG_MANAGE='no'
```

4.10 AWS Roles and Policies

SUSE Linux Enterprise HAE cluster software and its agents need several AWS IAM privileges to operate the cluster. An IAM Security Role is required to be attached to the EC2 instance that are part of the cluster. A single IAM Role can be used across the cluster, and associated to all EC2 instances.

This IAM Security Role will require the IAM Security Policies detailed below.

4.10.1 AWS Data Provider Policy

SAP systems on AWS require the installation of the "AWS Data Provider for SAP", which needs a policy to access AWS resources. Use the policy shown in the "AWS Data Provider for SAP Installation and Operations Guide" (<https://docs.aws.amazon.com/sap/latest/general/aws-data-provider.html>) and attach it to the IAM Security Role to be used by the cluster EC2 instance. This policy can be used by all SAP systems. Only one policy per AWS account is required for "AWS Data Provider for SAP". Therefore, if an IAM Security Policy for "AWS Data Provider for SAP" already exists, it can be used.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
```

```

        "EC2:DescribeInstances",
        "EC2:DescribeVolumes"
    ],
    "Resource": "*"
  },
  {
    "Effect": "Allow",
    "Action": "cloudwatch:GetMetricStatistics",
    "Resource": "*"
  },
  {
    "Effect": "Allow",
    "Action": "s3:GetObject",
    "Resource": "arn:aws:s3:::aws-data-provider/config.properties"
  }
]
}

```

4.10.2 Overlay IP Agent Policy

The Overlay IP agent will change a routing entry in an AWS routing table. Create a policy with a name like *Manage-Overlay-IP-Policy* and attach it to the IAM Security Role of the cluster instances:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1424870324000",
      "Action": "ec2:ReplaceRoute",
      "Effect": "Allow",
      "Resource": "arn:aws:ec2:region-name:account-id:route-table/rtb-XYZ"
    },
    {
      "Sid": "Stmt1424870324000",
      "Action": "ec2:DescribeRouteTables",
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}

```

This policy allows the agent to update the routing tables which get used. Replace the following variables with the appropriate names:

- `region-name` : the name of the AWS region
- `account-id` : The name of the AWS account in which the policy is getting used
- `rtb-XYZ` : The identifier of the routing table which needs to be updated

4.11 Add Overlay IP Addresses to Routing Table

Manually create a route entry in the routing table which is assigned to the two subnets used by the EC2 cluster instances. This IP address is the virtual service IP address of the HANA cluster. The Overlay IP address needs to be outside of the CIDR range of the VPC. Use the AWS console and search for “VPC”.

- Select VPC
- Click “Route Tables” in the left column
- Select route table used the subnets from one of your SAP EC2 instances and their application servers
- Click the tabulator “Routes”
- Click “Edit”
- Scroll to the end of the list and click “Add another route”

Add the overlay IP address of the HANA database. Use as filter /32 (example: 192.168.10.1/32). Add the Elastic Network Interface (ENI) name to EC2 instance to be the SAP HANA Master. The resource agent will modify this latter one automatically as required. Save your changes by clicking “Save”.



Note

It is important that the routing table, which will contain the routing entry, needs to be inherited to all subnets in the VPC which have consumers of the service. Check the AWS VPC documentation at http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Introduction.html for more details on routing table inheritance.

4.12 Create EFS File Systems

Each set of SAP HANA Scale-Out clusters needs to have its own Amazon Elastic Filesystem (EFS). To create an EFS file system review the AWS public documentation which contains a step-by-step guide of how to create and mount it (see <https://docs.aws.amazon.com/efs/latest/ug/getting-started.html>).

4.13 Configure the Operating System for SAP HANA

Consider these SAP notes to configure the operating system (modules, packages, kernel settings, etc.) for your version of SAP HANA:

If using SUSE Linux Enterprise Server for SAP Applications 12: - **1984787** SUSE LINUX Enterprise Server 12: Installation Notes - **2205917** SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12

If using SUSE Linux Enterprise Server for SAP Applications 15: - **2578899** SUSE Linux Enterprise Server 15: Installation Notes - **2684254** SAP HANA DB: Recommended OS settings for SLES 15 / SLES for SAP Applications 15

Other related SAP Notes: - **1275776** Linux: Preparing SLES for SAP environments - **2382421** Optimizing the Network Configuration on HANA- and OS-Level

4.13.1 Install SAP HANA Scale-Out Cluster Agent

SUSE delivers with SLES for SAP special resource agents for SAP HANA. With the installation of pattern "sap-hana" the resource agent for SAP HANA Scale-Up is installed, but for the Scale-Out scenario we need a special resource agent. Follow these instructions on each node if you have installed the systems based on the existing AWS Amazon Machine Images (AMIs), or deployed the SAP HANA nodes using the AWS QuickStart. The pattern *High Availability* summarizes all tools what we recommend to install on **all** nodes including the majority maker node.

- Remove packages: SAPHanaSR SAPHanaSR-doc yast2-sap-ha
- Install packages: SAPHanaSR-ScaleOut SAPHanaSR-ScaleOut-doc

EXAMPLE 1: UNINSTALL THE SAPHANASR AGENT FOR SCALE-UP

As user root:

```
zypper remove SAPHanaSR SAPHanaSR-doc yast2-sap-ha
```


EXAMPLE 2: INSTALL THE SAPHANASR AGENT FOR SCALE-OUT

As user root:

```
zypper in SAPHanaSR-ScaleOut SAPHanaSR-ScaleOut-doc
```

If the package is not installed yet. You should get an output like:

```
Refreshing service 'Advanced_Systems_Management_Module_12_x86_64'.
Refreshing service
'SUSE_Linux_Enterprise_Server_for_SAP_Applications_12_SP3_x86_64'.
Loading repository data...
Reading installed packages...
Resolving package dependencies...

The following 2 NEW packages are going to be installed:
  SAPHanaSR-ScaleOut SAPHanaSR-ScaleOut-doc

2 new packages to install.
Overall download size: 539.1 KiB. Already cached: 0 B. After the operation,
additional 763.1 KiB will be used.
Continue? [y/n/...? shows all options] (y): y
Retrieving package SAPHanaSR-ScaleOut-0.161.1-1.1.noarch
                                                                    (1/2), 48.7 KiB (211.8 KiB
unpacked)
Retrieving: SAPHanaSR-
ScaleOut-0.161.1-1.1.noarch.rpm .....[done]
Retrieving package SAPHanaSR-ScaleOut-doc-0.161.1-1.1.noarch
                                                                    (2/2), 490.4 KiB (551.3 KiB
unpacked)
Retrieving: SAPHanaSR-ScaleOut-
doc-0.161.1-1.1.noarch.rpm .....[done (48.0 KiB/s)]
Checking for file
  conflicts: .....[done]
(1/2) Installing: SAPHanaSR-
ScaleOut-0.161.1-1.1.noarch ..... [done]
(2/2) Installing: SAPHanaSR-ScaleOut-
doc-0.161.1-1.1.noarch .....[done]
```

Install SUSE's High Availability Pattern

```
zypper in --type pattern ha_sles
```

4.13.2 Install the Latest Available Updates from SUSE

If you have the packages installed before, make sure to install the latest package updates on **all** machines to have the latest versions of the resource agents and other packages. There are multiple ways to get updates like SUSE Manager, SMT, or directly connected to SCC (SUSE Customer Center).

Depending on your company or customer rules use *zypper update* or *zypper patch*.

EXAMPLE 3: SOFTWARE UPDATE MUST BE TRIGGERED FROM EACH NODE

Zypper patch will install all available needed patches. As user root:

```
zypper patch
```

Zypper update will update all or specified installed packages with newer versions, if possible. As user root:

```
zypper update
```

4.14 Configure SLES for SAP to Run SAP HANA

4.14.1 Tuning / Modification

All needed operating system tuning are described in SAP Note 2684254 and in SAP Note 2205917. It is recommended to manually verify each parameter mentioned in the SAP Notes. This is to ensure all performance tunings for SAP HANA are correctly set.

The SAP note covers:

- SLES 15 or SLES 12
- Additional 3rd-party kernel modules
- Configure sapconf or saptune
- Turn off NUMA balancing
- Disable transparent hugepages
- Configure C-States for lower latency in Linux (applies to Intel-based systems only)
- CPU Frequency/Voltage scaling (applies to Intel-based systems only)
- Energy Performance Bias (EPB, applies to Intel-based systems only)

- Turn off kernel samepage merging (KSM)
- Linux Pagecache Limit

4.14.2 Enabling SSH access via public key (optional)

Public key authentication provides SSH users access to their servers without entering their passwords. SSH keys are also more secure than passwords, because the private key used to secure the connection is never shared. Private keys can also be encrypted. Their encrypted contents cannot easily be read. For the document at hand, a very simple but useful setup is used. This setup is based on only one ssh-key pair which enables SSH access to all cluster nodes.



Note

Follow your company security policy to set up access to the systems.

EXAMPLE 4: SSH KEY CREATION AND EXCHANGE

As user root create an SSH key on one node.

