SAP HANA System Replication Scale-Up - Performance Optimized Scenario

with SAPHanaSR-angi

SUSE Linux Enterprise Server for SAP applications 16

Fabian Herschel, Distinguished Architect SAP (SUSE)
Lars Pinne, System Engineer (SUSE)

SAP HANA System Replication Scale-Up - Performance Optimized Scenario

with SAPHanaSR-angi

Date: 2026-11-12

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 system replication in the performance optimized scenario. The document focuses on the steps to integrate an already installed and working SAP HANA with system replication. It is based on SUSE Linux Enterprise Server for SAP applications 16.0.

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 Supported setups and prerequisites 8

3 Scope of this document 12

4 Planning the installation 13

5 Setting up the operating system 16

6 Installing the SAP HANA Databases on both cluster nodes 18

7 Setting up SAP HANA System Replication 20

8 Setting up SAP HANA HA/DR providers 24

9 Configuring the cluster 31

10 Testing the cluster 43

11 Administration 57

12 Examples 69

13 References 75

14 Legal notice 85

15 GNU Free Documentation License 86

16 GNU Free Documentation License 94

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 system replication in the scale-up performance optimized scenario**.

SUSE supports customers by offering SUSE Linux Enterprise Server for SAP applications, the recommended and supported operating system for SAP HANA. In close collaboration with SAP, cloud service and hardware partners, SUSE provides solutions for customers to ensure high availability for SAP HANA data bases.

1.1.1 Abstract

This guide describes planning, setup, and basic testing of SUSE Linux Enterprise Server for SAP applications based on the high availability solution scenario "SAP HANA scale-up performance optimized".

From an application perspective, the following variations are covered:

- Plain system replication
- System replication with secondary site read-enabled
- Multi-tier (chained) system replication
- Multi-target system replication
- Multi-tenant database containers for all of the above

From an infrastructure perspective, the following scenario variations are covered:

- 2-node cluster with disk-based SBD
- 3-node cluster with diskless SBD
- On-premises deployment on physical and virtual machines
- Public cloud deployment (usually needs additional documentation focusing on the cloud specific implementation details)

Deployment automation simplifies roll-out. There are several options available, particularly on public cloud platforms. Ask your public cloud provider or your SUSE contact for more information.

See *Section 2, "Supported setups and prerequisites"* for details.

Note

In this guide the software package `SAPHanaSR-angi` is used. This package replaces the two packages `SAPHanaSR` and `SAPHanaSR-ScaleOut`. Therefore, deploy new deployment with SAPHanaSR-angi only. For upgrading existing clusters to SAPHanaSR-angi, read the blog article at https://www.suse.com/c/how-to-upgrade-to-saphanasr-angi/ .

1.1.2 The concept of the performance optimized scenario

This document describes the **performance optimized** scenario [$A' \Rightarrow B$]. An overview of available scenarios is decribed in our white paper.

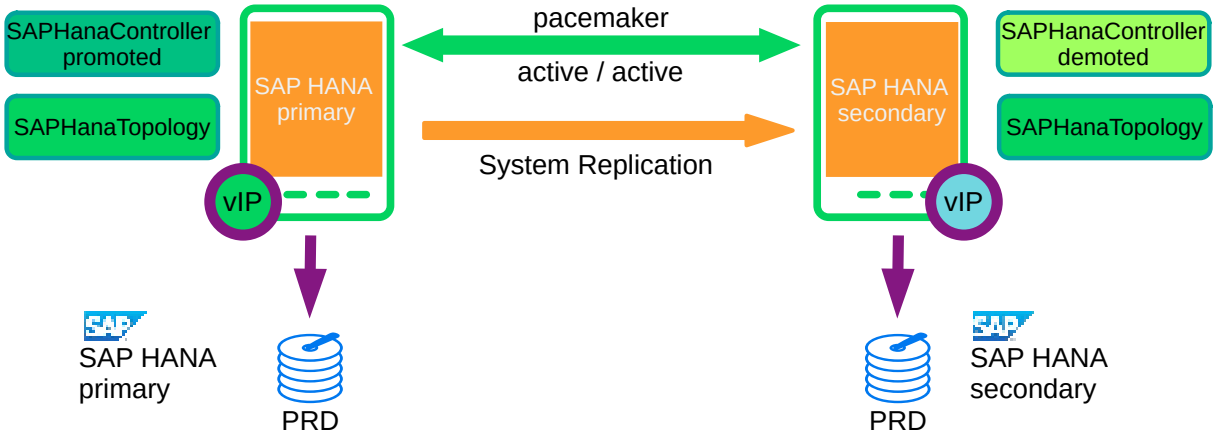

FIGURE 1: SAP HANA SYSTEM REPLICATION SCALE-UP IN THE CLUSTER - PERFORMANCE OPTIMIZED

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

One big advantage of the performance optimized scenario of SAP HANA is the option to allow read access on the secondary database site. To support this **read enabled** scenario, a second virtual IP address is added to the cluster and bound to the secondary role of the system replication [$A' \Rightarrow B''$].

In case of failure of the primary SAP HANA on node 1 (node or database instance), the cluster first tries to start the takeover process. This allows to use the already loaded data at the secondary site. Typically, the takeover is much faster than the local restart.

To achieve an automation of this resource handling process, you must use the SAP HANA resource agents included in SAPHanaSR-angi. System replication of the productive database is automated with SAPHanaController and SAPHanaTopology.

The cluster allows a takeover to the secondary site only if the SAP HANA system replication was in sync until the service of the primary site was lost. This ensures that the last commits processed on the primary site are already available at the secondary site.

SAP did improve the interfaces between SAP HANA and external software, such as cluster frameworks. These improvements also include the implementation of SAP HANA call-outs in case of special events, such as status changes for services or system replication channels. These call-outs are also called HA/DR providers. These interfaces can be used by implementing SAP HANA hooks written in python. SUSE has enhanced the SAPHanaSR-angi package to include such SAP HANA hooks to optimize the cluster interface. Using the SAP HANA hooks described in this document allows to inform the cluster immediately if the SAP HANA system replication is broken. In addition to the SAP HANA hook status, the cluster continues to poll the system replication status on a regular basis.

Adjust the level of automation by setting the parameter `AUTOMATED_REGISTER`. If automated registration is activated, the cluster will automatically register a former failed primary to become the new secondary. Refer to the manual pages SAPHanaSR(7) and ocf_suse_SAPHanaController(7) for details on all supported parameters and features.

Important

The solution is not designed to manually 'migrate' the primary or secondary instance using the HAWK Web console or any cluster client commands. In the *Administration* section of this document, we describe how to 'migrate' the primary to the secondary site using SAP and cluster commands.

1.2 Ecosystem of the document

1.2.1 Additional documentation and resources

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

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

Find numerous whitepapers, best practices, setup guides, and other resources on the SUSE Linux Enterprise Server for SAP applications best practices Web page: https://documentation.suse.com/sbp/sap-16/. In particular, there is an overview of all SUSE high availability solutions for SAP HANA and SAP S/4HANA workloads. Find the overview of high availability solutions supported by SUSE Linux Enterprise Server for SAP applications here:

https://documentation.suse.com/sles-sap/sap-ha-support/html/sap-ha-support/article-sap-ha-support.html

SUSE also publishes blog articles about SAP and high availability. Join us by using the hashtag #TowardsZeroDowntime. Use the following link: https://www.suse.com/c/tag/TowardsZeroDowntime/.

Finally, there are manual pages shipped with the product.

1.2.2 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 Supported setups and prerequisites

For the `SAPHanaSR-angi` package configuration as decribed in this document, we limit the support to scale-up (single-box to single-box) system replication with the following configurations and parameters:

- Two-node clusters are standard. Three node clusters are fine if you install the resource agents also on that third node. But define in the cluster that SAP HANA resources must never run on that third node. In this case the third node is an additional majority maker in case of cluster separation.
- The cluster must include a valid fencing method.
  - Any fencing mechanism supported for production use by SUSE Linux Enterprise High Availability 16 (like SBD, IPMI) is supported with SAPHanaSR-angi.
  - This guide is focusing on the SBD fencing method as this is hardware independent.
  - If you use disk-based SBD as the fencing mechanism, you need one or more shared drives. For productive environments, we recommend SBD devices. For details on disk-based SBD, read the product documentation for SUSE Linux Enterprise High Availability and the manual pages sbd(8) and stonith_sbd(7).
  - For diskless SBD, you need at least three cluster nodes. The diskless SBD mechanism has the benefit that you do not need a shared drive for fencing. Since diskless SBD is based on self-fencing, reliable detection of lost quorum is absolutely crucial.
  - Priority fencing is an optional improvement for two nodes, but does not work for three nodes.
- Both nodes are in the same network segment (layer 2). Similar methods provided by cloud environments such as overlay IP addresses and load balancer functionality are also fine. Follow the cloud specific guides to set up your SUSE Linux Enterprise Server for SAP applications cluster.
- Technical users and groups, such as *<sid>adm* are defined locally in the Linux system. If that is not possible, additional measures are needed to ensure reliable resolution of users, groups and permissions at any time. This might include caching.
- Name resolution of the cluster nodes and the virtual IP address must be done locally on all cluster nodes. If that is not possible, additional measures are needed to ensure reliable resolution of host names at any time.
- Time synchronization between the cluster nodes, such as NTP, is required.
- Both SAP HANA instances of the system replication pair (primary and secondary) have the same SAP Identifier (SID) and instance number.

- If the cluster nodes are installed in different data centers or data center areas, the environment must match the requirements of the SUSE Linux Enterprise High Availability cluster product. Of particular concern are the network latency and recommended maximum distance between the nodes. Review the product documentation for SUSE Linux Enterprise High Availability about those recommendations.
- Automated registration of a failed primary after takeover prerequisites need to be defined.
  - As a good starting configuration for projects, we recommend to switch off the automated registration of a failed primary. The setup `AUTOMATED_REGISTER="false"` is set as default. In this case, you need to register a failed primary after a takeover manually. For re-registration, use precisely the site names that are already known by the cluster. Use SAP tools like SAP HANA cockpit or *hdbnsutil*.
  - For optimal automation, we recommend to set `AUTOMATED_REGISTER="true"`.
- Automated start of SAP HANA instances during system boot must be switched off.
- Multi-tenancy (MDC) databases are supported.
  - Multi-tenancy databases can be used in combination with any other setup (performance-optimized, cost-optimized, multi-tier, multi-target and read-enabled).
  - In MDC configurations, the SAP HANA RDBMS is treated as a single system including all database containers. Therefore, cluster takeover decisions are based on the complete RDBMS status independent of the status of individual database containers.
  - Tests on Multi-tenancy databases can force a different test procedure if you are using strong separation of the tenants. As an example, killing the complete SAP HANA instance using *HDB kill* does not work, because the tenants are running with different Linux user UIDs. *<sid>adm* is not allowed to terminate the processes of the other tenant users.
- Only one system replication between the two SAP HANA database in the Linux cluster. Maximum one system replication to an SAP HANA database outside the Linux cluster.
  - When an SAP HANA system replication site is known to the Linux cluster, that exact site name needs to be used whenever the site is registered manually.
  - If a third SAP HANA site is connected by system replication, that SAP HANA is not controlled by another Linux cluster. If that third site should work as part of a fall-back HA cluster in DR case, that HA cluster needs to be in standby.