```
# ssh-keygen -t rsa
```

The ssh-key generation asks for missing parameters.

```
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
SHA256:ip/8kdTbYZNuuEUAdsYOAErkwnkAPBR7d2SQIpIZCU root@<host1>
The key's randomart image is:
+---[RSA 2048]-----+
|XEooo+ooo+o      |
|=+.= o=.o+.      |
|..B o. + o.      |
|  o  . +... .    |
|          S.. *   |
|   . o . B o     |
|   . . o o =     |
|   o . . +       |
|   +.. .         |
+-----[SHA256]-----+
```

After the `ssh-keygen` is set up, you will have two new files under `/root/.ssh/`.

```
# ls /root/.ssh/  
id_rsa id_rsa.pub
```

To allow password-free access for the user `root` between nodes in the cluster copy `id_rsa.pub` to `authorized_keys` and set the required permissions.

```
# cp /root/.ssh/id_rsa.pub /root/.ssh/authorized_keys  
# chmod 600 /root/.ssh/authorized_keys
```

Collect the public host keys from all other node. For the document at hand, the `ssh-keyscan` command is used.

```
# ssh-keyscan
```

The SSH host key is automatically collected and stored in the file `/root/.ssh/known_host` during the first SSH connection. To avoid to confirm the first login with "yes" which accept the host key, we collect and store them beforehand.

```
# ssh-keyscan -t ecdsa-sha2-nistp256 <host1>,<host1 ip> >>.ssh/known_hosts  
# ssh-keyscan -t ecdsa-sha2-nistp256 <host2>,<host2 ip> >>.ssh/known_hosts  
# ssh-keyscan -t ecdsa-sha2-nistp256 <host3>,<host3 ip> >>.ssh/known_hosts  
...
```

After collecting all host keys push the entire directory `/root/.ssh/` from the first node to all further cluster members.

```
# rsync -ay /root/.ssh/ <host2>:/root/.ssh/  
# rsync -ay /root/.ssh/ <host3>:/root/.ssh/  
# rsync -ay /root/.ssh/ <host4>:/root/.ssh/  
...
```

4.14.3 Set up Disk Layout for SAP HANA

We highly recommend to follow the storage layout described at <https://docs.aws.amazon.com/quickstart/latest/sap-hana/planning.html>. The table on this Web site lists the minimum required number of EBS volumes, volume size and IOPS (for IO1) for your desired EC2 instance type. You can choose more storage or more IOPS depending on your workload's requirements. Configure the EBS volumes for:

- `/hana/data`
- `/hana/log`

File systems:

`/hana/shared/<SID>`

On SAP HANA Scale-Out this directory is mounted on all nodes of the same Scale-Out cluster, and in AWS this directory uses EFS. It contains shared files, like binaries, trace, and log files.

`/hana/log/<SID>`

This directory contains the redo log files of the SAP HANA host. On AWS this directory is local to the instance.

`/hana/data/<SID>`

This directory contains the data files of the SAP HANA host. On AWS this directory is local to the instance.

`/usr/sap`

This is the path to the local SAP system instance directories. It is recommended to have a separate volume for `/usr/sap`.

4.15 Configure Host Name Resolution

To configure host name resolution, you can either use a DNS server or modify the `/etc/hosts` on **all** cluster nodes.

With maintaining the `/etc/hosts` file, you minimize the impact of a failing DNS service. Edit the `/etc/hosts` file on **all** cluster nodes and add all cluster nodes' host name and IPs to it.

```
vi /etc/hosts
```

Insert the following lines to `/etc/hosts`. Change the IP address and host name to match your environment.

```
192.168.201.151 suse01
192.168.201.152 suse02
...
```

4.16 Enable Chrony/NTP Service on All Nodes

By default all nodes should automatically synchronize with Amazon Time Sync Service. Check the NTP/chrony configuration `/etc/chrony.conf` of all nodes and confirm that the time source server is 169.254.169.123

5 Installing the SAP HANA Databases on both sites

As now the infrastructure is set up, we can install the SAP HANA database on both sites. In a cluster a machine is also called a *node*.




FIGURE 7: SECTION 3, "PLANNING THE INSTALLATION" SECTION 4, "USING AWS WITH SUSE LINUX ENTERPRISE HIGH AVAILABILITY EXTENSION CLUSTERS" SAPHANAINST SECTION 6, "SET UP THE SAP HANA SYSTEM REPLICATION" SECTION 7, "INTEGRATION OF SAP HANA WITH THE CLUSTER" SECTION 8, "CONFIGURATION OF THE CLUSTER AND SAP HANA RESOURCES" SECTION 9, "TESTING THE CLUSTER"

In our example here and to make it more easy to follow the documentation, we name the machines (or nodes) *suse01*, ... *suseXX*. The nodes with odd numbers (*suse01*, *suse03*, *suse05*, ...) will be part of site "A" (WDF1) and the nodes with even (*suse02*, *suse04*, *suse06*, ...) will be part of site "B" (ROT1) .

The following users are automatically created during the SAP HANA installation:

<sid>adm

The user *<sid> adm* is the operating system user required for administrative tasks, such as starting and stopping the system.

sapadm

The SAP Host Agent administrator.

SYSTEM

The SAP HANA database superuser

5.1 Preparation

- Read the SAP Installation and Setup guides available at the SAP Marketplace.
- Download the SAP HANA Software from SAP Marketplace.
- Mount the file systems to install SAP HANA database software and database content (data, log and shared).

5.2 Installation

Install the SAP HANA Database as described in the SAP HANA Server Installation Guide. To do this install the HANA primary master server first as a single scale-up system. Once that is done change the `global.ini` parameter `persistence/basepath_shared` to "no".

```
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET
('persistence', 'basepath_shared')='no';
```

This way the HANA database will not expect shared log/data directories across all nodes. After this setting is applied you can add hosts to the database. Add all HANA nodes of the scale-out cluster within the same Availability Zone (primary site).

Follow SAP note **2369981 - Required configuration steps for authentication with HANA System Replication** to exchange the encryption keys with the secondary site!

Now repeat the same procedure to install the SAP HANA database on the master of the secondary site. Change the `persistence/basepath_shared` parameter and add nodes to the secondary scale-out cluster.

5.3 Checks

Verify that **both** database sites are up and all processes of these databases are running correctly.

1. As Linux user `<sid>adm` use the SAP command line tool `HDB` to get an overview of running SAP HANA processes. The output of `HDB info` should look similar to the following screenshot for **both** sites:

EXAMPLE 5: CALLING HDB INFO (AS USER <SID>ADM)

```
HDB info
```

The `HDB info` command lists the processes currently running for that SID.

```
USER      PID  ...  COMMAND
haladm    6561 ...  -csh
haladm    6635 ...  \_ /bin/sh /usr/sap/HA1/HDB00/HDB info
haladm    6658 ...  \_ ps fx -U HA1 -o
          user,pid,ppid,pcpu,vsz,rss,args
haladm    5442 ...  sapstart pf=/hana/shared/HA1/profile/HA1_HDB00_suse01
haladm    5456 ...  \_ /usr/sap/HA1/HDB00/suse01/trace/hdb.sapha1_HDB00
          -d -nw -f /usr/sap/ha1/HDB00/suse
haladm    5482 ...  \_ hdbnameserver
```

```

haladm      5551 ...      \_ hdbpreprocessor
haladm      5554 ...      \_ hdbcompileserver
haladm      5583 ...      \_ hdbindexserver
haladm      5586 ...      \_ hdbstatisticsserver
haladm      5589 ...      \_ hdbxsengine
haladm      5944 ...      \_ sapwebdisp_hdb pf=/usr/sap/HA1/HDB00}/suse01/
wdisp/sapwebdisp.pfl -f /usr/sap/SL
haladm      5363 ...      /usr/sap/HA1/HDB00/exe/sapstartsrv pf=/hana/shared/
HA1/profile/HA1_HDB00_suse02 -D -u s

```

2. Use the python script *landscapeHostConfiguration.py* to show the status of an entire SAP HANA site.

EXAMPLE 6: QUERY THE HOST ROLES (AS USER <SID>ADM)

```
HDBSettings.sh landscapeHostConfiguration.py
```

The landscape host configuration is shown with a line per SAP HANA host.

| Host | Host | ... NameServer | NameServer | IndexServer | IndexServer |
|--------|--------|-----------------|-------------|-------------|-------------|
| | Active | ... Config Role | Actual Role | Config Role | Actual Role |
| ----- | ----- | ... ----- | ----- | ----- | ----- |
| suse01 | yes | ... master 1 | master | worker | master |
| suse03 | yes | ... master 2 | slave | worker | slave |
| suse05 | yes | ... master 3 | slave | standby | standby |

overall host status: ok

3. Get an overview of instances of that site (as user <sid> adm)

EXAMPLE 7: GET THE LIST OF INSTANCES

```
sapcontrol -nr <Inst> -function GetSystemInstanceList
```

You should get a list of SAP HANA instances belonging to that site.