    - The replication mode is either sync or syncmem for the controlled replication. Replication mode async is not supported. The operation modes delta_datashipping, logreplay and logreplay_readaccess are supported.
    - See also the dedicated section on requirements for susHanaSR.py.
- The current resource agent supports SAP HANA in system replication beginning with SAP HANA version 2.0 SPS05 revision 59.04. Even in SAP HANA multi-target environments, the current resource agent manages only two sites. Thus only two SAP HANA sites are part of the Linux cluster.
- Besides SAP HANA you need SAP hostagent installed and started on your system.
- On SLES 16 only systemd style is supported. The service SAP<SID>_<INO> can stay enabled. The systemd enabled saphostagent and instance´s sapstartsrv is supported. Refer to the OS documentation for the systemd version. SAP HANA comes with native systemd integration as default starting with version 2.0 SPS07. Refer to SAP documentation for information on other SAP HANA versions.
- The RA’s monitoring operations need to be active.
- Using HA/DR provider hook for srConnectionChanged() by enabling susHanaSR.py is mandatory.
- RA and HA/DR provider hook script´s runtime almost completely depends on call-outs to controlled resources, OS and Linux cluster. The infrastructure needs to allow these call-outs to return in time.
- Colocation constraints between the SAPHanaController RA and other resources are allowed only if they do not affect the RA’s scoring. The location scoring finally depends on system replication status and must not be overruled by additional constraints. Thus it is not allowed to define rules forcing an SAPHanaController promoted clone to follow another resource.
- Reliable access to the /hana/shared/ file system is crucial for SAP HANA and the Linux cluster.
- SAP HANA feature Secondary Time Travel is not supported.
- The SAP HANA Fast Restart feature on RAM-tmpfs and SAP HANA on persistent memory can be used, as long as they are supported by the OS and transparent to Linux HA.
- No manual actions must be performed on the SAP HANA database while it is controlled by the Linux cluster. All administrative actions need to be aligned with the cluster.
- No firewall rules must block any needed port.
- No SELinux rules must block any needed action.
- Sizing of both SAP HANA sites needs to be done according to SAP rules. The scale-up scenarios require both sites to be prepared to run the primary SAP HANA database.

For the HA/DR provider hook scripts *susHanaSR.py* and *susTkOver.py*, the following requirements apply:

- SAP HANA 2.0 SPS05 revision 059.04 and later provides Python3 and the HA/DR provider hook method **srConnectionChanegd()** with multi-target-aware parameters. Python 3 and multi-target-aware parameters are needed for the `SAPHanaSR-angi` package.
- SAP HANA 2.0 SPS05 and later provides the HA/DR provider hook method **preTakeover()**.
- The user <*sid*>*adm* needs execution permission as user root for the command *crm_attribute*.
- The hook provider needs to be added to the SAP HANA global configuration, in memory and on disk (in persistence).

For the HA/DR provider hook script *susChkSrv.py*, the following requirements apply:

- SAP HANA 2.0 SPS05 or later provides the HA/DR provider hook method srServiceStateChanged() with needed parameters.
- No other HA/DR provider hook script should be configured for the **srServiceStateChanged()** method. Hook scripts for other methods, provided in SAPHanaSR-angi can be used in parallel to *susChkSrv.py*, if not documented otherwise.
- The user <sid>adm needs execution permission as user root for the command SAPHanaSR-hookHelper.
- The hook provider needs to be added to the SAP HANA global configuration, in memory and on disk (in persistence).
- The hook script runs in the SAP HANA name server. It runs on the node where the event **srServiceStateChanged()** occurs.
- If susChkSrv.py parameter action_on_lost=stop is set and the RA SAPHanaController parameter AUTOMATED_REGISTER=true is set, it depends on HANA to release all OS resources prior to the registering attempt.

See also manual pages SAPHanaSR(7), susHanaSR.py(7), susTkOver.py(7) and susChkSrv.py(7) for more details and requirements.

Important

Without a valid fencing method, the complete cluster is unsupported and will not work properly.

If you need to implement a different scenario, we strongly recommend to define a Proof of Concept (PoC) with SUSE. This PoC will focus on testing the existing solution in your scenario. Most of the above mentioned limitations are set because careful testing is needed.

3 Scope of this document

This document describes how to set up the cluster to control the database instances and system replication for the scenario SAP HANA scale-up performance optimized.

The document focuses on the steps to integrate an already installed and working SAP HANA with System Replication. This guide can be used with SUSE Linux Enterprise Server for SAP applications 16.0 or newer.

The described example setup builds an SAP HANA HA cluster in two data centers in Walldorf (WDF) and in Rot (ROT), installed on two SLES for SAP 16.0 systems.

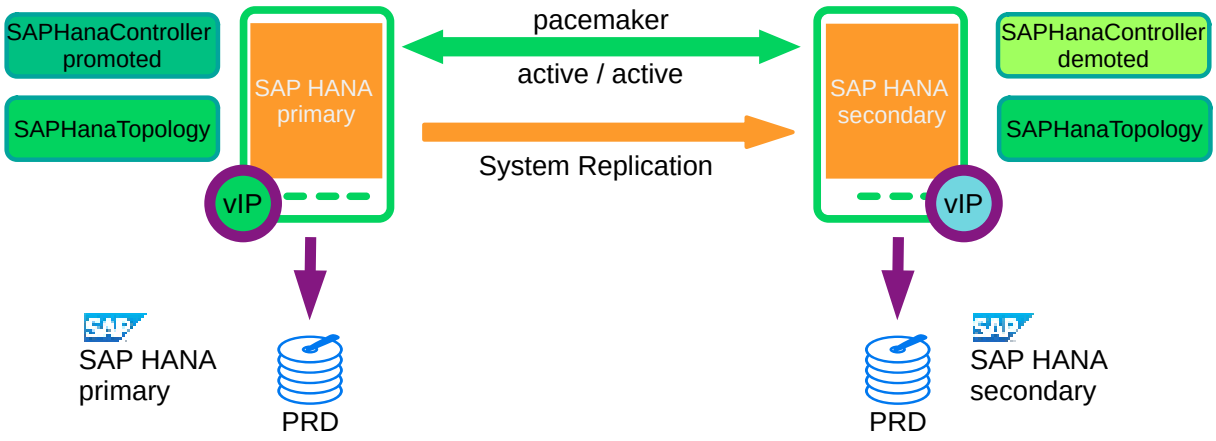

FIGURE 2: SAP HANA SYSTEM REPLICATION SCALE-UP IN THE CLUSTER - PERFORMANCE OPTIMIZED

This guide focuses on the manual setup of the cluster to explain the details and to give you the possibility to create your own automation.