```

12.06.2018 17:25:16
GetSystemInstanceList
OK
hostname, instanceNr, httpPort, httpsPort, startPriority, features, dispstatus
suse01, 00, 50013, 50014, 0.3, HDB|HDB_WORKER, GREEN
suse05, 00, 50013, 50014, 0.3, HDB|HDB_WORKER, GREEN
suse03, 00, 50013, 50014, 0.3, HDB|HDB_WORKER, GREEN

```


6 Set up the SAP HANA System Replication

This section describes the setup of the system replication (HSR) after SAP HANA has been installed properly.

Procedure

1. Back up the primary database
2. Enable primary database
3. Register the secondary database
4. Verify the system replication




FIGURE 8: SECTION 3, "PLANNING THE INSTALLATION" SECTION 4, "USING AWS WITH SUSE LINUX ENTERPRISE HIGH AVAILABILITY EXTENSION CLUSTERS" SECTION 5, "INSTALLING THE SAP HANA DATABASES ON BOTH SITES" SAPHANAHSR SECTION 7, "INTEGRATION OF SAP HANA WITH THE CLUSTER" SECTION 8, "CONFIGURATION OF THE CLUSTER AND SAP HANA RESOURCES" SECTION 9, "TESTING THE CLUSTER"

For more information read the Section *Setting Up System Replication* of the SAP HANA Administration Guide.

6.1 Back Up the Primary Database

Please, first back up the primary database as described in the *SAP HANA Administration Guide, Section SAP HANA Database Backup and Recovery*.

We provide some examples to back up SAP HANA with SQL Commands:

EXAMPLE 8: SIMPLE BACKUP FOR THE SYSTEM DATABASE AND ALL TENANTS WITH ONE SINGLE BACKUP CALL

As user <sid> adm enter the following command:

```
hdbsql -u SYSTEM -d SYSTEMDB \  
"BACKUP DATA FOR FULL SYSTEM USING FILE ('backup')"
```

You get the following command output (or similar):

```
0 rows affected (overall time 15.352069 sec; server time 15.347745 sec)
```

EXAMPLE 9: SIMPLE BACKUP FOR A SINGLE CONTAINER (NON MDC) DATABASE

Enter the following command as user `<sid> adm`:

```
hdbsql -i <Inst> -u <dbuser> \  
"BACKUP DATA USING FILE ('backup')"
```



Important

Without a valid backup, you **cannot** bring SAP HANA into a system replication configuration.

6.2 Enable Primary Database

As Linux user `<sid> adm` enable the system replication at the primary node. You need to define a site name (like `WDF1`) which must be unique for all SAP HANA databases which are connected via system replication. This means the secondary must have a different site name.

EXAMPLE 10: ENABLE THE SYSTEM REPLICATION ON THE PRIMARY SITE

As user `<sid> adm` enable the primary:

```
hdbnsutil -sr_enable --name=HAP
```

Check, if the command output is similar to:

```
nameserver is active, proceeding ...  
successfully enabled system as system replication source site  
done.
```

The command line tool `hdbnsutil` can be used to check the system replication mode and site name.

EXAMPLE 11: CHECK THE SYSTEM REPLICATION CONFIGURATION STATUS AS USER `<SID>ADM` ON THE PRIMARY

```
hdbnsutil -sr_stateConfiguration
```

If the system replication enablement was successful at the primary, the output should look like the following:

```
checking for active or inactive nameserver ...
```

```
System Replication State
~~~~~

mode: primary
site id: 1
site name: HAP
done.
```

The mode has changed from “none” to “primary”. The site now has a site name and a site ID.

6.3 Register the Secondary Database

The SAP HANA database instance on the secondary side must be stopped before the system can be registered for the system replication. You can use your preferred method to stop the instance (like *HDB* or *sapcontrol*). After the database instance has been stopped successfully, you can register the instance using *hdbnsutil*.

EXAMPLE 12: STOP THE SECONDARY AS LINUX USER <SID>ADM:

```
sapcontrol -nr <Inst> -function StopSystem
```

EXAMPLE 13: REGISTER THE SECONDARY AS LINUX USER <SID>ADM:

```
hdbnsutil -sr_register --name=<site2> \
  --remoteHost=<node1-siteA> --remoteInstance=<Inst> \
  --replicationMode=syncmem --operationMode=logreplay
```

```
adding site ...
checking for inactive nameserver ...
nameserver suse02:30001 not responding.
collecting information ...
updating local ini files ...
done.
```

The *remoteHost* is the primary node in our case, the *remoteInstance* is the database instance number (here 00).

Now start the secondary database instance again and verify the system replication status. On the secondary site, the mode should be one of „SYNC“, „SYNCMEM“ or „ASYN““. The mode depends on the sync option defined during the registration of the secondary.

EXAMPLE 14: START THE SYSTEM ON THE SECONDARY SITE AS USER <SID>ADM

```
sapcontrol -nr <Inst> -function StartSystem
```

Wait till the SAP HANA database is started completely.

EXAMPLE 15: CHECK THE SYSTEM REPLICATION CONFIGURATION AS LINUX USER <SID>ADM

```
hdbnsutil -sr_stateConfiguration
```

The output should look like:

```
System Replication State
~~~~~
mode: sync
site id: 2
site name: HAS
active primary site: 1

primary masters: suse01 suse03 suse05
done.
```

6.4 Verify the System Replication

To view the replication state of the whole SAP HANA cluster, use the following command as <sid>adm user on the primary site.

EXAMPLE 16: CHECK THE SYSTEM REPLICATION STATUS AT THE PRIMARY SITE (AS <SID>ADM)

```
HDBSettings.sh systemReplicationStatus.py
```

This script prints a human readable table of the system replication channels and their status. The most interesting column is the **Replication Status**, which should be **ACTIVE**.

```
| Database | Host | .. Site Name | Secondary | .. Secondary | .. Replication
|          |     | ..          | Host      | .. Site Name | .. Status
| - - - - - | - - - - - | .. - - - - - | - - - - - | .. - - - - - | .. - - - - -
| SYSTEMDB | suse01 | .. HAP      | suse02    | .. HAS      | .. ACTIVE
| HA1      | suse01 | .. HAP      | suse02    | .. HAS      | .. ACTIVE
| HA1      | suse01 | .. HAP      | suse02    | .. HAS      | .. ACTIVE
| HA1      | suse03 | .. HAP      | suse04    | .. HAS      | .. ACTIVE

status system replication site "2": ACTIVE
overall system replication status: ACTIVE

Local System Replication State
~~~~~
mode: PRIMARY
```

```
site id: 1
site name: HAP
```

7 Integration of SAP HANA with the Cluster



FIGURE 9: SECTION 3, "PLANNING THE INSTALLATION" SECTION 4, "USING AWS WITH SUSE LINUX ENTERPRISE HIGH AVAILABILITY EXTENSION CLUSTERS" SECTION 5, "INSTALLING THE SAP HANA DATABASES ON BOTH SITES" SECTION 6, "SET UP THE SAP HANA SYSTEM REPLICATION" INTEGRATION SECTION 8, "CONFIGURATION OF THE CLUSTER AND SAP HANA RESOURCES" SECTION 9, "TESTING THE CLUSTER"

We need to proceed the following steps:

Procedure

1. Implement the python hook SAPHanaSR
2. Configure system replication operation mode
3. Allow <sid> adm to access the cluster
4. Start SAP HANA
5. Test the hook integration

7.1 Implement the Python Hook SAPHanaSR

SUSE's SAPHanaSR-ScaleOut resource agent includes an SAP HANA integration script to handle failures on the SAP HANA replication and prevent a cluster failover to an out of sync SAP HANA node and avoid data loss.

This integration script will watch SAP HANA's "srConnectionChanged" hook. The method srConnectionChanged() is called on the master name server when one of the replicating services loses or establishes the system replication connection and inform the cluster.

This step must be done on both sites. SAP HANA must be stopped to change the global.ini and allow SAP HANA to integrate the HA/DR hook script during start.

- Install the HA/DR hook script into a read/writable directory
- Integrate the hook into `global.ini` (SAP HANA needs to be stopped for doing that offline)
- Check integration of the hook during start-up

Take the hook from the `SAPHanaSR-ScaleOut` package and copy it to your preferred directory like `/hana/share/myHooks`. The hook must be available on all SAP HANA nodes.

```
suse01~ # mkdir -p /hana/shared/myHooks
suse01~ # cp /usr/share/SAPHanaSR-ScaleOut/SAPHanaSR.py /hana/shared/myHooks
suse01~ # chown -R <sid>adm:sapsys /hana/shared/myHooks
```

Stop SAP HANA

```
sapcontrol -nr <Inst> -function StopSystem
```

EXAMPLE 17: ADDING SAPHANASR VIA GLOBAL.INI

```
[ha_dr_provider_SAPHanaSR]
provider = SAPHanaSR
path = /hana/shared/myHooks
execution_order = 1

[trace]
ha_dr_saphanasr = info
```

7.2 Configure System Replication Operation Mode

When your system is connected as an SAPHanaSR target you can find an entry in the `global.ini` which defines the operation mode. Up to now there are two modes available.

- *delta_datashipping*
- *logreplay*

Until a takeover and re-registration in the opposite direction the entry for the operation mode is missing on your primary site. The "classic" operation mode is `delta_datashipping`. The preferred mode for HA is `logreplay`. Using the operation mode `logreplay` makes your secondary site in the SAP HANA system replication a HotStandby system. For more details regarding both modes check the available SAP documentation like "How To Perform System Replication for SAP HANA".

Check both `global.ini` files and add the operation mode, if needed.

section

```
[ system_replication ]
```

key

```
operation_mode = logreplay
```

Path for the *global.ini*: /hana/shared/ <SID> /global/hdb/custom/config/

```
[system_replication]
operation_mode = logreplay
```

7.3 Allow <sid>adm to Access the Cluster

The current version of the SAPHanaSR python hook uses the command 'sudo' to allow the <sid>adm user to access the cluster attributes. In Linux you can use 'visudo' to start the vi editor for the '/etc/sudoers' configuration file.

The user <sid>adm must be able to set the cluster attribute hana_<sid>_glob_srHook_*. The SAP HANA system replication hook needs password free access. The following example limits the sudo access to exactly setting the needed attribute.

Replace the <sid> > by the **lowercase** SAP system ID.

This change is required in all cluster nodes.

EXAMPLE 18: ENTRY IN SUDO PERMISSIONS /ETC/SUDOERS FILE

Basic parameter option to allow <sidadm> to use the srHook.

```
# SAPHanaSR-ScaleOut needs for srHook
<sid>adm ALL=(ALL) NOPASSWD: /usr/sbin/crm_attribute -n hana_<sid>_glob_srHook -v *
```

More specific parameters option to meet a high security level.

```
# SAPHanaSR-ScaleOut needs for srHook
Cmnd_Alias SOK = /usr/sbin/crm_attribute -n hana_<sid>_glob_srHook -v SOK -t
crm_config -s SAPHanaSR
Cmnd_Alias SFAIL = /usr/sbin/crm_attribute -n hana_<sid>_glob_srHook -v SFAIL -t
crm_config -s SAPHanaSR
<sid>adm ALL=(ALL) NOPASSWD: SOK, SFAIL
```

EXAMPLE 19: RESULT OF REPLACING <SID> WITH HA1

```
# SAPHanaSR-ScaleOut needs for srHook
hd1adm ALL=(ALL) NOPASSWD: /usr/sbin/crm_attribute -n hana_ha1_glob_srHook -v *
```

7.4 Start SAP HANA

After SAP HANA integration has been configured and the communication between SAP HANA is working the cluster can now start the SAP HANA databases on both sites.

EXAMPLE 20: STARTING A COMPLETE SAP HANA SITE AS USE <SID>ADM

```
sapcontrol -nr <Inst> -function StartSystem
```

The sapcontrol service commits the request with OK.

```
11.06.2018 18:30:16
StartSystem
OK
```

Check if SAP HANA has finished starting:

```
sapcontrol -nr <Inst> -function WaitForStarted 300 20
```

7.5 Test the Hook Integration

When the SAP HANA database has been restarted after the changes, check if the hook script is called correctly.

Check the SAP HANA trace files as <sid> adm:

```
suse01:haladm> cdtrace
suse01:haladm> awk '/ha_dr_SAPHanaS.*crm_attribute/ \
    { printf "%s %s %s %s\n", $2, $3, $5, $16 }' nameserver_suse01.*
2018-05-04 12:34:04.476445 ha_dr_SAPHanaS...SFMAIL
2018-05-04 12:53:06.316973 ha_dr_SAPHanaS...SOK
```

If you can spot "ha_dr_SAPHanaSR" messages the hook script is called and executed.

8 Configuration of the Cluster and SAP HANA Resources

This chapter describes the configuration of the _SUSE Linux Enterprise High Availability (SLE HA) cluster. The SUSE Linux Enterprise High Availability is part of the SUSE Linux Enterprise Server for SAP Applications. Further, the integration of SAP HANA System Replication with the

SUSE Linux Enterprise High Availability cluster is explained. The integration is done by using the SAPHanaSR-ScaleOut package which also is part of the SUSE Linux Enterprise Server for SAP Applications.




FIGURE 10: SECTION 3, "PLANNING THE INSTALLATION" SECTION 4, "USING AWS WITH SUSE LINUX ENTERPRISE HIGH AVAILABILITY EXTENSION CLUSTERS" SECTION 5, "INSTALLING THE SAP HANA DATABASES ON BOTH SITES" SECTION 6, "SET UP THE SAP HANA SYSTEM REPLICATION" SECTION 7, "INTEGRATION OF SAP HANA WITH THE CLUSTER" CLUSTER SECTION 9, "TESTING THE CLUSTER"

Procedure

1. Basic Cluster Configuration
2. Configure Cluster Properties and Resources
3. Final steps

8.1 Basic Cluster Configuration

8.1.1 Set up Watchdog for SBD Fencing

All instances will use SUSE's Diskless SBD fencing mechanism. This method does not require additional AWS permissions because SBD does not issue AWS API calls. Instead SBD relies on (hardware/software) watchdog timers.

Most AWS bare metal instances feature a hardware watchdog, and on these instances no additional action is required to use the hardware watchdog, and non-bare metal instances will use a software watchdog.

Whenever SBD is used, a correctly working watchdog is crucial. Modern systems support a watchdog that needs to be "tickled" or "fed" by a software component. The software component (usually a daemon) regularly writes a service pulse to the watchdog. If the daemon stops feeding the watchdog, the hardware will enforce a system restart. This protects against failures of the SBD process itself, such as dying, or becoming stuck on an I/O error.

Determine the right watchdog module. Alternatively, you can find a list of installed drivers with your kernel version.

```
ls -l /lib/modules/$(uname -r)/kernel/drivers/watchdog
```

Check if any watchdog module is already loaded.

```
lsmod | egrep "(wd|dog|i6|iT)"
```

If you get a result, the system has already a loaded watchdog.

Check if any software is using the watchdog

```
lsof /dev/watchdog
```

If no watchdog is available (on virtualized EC2 instances), you can enable the *softdog*.

To enable softdog persistently across reboots, execute the following step in **all** EC2 instances that are going to be part of the cluster (including the majority maker node)

```
echo softdog > /etc/modules-load.d/watchdog.conf  
systemctl restart systemd-modules-load
```

This will also load the software watchdog kernel module during system boot.

Check if the watchdog module is loaded correctly.

```
lsmod | grep dog
```

Testing the watchdog can be done with a simple action. Take care to switch off your SAP HANA first because the watchdog will force an unclean reset / shutdown of your system.

In case of a hardware watchdog, a desired action is predefined after the timeout of the watchdog has reached. If your watchdog module is loaded and not controlled by any other application, run the command below.



Important

Triggering the watchdog without continuously updating the watchdog does reset/switches off the system. This is the intended mechanism. The following commands will force your system to be reset/switched off.

```
touch /dev/watchdog
```

In case of the softdog module is used the following can be done.

```
echo 1> /dev/watchdog
```

After your test was successful you can implement the watchdog on **all** cluster members. The example below applies to the softdog module.

```
for i in suse{02,03,04,05,06,-mm}; do
  ssh -T $i <<EOSSH
  hostname
  echo softdog > /etc/modules-load.d/watchdog.conf
  systemctl restart systemd-modules-load
  lsmod |grep -e dog
EOSSH
done
```

Once all cluster nodes have access to hardware/software watchdog devices at /dev/watchdog check the following attributes of the SBD configuration at /etc/sysconfig/sbd on **all** cluster nodes:

```
#SBD_DEVICE=""
SBD_PACEMAKER=yes
SBD_STARTMODE=always
SBD_DELAY_START=no
SBD_WATCHDOG_DEV=/dev/watchdog
SBD_WATCHDOG_TIMEOUT=5
SBD_TIMEOUT_ACTION=flush,reboot
SBD_OPTS=
```

Now enable diskless SBD on all cluster nodes:

```
systemctl enable sbd
```

8.1.2 Initial Cluster Setup

Since AWS VPC does not support multicast traffic corosync communication requires unicast UDP, On the first cluster node create an encryption key for the cluster communication:

```
corosync-keygen
```

The above command will generate the file */etc/corosync/authkey*. Copy this key over to **all** nodes while keeping the Unix file owner and permissions unchanged:

```
ls -l /etc/corosync/authkey
```

```
-r----- 1 root root 128 Feb  5 19:47 /etc/corosync/authkey
```

After distributing the encryption key, create an initial `/etc/corosync/corosync.conf` configuration using these parameters for cluster timings, transport protocol and encryption. Exchange the `bindnetaddr` and the `ring0_addr` (IPv4 addresses) of all cluster nodes in the nodelist according to your network topology.