The seven main setup steps are:

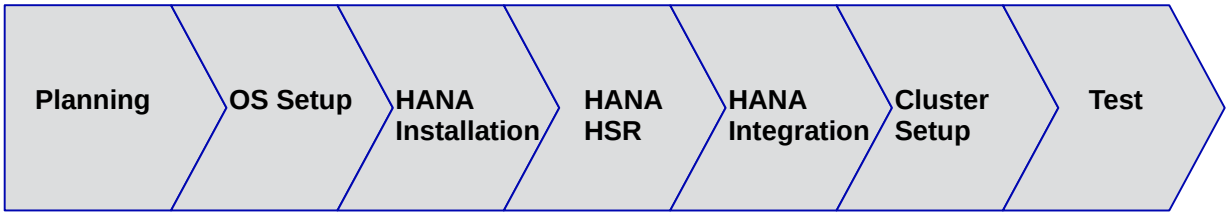

- Planning (see *Section 4, “Planning the installation”*)
- OS installation (see *Section 5, “Setting up the operating system”*)
- Database installation (see *Section 6, “Installing the SAP HANA Databases on both cluster nodes”*)
- SAP HANA system replication setup (see *Section 7, “Setting up SAP HANA System Replication”*
- SAP HANA HA/DR provider hooks (see *Section 8, “Setting up SAP HANA HA/DR providers”*)
- Cluster configuration (see *Section 9, “Configuring the cluster”*)
- Testing (see *Section 10, “Testing the cluster”*)

4 Planning the installation

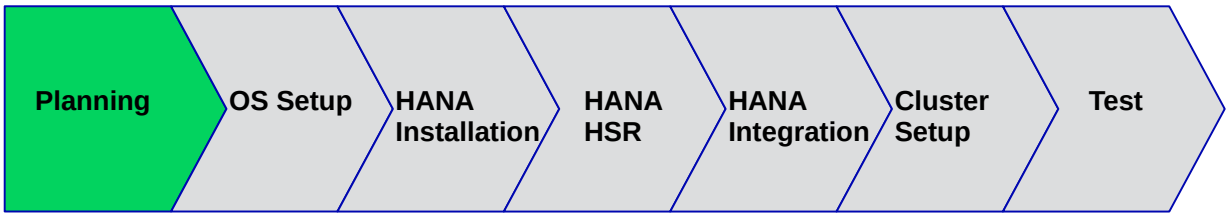

Planning the installation is essential for a successful SAP HANA cluster setup.

Before you start, you need the following:

- Software from SUSE: SUSE Linux Enterprise Server for SAP applications installation media, a valid subscription, and access to update channels
- Software from SAP: SAP HANA installation media
- Physical or virtual systems including disks
- Filled parameter sheet (see below *Section 4.2, “Parameter sheet”* (page 15))

4.1 Minimum lab requirements and prerequisites

Note

The minimum lab requirements mentioned here are by no means SAP sizing information. These data are provided only to rebuild the described cluster in a lab for test purposes. Even for tests the requirements can increase, depending on your test scenario. For productive systems ask your hardware vendor or use the official SAP sizing tools and services.

Note

Refer to SAP HANA TDI documentation for allowed storage configuration and file systems.

Requirements with one SAP system replication instance per site (1 : 1) - without a majority maker (2-node cluster):

- 2 VMs with each 32GB RAM, 50GB disk space for the system
- 1 shared disk for SBD with 10 MB disk space
- 2 data disks (one per site) with a capacity of each 96GB for SAP HANA
- 1 additional IP address for takeover
- 1 optional IP address for the read-enabled setup

Requirements with one SAP system replication instance per site (1 : 1) - with a majority maker (3-node cluster):

- 2 VMs with each 32GB RAM, 50GB disk space for the system
- 1 VM with 2GB RAM, 50GB disk space for the system
- 2 data disks (one per site) with a capacity of each 96GB for SAP HANA
- 1 additional IP address for takeover
- 1 optional IP address for the read-enabled setup

4.2 Parameter sheet

Even if the setup of the cluster organizing two SAP HANA sites is quite simple, the installation should be planned properly. You should have all needed parameters like SID, IP addresses and much more in place. It is good practice to first fill out the parameter sheet and then begin with the installation.

TABLE 1: PARAMETER SHEET FOR PLANNING

Parameter	Value	Role
Node 1		Cluster node name and IP address.
Node 2		Cluster node name and IP address.
Site A		Site name of the primary replicating SAP HANA database
Site B		Site name of the secondary replicating and the non-replicating SAP HANA database
SID		SAP System Identifier
Instance Number		Number of the SAP HANA database. For system replication also Instance Number+1 is blocked.
Network mask		
vIP primary		Virtual IP address to be assigned to the primary SAP HANA site
vIP secondary		Virtual IP address to be assigned to the read-enabled secondary SAP HANA site (optional)
Storage		Storage for HDB data and log files is connected “locally” (per node; not shared)
SBD		SBD fencing device (three for production), or diskless SBD
NTP Server		Address or name of your time server