Example `corosync.conf` file:

```
# Read the corosync.conf.5 manual page
totem {
  version: 2
  token: 30000
  consensus: 32000
  token_retransmits_before_loss_const: 6
  secauth: on
  crypto_hash: sha1
  crypto_cipher: aes256
  clear_node_high_bit: yes
  interface {
    ringnumber: 0
    bindnetaddr: <<local-node-ip-address>>
    mcastport: 5405
    ttl: 1
  }
  transport: udpu
}
logging {
  fileline: off
  to_logfile: yes
  to_syslog: yes
  logfile: /var/log/cluster/corosync.log
  debug: off
  timestamp: on
  logger_subsys {
    subsys: QUORUM
    debug: off
  }
}
nodelist {
  node {
    ring0_addr: <<ip-node01-AZ-a>>
    nodeid: 1
  }
  node {
    ring0_addr: <<ip-node02-AZ-a>>
    nodeid: 2
  }
}
```

```

}
node {
  ring0_addr: <<ip-node03-AZ-a>>
  nodeid: 3
}
node {
  ring0_addr: <<ip-node01-AZ-b>>
  nodeid: 4
}
node {
  ring0_addr: <<ip-node02-AZ-b>>
  nodeid: 5
}
node {
  ring0_addr: <<ip-node03-AZ-b>>
  nodeid: 6
}
node {
  ring0_addr: <<ip-majority-maker-node>>
  nodeid: 7
}
}
quorum {
# Enable and configure quorum subsystem (default: off)
# see also corosync.conf.5 and votequorum.5
  provider: corosync_votequorum
}

```

Distribute the configuration to **all** nodes at `/etc/corosync/corosync.conf`

8.1.3 Start the Cluster

Start the cluster on all nodes

```
systemctl start pacemaker
```



Note

All nodes should be started in parallel. Otherwise unseen nodes might get fenced.

Check the cluster status with `crm_mon`. We use the option `"-r"` to also see resources, which are configured but stopped.

```
crm_mon -r
```

The command will show the "empty" cluster and will print something like the following screen output. The most interesting information for now is that there are two nodes in the status "online" and the message "partition with quorum".

```
Stack: corosync
Current DC: suse05 (version 1.1.18+20180430.b12c320f5-3.15.1-b12c320f5) - partition with
quorum
Last updated: Fri Nov 29 14:23:19 2019
Last change: Fri Nov 29 12:31:06 2019 by hacluster via crmd on suse03

7 nodes configured
0 resource configured

Online: [ suse-mm suse01 suse02 suse03 suse04 suse05 suse06 ]

No resources
```

8.2 Configure Cluster Properties and Resources

This section describes how to configure cluster resources, STONITH, and constraints using the `crm` configure shell command. This is also detailed in section *Configuring and Managing Cluster Resources (Command Line)*, SUSE Linux Enterprise High Availability of the SUSE Linux Enterprise High Availability Administration Guide at https://www.suse.com/documentation/sle-ha-12/singlehtml/book_sleha/book_sleha.html#cha.ha.manual_config.

Use the command `crm` to add the objects to the Cluster Resource Management (CRM). Copy the following examples to a local file and then load the configuration to the Cluster Information Base (CIB). The benefit is here that you have a scripted setup and a backup of your configuration.

Perform all `crm` commands only on **one** node, for example on machine `suse01`

First create a text file with the configuration, which you load into our cluster in a second step. This step is as follow:

```
vi crm-file<XX>
crm configure load update crm-file<XX>
```

8.2.1 Cluster Bootstrap and More

The first example defines the cluster bootstrap options including the resource and operation defaults.

The stonith-timeout should be greater than 1.2 times the SBD msgwait timeout.

```
vi crm-bs.txt
```

Enter the following to crm-bs.txt

```
property $id="cib-bootstrap-options" \  
    no-quorum-policy="suicide" \  
    stonith-enabled="true" \  
    stonith-action="reboot" \  
    stonith-watchdog-timeout="10" \  
  
op_defaults $id="op-options" \  
    timeout="600"  
  
rsc_defaults rsc-options: \  
    resource-stickiness="1000" \  
    migration-threshold="5"
```

Now add the configuration to the cluster.

```
crm configure load update crm-bs.txt
```

8.2.2 STONITH

As previously explained in the requirements section of this document, STONITH is crucial for a supported cluster setup. Without a valid fencing mechanism your cluster is unsupported.

As standard STONITH mechanism diskless SBD fencing is implemented. The SBD STONITH method is very stable, reliable and has proven very good road capabilities.

You can use other fencing methods available. However, intensive testing the server fencing under all circumstances is crucial.

If diskless SBD has been configured and enabled the SBD daemon will be started automatically with the cluster. You can verify this with:

```
# systemctl status sbd  
● sbd.service - Shared-storage based fencing daemon  
   Loaded: loaded (/usr/lib/systemd/system/sbd.service; enabled; vendor preset: disabled)  
   Active: active (running) since Fri 2019-02-15 08:12:57 UTC; 1 months 22 days ago  
     Docs: man:sbd(8)  
  Process: 10366 ExecStart=/usr/sbin/sbd $SBD_OPTS -p /var/run/sbd.pid watch  
 (code=exited, status=0/SUCCESS)  
 Main PID: 10374 (sbd)  
    Tasks: 3 (limit: 4915)
```

```
CGroup: /system.slice/sbd.service
├─10374 sbd: inquisitor
├─10375 sbd: watcher: Pacemaker
└─10376 sbd: watcher: Cluster
```

8.2.3 Cluster in Maintenance Mode

We will load the configuration for the resources and the constraints step-by-step to the cluster. The best way to avoid unexpected cluster reactions is to first set the complete cluster to maintenance mode. Then, after all needed changes have been made, as last step, the cluster can be removed from maintenance mode.

```
crm configure property maintenance-mode=true
```

8.2.4 SAPHanaTopology

Next, define the group of resources needed, before the SAP HANA instances can be started. Prepare the changes in a text file, for example *crm-saphanatop.txt*, and load these with the *crm* command.

You need to change maybe the **SID** and **instance number** (bold) to your values.

EXAMPLE 21: CONFIGURE SAPHANATOPOLOGY

```
suse01:~ # vi crm-saphanatop.txt
```

Enter the following to *crm-saphanatop.txt*

```
primitive rsc_SAPHanaTop_<SID>_HDB<Inst> ocf:suse:SAPHanaTopology \
    op monitor interval="10" timeout="600" \
    op start interval="0" timeout="600" \
    op stop interval="0" timeout="300" \
    params SID="<SID>" InstanceNumber="<Inst>"

clone cln_SAPHanaTop_<SID>_HDB<Inst> rsc_SAPHanaTop_<SID>_HDB<Inst> \
    meta clone-node-max="1" interleave="true"
```

```
primitive rsc_SAPHanaTop_HA1_HDB00 ocf:suse:SAPHanaTopology \
    op monitor interval="10" timeout="600" \
    op start interval="0" timeout="600" \
    op stop interval="0" timeout="300" \
    params SID="HA1" InstanceNumber="00"

clone cln_SAPHanaTop_HA1_HDB00 rsc_SAPHanaTop_HA1_HDB00 \
```



```
meta clone-node-max="1" interleave="true"
```

For additional information about all parameters could be found with the command *man ocf_suse_SAPHanaTopology*.

Again, add the configuration to the cluster.

```
crm configure load update crm-saphanatop.txt
```

The most important parameters here are *SID* (HA1) and *InstanceNumber* (00), which are self explaining in an SAP context.

Beside these parameters, the timeout values or the operations (start, monitor, stop) are typical values to be adjusted to your environment.