TABLE 2: PARAMETER SHEET WITH VALUES USED IN THIS DOCUMENT

Parameter	Value	Role
Node 1	`suse01`, `192.168.1.11`	Cluster node name and IP address.
Node 2	`suse02`, `192.168.1.12`	Cluster node name and IP address.
SID	`HA1`	SAP System Identifier
Instance Number	`10`	Instance number of the SAP HANA database. For system replication also Instance Number+1 is blocked.
Network mask	`255.255.255.0`	
vIP primary	`192.168.1.20`	
vIP secondary	`192.168.1.21`	(optional)
Storage		Storage for HDB data and log files is connected “locally” (per node; not shared)
SBD	`/dev/disk/by-id/SBDA`	SBD fencing device (three for production), or diskless
NTP Server	pool pool.ntp.org	Address or name of your time server

5 Setting up the operating system

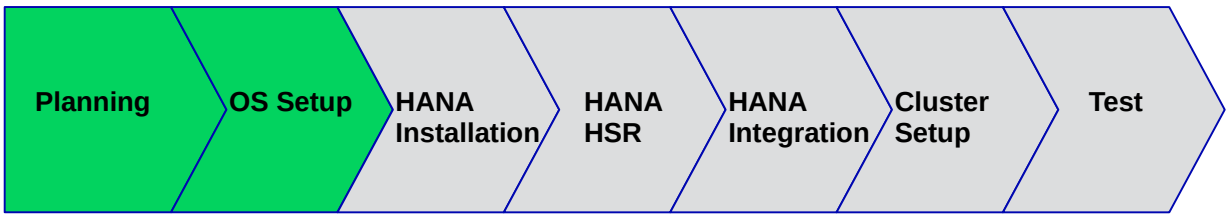

This section contains information you should consider during the installation of the operating system.

For the scope of this document, first SUSE Linux Enterprise Server for SAP applications is installed and configured. Then the SAP HANA database including the system replication is set up. Finally the automation with the cluster is set up and configured.

5.1 Installing SUSE Linux Enterprise Server for SAP applications

Multiple installation guides already exist, for different purposes and with different reasons to set up the server in a certain way. Below it is outlined where this information can be found. In addition, you will find important details you should consider to get a well-working system in place.

5.1.1 Installing the base operating system

Depending on your infrastructure and the hardware used, you need to adapt the installation. All supported installation methods and minimum requirement are described in the *Deployment Guide* for SUSE Linux Enterprise Server (https://documentation.suse.com/sles/16.0/html/SLES-deployment-pxe-client/). In case of automated installations you can find further information in the *Agama Guide* (https://documentation.suse.com/sles-sap/16.0/html/SAP-x86-64-agama-automated-installation/). The main installation guides for SUSE Linux Enterprise Server for SAP applications that fit all requirements for SAP HANA are available from the SAP notes:

- 3565382 SUSE Linux Enterprise Server 16: Installation Note
- 3577842 SAP HANA DB: Recommended OS settings for SLES 16 / SLES for SAP applications 16

5.1.2 Installing additional software

With SUSE Linux Enterprise Server for SAP applications, SUSE delivers special resource agents for SAP HANA. The pattern *sles_sap_HADB* already includes the `SAPHanaSR-angi` package. The pattern *ha_sles* summarizes all tools recommended to be installed on **all** nodes, including the *majority maker*.

If not already done, install the patterns *ha_sles* and *sles_sap_HADB* on all nodes.

```
suse01:~ # zypper in --type pattern ha_sles sles_sap_HADB
```

For more information, see section *Installation and Setup* of the SUSE Linux Enterprise High Availability Administration Guide.

6 Installing the SAP HANA Databases on both cluster nodes

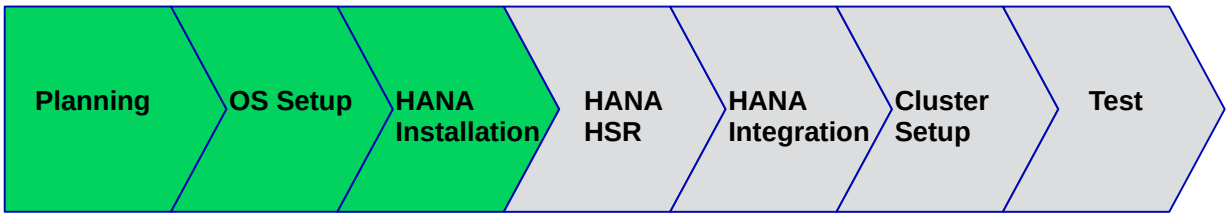

Even though this document focuses on the integration of an installed SAP HANA with system replication already set up into the Linux cluster, this chapter summarizes the test environment. Always use the official documentation from SAP to install SAP HANA and to set up the system replication.

This guide shows SAP HANA and saphostagent with native systemd integration.

PROCEDURE

1. Install the SAP HANA databases.
2. Check if the SAP hostagent is installed on all cluster nodes. If this SAP service is not installed, install it now.
3. Verify that both databases are up and running.

6.1 Installing the SAP HANA databases

- Read the SAP Installation and Setup Manuals available at "SAP for Me".
- Download the SAP HANA Software from the portal "SAP for Me".
- Install the SAP HANA database as described in the SAP HANA Server Installation Guide. The SAP HANA database client will be installed together with the server by default.

6.2 Checking if the SAP hostagent is installed on all cluster nodes

Check if the native `systemd`-enabled SAP hostagent and instance `sapstartsrv` are installed on all cluster nodes. If not, install and enable them now.

As Linux user *root*, use the command `systemctl` and `systemd-cgls` to check the SAP hostagent and instance services:

```
# systemctl list-unit-files "sap*"
UNIT FILE                         STATE           PRESET
saphostagent.service              enabled         disabled
saptune.service                   enabled         enabled
```

The mandatory `saphostagent` service is enabled. This is the installation default. Some more SAP related services might be enabled. For correct tuning the service `saptune` is required.

```
# systemctl list-unit-files "SAP*"
UNIT FILE                         STATE           PRESET
SAPHA1_10.service                 enabled         disabled
SAP.slice                         static          -
```

The SAP instance service SAPHA1_10.service is enabled and SAP.slice is available, as required.

6.3 Verifying both databases are up and running

```
# systemd-cgls -u SAP.slice
Unit SAP.slice (/SAP.slice):
├─saphostagent.service
│ ├─2630 /usr/sap/hostctrl/exe/saphostexec pf=/usr/sap/hostctrl/exe/host_profile -systemd
│ ├─2671 /usr/sap/hostctrl/exe/sapstartsrv pf=/usr/sap/hostctrl/exe/host_profile -D
│ └─3591 /usr/sap/hostctrl/exe/saposcol -l -w60 pf=/usr/sap/hostctrl/exe/host_profile
└─SAPHA1_10.service
  ├─ 1257 hdbcompileserver
  ├─ 1274 hdbpreprocessor
  ├─ 1353 hdbindexserver -port 31003
  ├─ 1356 hdbxsengine -port 31007
  ├─ 2077 hdbwebdispatcher
  ├─ 2300 hdbrsutil --start --port 31003 --volume 3 --volumesuffix mnt00001/
hdb00003.00003 --identifier 1644426276
  ├─28462 /usr/sap/HA1/HDB10/exe/sapstartsrv pf=/usr/sap/HA1/SYS/profile/HA1_HDB10_suse01
  ├─31314 sapstart pf=/usr/sap/HA1/SYS/profile/HA1_HDB10_suse01
  ├─31372 /usr/sap/HA1/HDB10/suse01/trace/hdb.sapHA1_HDB10 -d -nw -f /usr/sap/HA1/HDB10/
suse01/daemon.ini pf=/usr/sap/HA1/SYS/profile/HA1_HDB10_suse01
  ├─31479 hdbnameserver
  └─32201 hdbrsutil --start --port 31001 --volume 1 --volumesuffix mnt00001/hdb00001 --
identifier 1644426203
```

The SAP hostagent `saphostagent.service` and the instance´s `sapstartsrv` `SAPHA1_10.service` are running in the `SAP.slice`. See also manual pages systemctl(8) and systemd-cgls(8) for details.

7 Setting up SAP HANA System Replication

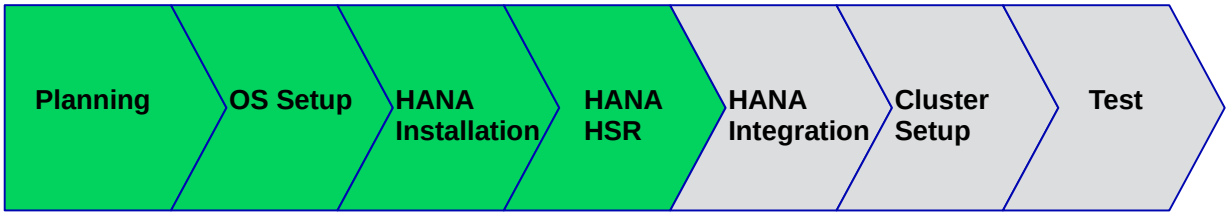

For more information read the section *Setting Up System Replication* of the SAP HANA Administration Guide.

Procedure

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

7.1 Backing up the primary database

Back up the primary database as described in the SAP HANA Administration Guide, section *SAP HANA Database Backup and Recovery*. We provide an example with SQL commands. You need to adapt these backup commands to match your backup infrastructure.

EXAMPLE 1: SIMPLE BACKUP FOR THE SYSTEM DATABASE AND ALL TENANTS WITH ONE SINGLE BACKUP CALL

> As user <*sid*>*adm* enter the following command:
>
> ```
> ~> hdbsql -i 10 -u SYSTEM -d SYSTEMDB \
>    "BACKUP DATA FOR FULL SYSTEM USING FILE ('backup')"
> ```
>
> You will get a command output similar to the following:
>
> ```
> 0 rows affected (overall time 15.352069 sec; server time 15.347745 sec)
> ```

EXAMPLE 2: SIMPLE BACKUP FOR A SINGLE CONTAINER (NON MDC) DATABASE

> Enter the following command as user <*sid*>*adm*:
>
> ```
> ~> hdbsql -i <instanceNumber> -u <dbuser> \
>    "BACKUP DATA USING FILE ('backup')"
> ```

Important

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

7.2 Enabling the primary node

As Linux user <*sid*>*adm*, enable the system replication at the primary node. You need to define a site name (like WDF). This site name must be unique for all SAP HANA databases which are connected via system replication. This means the secondary must have a different site name. The site names must not be changed later when the cluster has been activated.

Note

Do not use strings like "primary" and "secondary" as site names.

7.2.1 Enable primary node

Enable the primary using the `-sr_enable` option.

```
suse01:~> hdbnsutil -sr_enable --name=WDF
checking local nameserver:
checking for active nameserver ...
nameserver is running, proceeding ...
configuring ini files ...
successfully enabled system as primary site ...
done.
```

7.2.2 Check system replication configuration on primary node

Check the primary using the command `hdbnsutil -sr_stateConfiguration`.

```
suse01:~> hdbnsutil -sr_stateConfiguration --sapcontrol=1
SAPCONTROL-OK: <begin>
mode=primary
site id=1
site name=WDF
SAPCONTROL-OK: <end>
done.
```

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

7.3 Registering the secondary site

The SAP HANA database instance on the secondary side must be stopped before the instance 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`. Again, use the Linux user *<sid>adm*:

7.3.1 Stop the secondary database

To stop the secondary, you can use the command line tool *HDB*.

```
suse02:~> HDB stop
```