8.2.5 SAPHanaController

Next, define the group of resources needed before the SAP HANA instances can be started. Edit the changes in a text file, for example *crm-saphanacon.txt* and load these with the command *crm*.

```
vi crm-saphanacon.txt
```

EXAMPLE 22: CONFIGURE SAPHANACONTROLLER

Enter the following to *crm-saphanacon.txt*

```
primitive rsc_SAPHanaCon_<SID>_HDB<Inst> ocf:suse:SAPHanaController \  
    op start interval="0" timeout="3600" \  
    op stop interval="0" timeout="3600" \  
    op promote interval="0" timeout="3600" \  
    op monitor interval="60" role="Master" timeout="700" \  
    op monitor interval="61" role="Slave" timeout="700" \  
    params SID="<SID>" InstanceNumber="<Inst>" \  
    PREFER_SITE_TAKEOVER="true" \  
    DUPLICATE_PRIMARY_TIMEOUT="7200" AUTOMATED_REGISTER="false"  
  
ms msl_SAPHanaCon_<SID>_HDB<Inst> rsc_SAPHanaCon_<SID>_HDB<Inst> \  
    meta clone-node-max="1" master-max="1" interleave="true"
```

The most important parameters here are *<SID>* (HA1) and *<Inst>* (00), which are in the SAP context quite self explaining. Beside these parameters, the timeout values or the operations (start, monitor, stop) are typical tunables.

```
primitive rsc_SAPHanaCon_HA1_HDB00 ocf:suse:SAPHanaController \  
    op start interval="0" timeout="3600" \  
    op stop interval="0" timeout="3600" \  
    op promote interval="0" timeout="3600" \  
    op monitor interval="60" role="Master" timeout="700" \  
    op monitor interval="61" role="Slave" timeout="700" \  
    params SID="HA1" InstanceNumber="00" \  
    PREFER_SITE_TAKEOVER="true" \  
    DUPLICATE_PRIMARY_TIMEOUT="7200" AUTOMATED_REGISTER="false"
```

```

op monitor interval="60" role="Master" timeout="700" \
op monitor interval="61" role="Slave" timeout="700" \
params SID="HA1" InstanceNumber="00" PREFER_SITE_TAKEOVER="true" \
DUPLICATE_PRIMARY_TIMEOUT="7200" AUTOMATED_REGISTER="false"

ms msl_SAPHanaCon_HA1_HDB00 rsc_SAPHanaCon_HA1_HDB00 \
meta clone-node-max="1" master-max="1" interleave="true"

```

Add the configuration to the cluster.

```
crm configure load update crm-saphanacon.txt
```

TABLE 2: TABLE DESCRIPTION OF IMPORTANT RESOURCE AGENT PARAMETER

| Name | Description |
|---------------------------|--|
| PREFER_SITE_TAKEOVER | Defines whether RA should prefer to takeover to the secondary instance instead of restarting the failed primary locally. |
| AUTOMATED_REGISTER | <p>Defines whether a former primary should be automatically registered to be secondary of the new primary. With this parameter, you can adapt the level of system replication automation.</p> <p>If set to false, the former primary must be manually registered. The cluster will not start this SAP HANA RDBMS until it is registered to avoid double primary up situations.</p> |
| DUPLICATE_PRIMARY_TIMEOUT | Time difference needed between two primary time stamps, if a dual-primary situation occurs. If the time difference is less than the time gap, the cluster holds one or both sites in a "WAITING" status. This is to give an administrator the chance to react on a failover. If the complete node of the former primary crashed, the former primary will be registered after the time difference is passed. If "only" the SAP HANA RDBMS has crashed, then the former primary will be registered immediately. After this registration to the new primary all data will be overwritten by the system replication. |

Additional information about all parameters could be found with the command `man ocf_suse_SAPHanaController`.

8.2.6 Overlay IP Address

The last resource to be added to the cluster is covering the Overlay IP address. Replace the bold string with your instance number, SAP HANA system id, the AWS VPC routing table(s), and the Overlay IP address.

EXAMPLE 23: CONFIGURE THE IP ADDRESS

```
vi crm-oip.txt
```

Enter the following to crm-oip.txt

```
primitive rsc_ip_<SID>_HDB<Inst> ocf:heartbeat:aws-vpc-move-ip \  
    op start interval=0 timeout=180 \  
    op stop interval=0 timeout=180 \  
    op monitor interval=60 timeout=60 \  
    params ip=<overlayIP> routing_table=<aws-route-table>[,<2nd-route-table>] \  
    interface=eth0 profile=<aws-cli-profile>
```

We load the file to the cluster.

```
crm configure load update crm-oip.txt
```

The Overlay IP address needs to be outside the CIDR range of the VPC.

8.2.7 Constraints

The constraints are organizing the correct placement of the virtual IP address for the client database access and the start order between the two resource agents *SAPHana* and *SAPHanaTopology*. The rules help to remove false positive messages from *crm_mon* command.

EXAMPLE 24: CONFIGURE NEEDED CONSTRAINTS

```
vi crm-cs.txt
```

Enter the following to crm-cs.txt

```
colocation col_saphana_ip_<SID>_HDB<Inst> 2000: rsc_ip_<SID>_HDB<Inst>:Started \  
    msl_SAPHanaCon_<SID>_HDB<Inst>:Master  
order ord_SAPHana_<SID>_HDB<Inst> Optional: cln_SAPHanaTop_<SID>_HDB<Inst> \  
    msl_SAPHanaCon_<SID>_HDB<Inst>  
location OIP_not_on_majority_maker rsc_ip_<SID>_HDB<Inst> -inf: <majority maker>  
location SAPHanaCon_not_on_majority_maker msl_SAPHanaCon_<SID>_HDB<Inst> -inf:  
    <majority maker>  
location SAPHanaTop_not_on_majority_maker cln_SAPHanaTop_<SID>_HDB<Inst> -inf:
```

```
<majority maker>
```

Replace "<SID>" by SAP SID, "<Inst>" by SAP HANA instance number, and "<majority maker>" by the majority maker node host name.

```
colocation col_saphana_ip_HA1_HDB00 2000: rsc_ip_HA1_HDB00:Started \  
  msl_SAPHanaCon_HA1_HDB00:Master  
order ord_SAPHana_HA1_HDB00 Optional: cln_SAPHanaTop_HA1_HDB00 \  
  msl_SAPHanaCon_HA1_HDB00  
location OIP_not_on_majority_maker rsc_ip_HA1_HDB00 -inf: suse-mm  
location SAPHanaCon_not_on_majority_maker msl_SAPHanaCon_HA1_HDB00 -inf: suse-mm  
location SAPHanaTop_not_on_majority_maker cln_SAPHanaTop_HA1_HDB00 -inf: suse-mm
```

We load the file to the cluster.

```
crm configure load update crm-cs.txt
```

8.3 Final Steps

8.3.1 End the Cluster Maintenance Mode

If maintenance mode has been enabled to configure the cluster then as last step it is required to remove the cluster from maintenance mode.

It may take several minutes for the cluster to stabilize as it may be required to start SAP HANA and other cluster services on the required nodes.

```
crm configure property maintenance-mode=false
```

8.3.2 Verify the Communication between the Hook and the Cluster

Now check if the HA/DR provider could set the appropriate cluster attribute `hana_<sid>_glob_srHook`. Replace the `<sid>` by the **lowercase** SAP system ID (like `ha1`).

EXAMPLE 25: QUERY THE SRHOOK CLUSTER ATTRIBUTE

```
crm_attribute -G -n hana_<sid>_glob_srHook
```

You should get an output like:

```
scope=crm_config name=hana_<sid>_glob_srHook value=SOK
```

In this case the HA/DR provider set the attribute to SOK to inform the cluster about SAP HANA System Replication status.

8.3.3 Using Special Virtual Host Names or FQHN During Installation of SAP HANA

If you have used special virtual host names or the fully qualified host name (FQHN) instead of the short node name, the resource agents needs to map these names. To be able to match the short node name with the used SAP 'virtual host name', the saphostagent needs to report the list of installed instances correctly:

EXAMPLE 26: IN OUR SETUP THE VIRTUAL HOST NAME MATCHES THE NODE NAME

```
suse01:haladm> /usr/sap/hostctrl/exe/saphostctrl -function ListInstances
Inst Info : HA1 - 00 - suse01 - 749, patch 418, changelist 1816226
```

9 Testing the Cluster

Testing is one of the most important project tasks for implementing clusters. Proper testing is crucial. Make sure that all test cases derived from project or customer expectations are defined and passed completely. **Without testing the project is likely to fail in production use.**



FIGURE 11: SECTION 3, "PLANNING THE INSTALLATION" SECTION 4, "USING AWS WITH SUSE LINUX ENTERPRISE HIGH AVAILABILITY EXTENSION CLUSTERS" SECTION 5, "INSTALLING THE SAP HANA DATABASES ON BOTH SITES" SECTION 6, "SET UP THE SAP HANA SYSTEM REPLICATION" SECTION 7, "INTEGRATION OF SAP HANA WITH THE CLUSTER" SECTION 8, "CONFIGURATION OF THE CLUSTER AND SAP HANA RESOURCES" TESTING

The test prerequisite, if not described differently, is always that all cluster nodes are booted, are already normal members of the cluster and the SAP HANA RDBMS is running. The system replication is in sync represented by 'SOK'. The cluster is idle, no actions are pending, no migration constraints left over, no failcounts left over.