7.3.2 Copy KEY and KEY-DATA file from the primary to the secondary site

Beginning with SAP HANA 2.0, the system replication is running encrypted. The key files need to be copied-over from the primary to the secondary site.

```
~> cd /usr/sap/<SID>/SYS/global/security/rsecssfs
~> rsync -va {,<node1-siteB>:}$PWD/data/SSFS_<SID>.DAT
~> rsync -va {,<node1-siteB>:}$PWD/key/SSFS_<SID>.KEY
```

7.3.3 Register secondary site

The registration of the secondary is triggered by calling *hdbnsutil -sr_register* ….

```
...
suse02:~> hdbnsutil -sr_register --name=ROT \
     --remoteHost=suse01 --remoteInstance=10 \
     --replicationMode=sync --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 10).

Now start the database instance again and verify the system replication status. On the secondary node, the mode should be one of "SYNC" or "SYNCMEM". "ASYNC" is **not supported with automated cluster takeover**. The mode depends on the **replicationMode** option defined during the registration of the secondary.

7.3.4 Start secondary and check system replication configuration

To start the new secondary, use the command line tool *HDB*. Then check the SR configuration using `hdbnsutil -sr_stateConfiguration`.

```
suse02:~> HDB start
...
suse02:~> hdbnsutil -sr_stateConfiguration --sapcontrol=1
SAPCONTROL-OK: <begin>
mode=sync
site id=2
site name=ROT
active primary site=1
primary masters=suse01
SAPCONTROL-OK: <end>
done.
```

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

7.3.5 Check system replication status details

The python script *systemReplicationStatus.py* provides details about the current system replication.

```
suse01:~> cdpy; python3 systemReplicationStatus.py --sapcontrol=1
...
site/2/SITE_NAME=ROT
site/2/SOURCE_SITE_ID=1
site/2/REPLICATION_MODE=SYNC
site/2/REPLICATION_STATUS=ACTIVE
overall_replication_status=ACTIVE
site/1/REPLICATION_MODE=PRIMARY
site/1/SITE_NAME=WDF
local_site_id=1
```

```
...
```

7.4 Manually testing the SAP HANA SR takeover

Before you integrate your SAP HANA system replication into the HA cluster, it is mandatory to do a manual takeover. Testing without the cluster helps to make sure that basic operation (takeover and registration) is working as expected.

- Stop SAP HANA on node 1.
- Takeover SAP HANA to node 2.
- Register node 1 as secondary.
- Start SAP HANA on node 1.
- Wait until sync state is active.

7.5 Optional: Manually re-establishing SAP HANA SR to original state

Bring the systems back to the original state:

- Stop SAP HANA on node 2.
- Take over SAP HANA to node 1.
- Register node 2 as secondary.
- Start SAP HANA on node2.
- Wait until sync state is active.

8 Setting up SAP HANA HA/DR providers

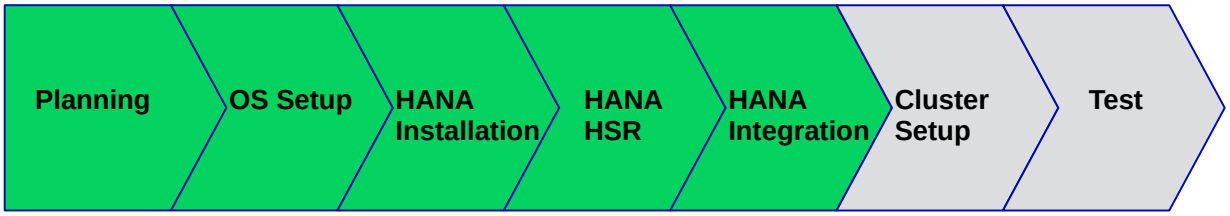

This step is mandatory to inform the cluster immediately if the secondary gets out of sync. The hook is called by SAP HANA using the HA/DR provider interface in point-of-time when the secondary gets out of sync. This is typically the case when the first commit pending is released. The hook is called by SAP HANA again when the system replication is back. This HA/DR provider method is **srConnectionChanged()**, the related SUSE hook script is *susHanaSR.py*. The hook script *susHanaSR.py* shall be used.

Another hook is called by SAP HANA before an SR takeover is processed. This method can be used to block a manual takeover during normal cluster operation. This HA/DR provider method is **preTakeover()**, the related SUSE hook script is *susTkOver.py*.

A third hook is called by SAP HANA when a service status changes. This method can be used to speed up the takeover in case the indexserver process fails. This HA/DR provider method is **srServiceStateChanged()**, the related SUSE hook script is *susChkSrv.py*.

Procedure

1. Implement the python hook script *susHanaSR.py* on both sites.
2. Implement the python hook script *susTkOver.py* on both sites.
3. Implement the python hook script *susChkSrv.py* on both sites.
4. Configure the system replication operation mode.
5. Allow <*sid*>*adm* to access the cluster.
6. Start SAP HANA.
7. Test the hook integration.

This will implement three SAP HANA HA/DR provider hook scripts. The hook script *susHanaSR.py* does not need any configuration parameters. The configuration for *susTkOver.py* normally does not need to be adapted. The configuration shown for *susChkSrv.py* is a good starting point. Any tuning should be aligned with the SAP experts.

Note

All hook scripts should be used directly from the `SAPHanaSR-angi` package. If the scripts are moved or copied, regular SUSE package updates will not work.

SAP HANA must be stopped to change the *global.ini* and allow SAP HANA to integrate the HA/DR hook scripts during start. Alternatively, `SAPHanaSR-manageProvider` might be used for adapting the *global.ini*. See manual page SAPHanaSR-manageProvider(8) for details.

8.1 Implementing susHanaSR hook for srConnectionChanged

Use the hook from the SAPHanaSR-angi package */usr/share/SAPHanaSR-angi/susHanaSR.py*. The hook must be configured on all SAP HANA cluster nodes. In *global.ini*, the section *[ha_dr_provider_sushanasr]* needs to be created. The section *[trace]* might be adapted. Refer to the manual page susHanaSR.py(7) for details on this HA/DR provider hook script, see also SAPHanaSR-manageProvider(8).

8.1.1 Stop SAP HANA

Stop SAP HANA either with *HDB* or using *sapcontrol*.