In this version of the setup guide we provide a plain list of test cases. We plan to describe the test cases more detailed in the future. Either we will provide these details in an update of this guide or we will extract the test cases to a separate test plan document.

9.1 Generic Cluster Tests

This kind of cluster tests covers the cluster reaction during operations. This includes starting and stopping the complete cluster or simulating SBD failures and much more.

- Parallel start of all cluster nodes (systemctl start pacemaker should be done in a short time frame).
- Stop of the complete cluster.
- Isolate ONE of the two SAP HANA sites.
- Power-off the majority maker.
- Simulate a maintenance procedure with cluster continuously running.
- Simulate a maintenance procedure with cluster restart.
- Kill the corosync process of one of the cluster nodes.

9.2 Tests on the Primary Site

This kind of tests are checking the reaction on several failures of the primary site.

9.2.1 Tests Regarding Cluster Nodes of the Primary Site

The tests listed here check the SAP HANA and cluster reaction if one or more nodes of the primary site are failing or re-joining the cluster.

- Power-off master name server of the primary.
- Power-off any worker node but not the master name server of the primary.
- Re-join of a previously powered-off cluster node.

9.2.2 Tests Regarding the Complete Primary Site

This test category is simulating a complete site failure.

- Power-off all nodes of the primary site in parallel.

9.2.3 Tests regarding the SAP HANA Instances of the Primary Site

The tests listed here are checks about the SAP HANA and cluster reactions triggered by application failures such as a killed SAP HANA instance.

- Kill the SAP HANA instance of the master name server of the primary.
- Kill the SAP HANA instance of any worker node but not the master name server of the primary.
- Kill sapstarttrv of any SAP HANA instance of the primary.

9.3 Tests on the Secondary Site

This kind of tests are checking the reaction on several failures of the secondary site.

9.3.1 Tests regarding Cluster Nodes of the Secondary Site

The tests listed here check the SAP HANA and cluster reaction if one or more nodes of the secondary site are failing or re-joining the cluster.

- Power-off master name server of the secondary.
- Power-off any worker node but not the master name server of the secondary.
- Re-join of a previously powered-off cluster node.

9.3.2 Tests Regarding the Complete Secondary Site

This test category is simulating a complete site failure.

- Power-off all nodes of the secondary site in parallel.

9.3.3 Tests Regarding the SAP HANA Instances of the Secondary Site

The tests listed here are checks about the SAP HANA and cluster reactions triggered by application failures such as a killed SAP HANA instance.

- Kill the SAP HANA instance of the master name server of the secondary.
- Kill the SAP HANA instance of any worker node but not the master name server of the secondary.
- Kill `sapstartrv` of any SAP HANA instance of the secondary.

10 Administration

10.1 Dos and Don'ts

In your project, you should **do**:

- Define (and test) **STONITH** **before** adding other resources to the cluster.
- Do **intensive** testing.
- **Tune** the timeouts of operations of `SAPHanaController` and `SAPHanaTopology`.
- Start with `PREFER_SITE_TAKEOVER = true`, `AUTOMATED_REGISTER = false` and `DUPLICATE_PRIMARY_TIMEOUT = "7200"`.
- Always make sure that the cluster configuration does not contain any left-over client-prefer location constraints or failcounts.
- Before testing or beginning maintenance procedures check, if the cluster is in idle state.

In your project, **avoid**:

- Rapidly changing/changing back cluster configuration, such as: Setting nodes to standby and online again or stopping/starting the master/slave resource.
- Creating a cluster without proper time synchronization or unstable name resolutions for hosts, users, and groups.

- Adding location rules for the clone, master/slave or IP resource. Only location rules mentioned in this setup guide are allowed.
- As "migrating" or "moving" resources in crm-shell, HAWK or other tools would add client-prefer location rules this activities are completely **forbidden!**.

10.2 Monitoring and Tools

You can use the High Availability Web Konsole (HAWK), SAP HANA Studio and different command line tools for cluster status requests.

10.2.1 HAWK – Cluster Status and More

You can use an Internet browser to check the cluster status. Use the following URL: <https://<node>:7630>

The login credentials are provided during installation dialog of ha-cluster-init. Keep in mind to change the default password of the Linux user *hacluster*.




FIGURE 12: CLUSTER STATUS IN HAWK

If you set up the cluster using *ha-cluster-init* and you have installed all packages as described above, your system will provide a very useful Web interface. You can use this graphical Web interface to get an overview of the complete cluster status, doing administrative tasks or even configure resources and cluster bootstrap parameters.

Read our product manuals for a complete documentation of this powerful user interface.

10.2.2 SAP HANA Studio

Database-specific administration and checks can be done with SAP HANA studio.

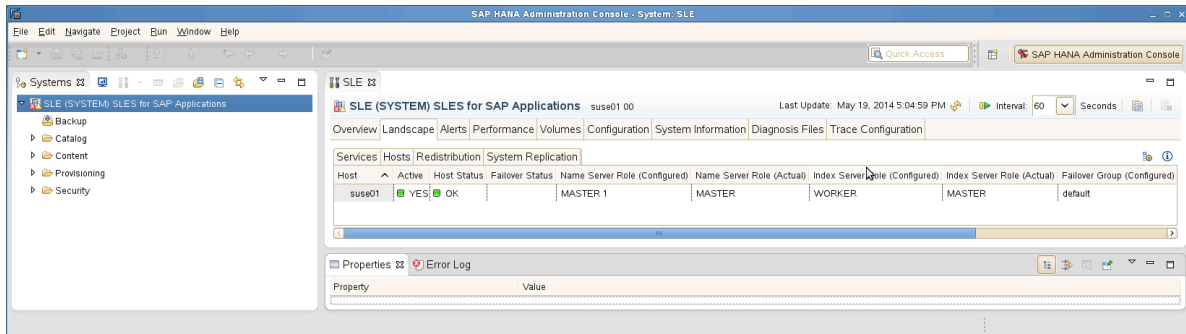


FIGURE 13: SAP HANA STUDIO – LANDSCAPE OF A SCALE-OUT SYSTEM

Be extremely careful with changing any parameters or topology of the system replication as it might get an interference with the cluster resource management.

A positive example would be to register a former primary as new secondary and you have set `AUTOMATED_REGISTER = false`.

A negative example would be to deregister a secondary, disable the system replication on the primary and such action.

For all actions which would change the system replication we recommend to first check for the maintenance procedure.

10.2.3 Cluster Command Line Tools

crm_mon

A simple overview can be obtained by calling `crm_mon`. Using option `-r` shows also stopped but already configured resources. Option `-l` tells `crm_mon` to output the status once instead of periodically.

```
Stack: corosync
Current DC: suse05 (version 1.1.16-4.8-77ea74d) - partition with quorum
Last updated: Mon Jun 11 16:55:04 2018
```

```
Last change: Mon Jun 11 16:53:58 2018 by root via crm_attribute on suse02
```

```
7 nodes configured
```

```
16 resources configured
```

```
Online: [ suse-mm suse01 suse02 suse03 suse04 suse05 suse06 ]
```

```
Full list of resources:
```

```
stonith-sbd      (stonith:external/sbd): Started suse-mm  
rsc_ip_HA1_HDB00 (ocf::heartbeat:IPaddr2):      Started suse02  
Master/Slave Set: msl_SAPHanaCon_HA1_HDB00 [rsc_SAPHanaCon_HA1_HDB00]  
  Masters: [ suse02 ]  
  Slaves: [ suse01 suse03 suse04 suse05 suse06 ]  
  Stopped: [ suse-mm ]  
Clone Set: cln_SAPHanaTop_HA1_HDB00 [rsc_SAPHanaTop_HA1_HDB00]  
  Started: [ suse01 suse02 suse03 suse04 suse05 suse06 ]  
  Stopped: [ suse-mm ]
```

See the manual page `crm_mon(8)` for details.

SAPHanaSR-showAttr

To show some of the SAPHanaController and SAPHanaTopology resource agent internal values, you can call the program *SAPHanaSR-showAttr*. The internal values, storage place and their parameter names may change in the next versions. The command *SAPHanaSR-showAttr* will always fetch the values from the correct storage place.



Important

Do **not** use cluster commands like *crm_attribute* to fetch the values directly from the cluster. In such cases your methods will be broken, when we need to move an attribute to a different storage place or even out of the cluster.

For first *SAPHanaSR-showAttr* is a test program only and should not be used for automated system monitoring.

EXAMPLE 27: CHECK SAPHANASR-SHOWATTR AS USER ROOT

```
suse-mm:~ # SAPHanaSR-showAttr --sid=<SID>
```

The tool display all interesting cluster attributes in three areas.

- The **global** section includes the information about the cib time stamp and the attributes covering the status of the system replication
- The **site** section includes the attributes per site and shows which site is the primary and the return code of the `landscapeHostConfiguration.py` script. In addition the active master name server is shown per site.
- The **hosts** section includes the node status, the roles of the host inside the SAP HANA database, the calculated score to get the primary master name server and the site name the host belongs to.

```

Global cib-time                prim sec  srHook sync_state
-----
global Tue Jun 12 15:02:58 2018 WDF1 ROT1 SOK    SOK

Site lpt          lss mns    srr
-----
WDF1 1528808568 4   suse02 P
ROT1 30          4   suse01 S

Hosts  clone_state node_state roles                                score site
-----
suse-mm online
suse01 DEMOTED    online   master1:master:worker:master 100 ROT1
suse02 PROMOTED    online   master1:master:worker:master 150 WDF1
suse03 DEMOTED    online   master3:slave:worker:slave   80  ROT1
suse04 DEMOTED    online   master2:slave:worker:slave   110 WDF1
suse05 DEMOTED    online   master2:slave:worker:slave   80  ROT1
suse06 DEMOTED    online   master3:slave:worker:slave   110 WDF1

```

The majority maker `suse-mm` does not run an SAP HANA instance. Therefore it neither has a role attribute nor a score or site value.

10.2.4 SAP HANA LandscapeHostConfiguration

To check the status of an SAP HANA database and to figure out if the cluster should react, you can use the script `landscapeHostConfiguration.py`.

EXAMPLE 28: CHECK THE LANDSCAPE STATUS AS USER <SID>ADM

```
HDBSettings.sh landscapeHostConfiguration.py
```

The landscape host configuration is shown with a line per SAP HANA host.

```

| Host | Host | ... NameServer | NameServer | IndexServer | IndexServer |
|      | Active | ... Config Role | Actual Role | Config Role | Actual Role |
| ----- | ----- | ... ----- | ----- | ----- | ----- |
| suse01 | yes | ... master 1 | master | worker | master |
| suse03 | yes | ... master 2 | slave | worker | slave |
| suse05 | yes | ... master 3 | slave | standby | standby |

overall host status: ok

```

Following the SAP HA guideline, the SAPHana resource agent interprets the return codes in the following way:

TABLE 3: TABLE INTERPRETATION OF RETURN CODES

| Return Code | Description |
|-------------|---|
| 4 | SAP HANA database is up and OK. The cluster does interpret this as correctly running database. |
| 3 | SAP HANA database is up and in status INFO. The cluster does interpret this as a correctly running database. |
| 2 | SAP HANA database is up and in status warning. The cluster does interpret this as a correctly running database. |
| 1 | SAP HANA database is down. If the database should be up and is not own by intention, this could trigger a takeover. |
| 0 | Internal Script Error – to be ignored. |

11 Useful Links, Manuals, and SAP Notes

11.1 SUSE Best Practices and More

Best Practices for SAP on SUSE Linux Enterprise

<https://www.suse.com/products/sles-for-sap/resource-library/sap-best-practices.html> 

Fail-Safe Operation of SAP HANA*: SUSE Extends Its High-Availability Solution

<http://scn.sap.com/community/hana-in-memory/blog/2014/04/04/fail-safe-operation-of-sap-hana-suse-extends-its-high-availability-solution>

HOW TO SET UP SAPHanaSR IN THE COST OPTIMIZED SAP HANA SR SCENARIO

<http://scn.sap.com/docs/DOC-65899>

11.2 SUSE Product Documentation

The SUSE product manuals and documentation can be downloaded at <https://www.suse.com/documentation>.

Current online documentation of SLES for SAP Applications

<https://www.suse.com/documentation/sles-for-sap-12/>

Current online documentation of SUSE Linux Enterprise High Availability

<https://www.suse.com/documentation/sle-ha-12/index.html>

Tuning guide for SUSE Linux Enterprise Server

https://www.suse.com/documentation/sles-12/book_sle_tuning/data/book_sle_tuning.html

Storage admin guide for SUSE Linux Enterprise Server

https://www.suse.com/documentation/sles-12/stor_admin/data/stor_admin.html

Release notes

<https://www.suse.com/releasenotes/>

TID aws-cli throwing Error on fresh SLES 15 Installation in AWS Cloud

<https://www.suse.com/support/kb/doc?id=7023686>

TID multipath system unable to boot after installing dracut-037-98.2.x86_64

<https://www.suse.com/support/kb/doc?id=7020912>

TID Systemd-udev-settle timing out

<https://www.suse.com/support/kb/doc?id=7022681>

TID How to load the correct watchdog kernel module

<http://www.suse.com/support/kb/doc.php?id=7016880>

TID rpcbind won't start after upgrade from SLES 11 to SLES 12

<https://www.suse.com/support/kb/doc?id=7017144>

TID Memory, I/O and DefaultTasksMax related considerations for SLES for SAP servers with huge memory

<https://www.suse.com/support/kb/doc?id=7021211> ↗

TID XFS metadata corruption and invalid checksum on SAP Hana servers

<https://www.suse.com/support/kb/doc?id=7022921> ↗

SUSE Linux Enterprise Server technical information

<https://www.suse.com/products/server/technical-information/> ↗

XFS file system

<https://www.suse.com/communities/conversations/xfs-the-file-system-of-choice/> ↗

11.3 SAP Product Documentation

SAP HANA Installation and Update Guide

http://help.sap.com/hana/SAP_HANA_Server_Installation_Guide_en.pdf ↗

SAP HANA Administration Guide

http://help.sap.com/hana/SAP_HANA_Administration_Guide_en.pdf ↗

11.4 SAP Notes

As SAP Notes are changing over time, this list is only a starting point

- 611361 Hostnames of SAP servers
- 1275776 Preparing SLES for Sap Environments
- 1514967 SAP HANA: Central Note
- 1523337 SAP In-Memory Database 1.0: Central Note
- 1501701 Single Computing Unit Performance and Sizing
- 1846872 "No space left on device" error reported from HANA
- 1876398 Network configuration for System Replication in HANA SP6
- 1888072 SAP HANA DB: Indexserver crash in strcmp sse42
- 1890444 Slow HANA system due to CPU power save mode

- [1944799 SAP HANA Guidelines for SLES Operating System Installation](#)
- [1984787 SUSE LINUX Enterprise Server 12: Installation notes and](#)
- [1999993 How-To: Interpreting SAP HANA Mini Check Results](#)
- [2000000 FAQ: SAP HANA Performance Optimization](#)
- [2100040 FAQ: SAP HANA CPU](#)
- [2205917 SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12.](#)
- [2382421 Optimizing the Network Configuration on HANA- and OS-Level](#)
- [2470289 FAQ: SAP HANA Non-Uniform Memory Access \(NUMA\)](#)
- [2578899 SUSE Linux Enterprise Server 15: Installation Note](#)
- [2684254 SAP HANA DB: Recommended OS settings for SLES 15 / SLES for SAP Applications 15](#)
- [2647673 HANA Installation Failure](#)

12 Legal notice

Copyright © 2006–2024 SUSE LLC and contributors. All rights reserved.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or (at your option) version 1.3; with the Invariant Section being this copyright notice and license. A copy of the license version 1.2 is included in the section entitled "GNU Free Documentation License".

SUSE, the SUSE logo and YaST are registered trademarks of SUSE LLC in the United States and other countries. For SUSE trademarks, see <https://www.suse.com/company/legal/>.

Linux is a registered trademark of Linus Torvalds. All other names or trademarks mentioned in this document may be trademarks or registered trademarks of their respective owners.

Documents published as part of the SUSE Best Practices series have been contributed voluntarily by SUSE employees and third parties. They are meant to serve as examples of how particular actions can be performed. They have been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. SUSE cannot verify that actions described in these documents do what is claimed or whether actions described have unintended consequences. SUSE LLC, its affiliates, the authors, and the translators may not be held liable for possible errors or the consequences thereof.

Below we draw your attention to the license under which the articles are published.

13 GNU Free Documentation License

Copyright © 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition. The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.

- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all

Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>. Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

Copyright (c) YEAR YOUR NAME.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2

```
or any later version published by the Free Software Foundation;  
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.  
A copy of the license is included in the section entitled "GNU  
Free Documentation License".
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “ with... Texts.” line with this:

```
with the Invariant Sections being LIST THEIR TITLES, with the  
Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.