```
~> sapcontrol -nr <instanceNumber> -function StopSystem
```

8.1.2 Add SAPHanaSR via *global.ini*

Best is to use the SAP HANA tools for changing *global.ini*. Alternatively you may use `SAPHanaSR-manageProvider`, see manual page SAPHanaSR-manageProvider(8).

```
[ha_dr_provider_sushanasr]
provider = susHanaSR
path = /usr/share/SAPHanaSR-angi/
execution_order = 1

[trace]
ha_dr_sushanasr = info
```

8.2 Implementing susTkOver hook for preTakeover

Use the hook from the SAPHanaSR-angi package */usr/share/SAPHanaSR-angi/susTkOver.py*. The hook must be configured on all SAP HANA cluster nodes. In *global.ini*, the section *[ha_dr_provider_sustkover]* needs to be created. The section *[trace]* might be adapted. Refer to the manual page susTkOver.py(7) for details on this HA/DR provider hook script, see also SAPHanaSR-manageProvider(8).

8.2.1 Stop SAP HANA

Stop SAP HANA either with *HDB* or using *sapcontrol*.

```
~> sapcontrol -nr <instanceNumber> -function StopSystem
```

8.2.2 Add *susTkOver* via *global.ini*

Best is to use the SAP HANA tools for changing *global.ini*. Alternatively you may use `SAPHanaSR-manageProvider`, see manual page SAPHanaSR-manageProvider(8).

```
[ha_dr_provider_sustkover]
provider = susTkOver
path = /usr/share/SAPHanaSR-angi/
execution_order = 2

[trace]
ha_dr_sustkover = info
...
```

8.3 Implementing susChkSrv hook for srServiceStateChanged

Use the hook from the SAPHanaSR-angi package */usr/share/SAPHanaSR-angi/susChkSrv.py*. The hook must be configured on all SAP HANA cluster nodes. In *global.ini*, the section *[ha_dr_provider_suschk-srv]* needs to be created. For the progressive configuration, parameter action_on_lost=**fence** is used. The section *[trace]* might be adapted. Refer to the manual page susChkSrv.py(7) for details on this HA/DR provider hook script, see also SAPHanaSR-manageProvider(8). For details about the progressive configuration see manual page SAPHanaSR-angi-scenarios(7).

8.3.1 Stop SAP HANA

Stop SAP HANA either with *HDB* or using *sapcontrol*.

```
~> sapcontrol -nr <instanceNumber> -function StopSystem
```

8.3.2 Add *susChkSrv* via *global.ini*

Best is to use the SAP HANA tools for changing *global.ini*. Alternatively you may use `SAPHanaSR-manageProvider`, see manual page SAPHanaSR-manageProvider(8).

```
[ha_dr_provider_suschksrv]
provider = susChkSrv
path = /usr/share/SAPHanaSR-angi/
```

```
execution_order = 3
action_on_lost=fence

[trace]
ha_dr_suschksrv = info
...
```

8.4 Configuring system replication operation mode

When your system is connected as SAP HANA system replication target, you can find an entry in the *global.ini* which defines the operation mode. Up to now there are the following modes available:

- *logreplay*
- *logreplay_readaccess*
- *delta_datashipping*

Until a takeover and re-registration in the opposite direction, the entry for the operation mode is missing on your primary site. The first operation mode which was available was *delta_datashipping*. Today the preferred modes for HA are *logreplay* or *logreplay_readaccess*. Using the operation mode *logreplay* makes your secondary site in the SAP HANA system replication a hot standby system. For more details regarding all operation modes, check the available SAP documentation such as the guide "How To Perform System Replication for SAP HANA " (https://help.sap.com/docs/SAP_HANA_PLATFORM/6b94445c94ae495c83a19646e7c3fd56/676844172c2442f0bf6c8b080db05ae7.html).

8.4.1 Check the Operation Mode

Check both *global.ini* files and add the operation mode if needed. Check the section *system_replication* for entry ´operation_mode = logreplay´.

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

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

8.5 Allowing <sid>adm to access the cluster

The current version of the susHanaSR python hook uses the command `sudo` to allow the *<sid>adm* user to access the cluster attributes.

The user *<sid>adm* must be able to set the cluster attributes **hana_<sid>_site_srHook_***. The SAP HANA system replication hook needs password free access. The following example limits the sudo access to exactly setting the needed attribute. The entries can be added to a new file */etc/sudoers.d/SAPHanaSR* so that the original */etc/sudoers* file does not need to be edited. See manual page sudoers(5) for details.

Replace the <sid> by the **lowercase** SAP system ID (like `ha1`).

- Entry in sudo permissions /etc/sudoers.d/SAPHanaSR file

Basic sudoers entry to allow <sid>adm to use the hooks susHanaSR and susTkOver.

```
# SAPHanaSR-ScaleUp entries for writing srHook cluster attribute and SAPHanaSR-hookHelper
<sid>adm ALL=(ALL) NOPASSWD: /usr/sbin/crm_attribute -n hana_<sid>_*
<sid>adm ALL=(ALL) NOPASSWD: /usr/bin/SAPHanaSR-hookHelper --sid=<SID> *
```

- More specific sudoers entries to meet a high security level

All Cmnd_Alias entries must be each defined as a single line entry. In our example, we have five separate lines with Cmnd_Alias entries and one line for the *<sid>adm* user permitting the Cmnd_Aliases. In the document at hand, however, the separate lines of the example might include a line-break forced by document formatting. The alias identifier (for example SOK_SITEA) needs to be in capitals.

Replace the <sid> by the **lowercase** SAP system ID (like `ha1`). Replace the <SID> by the **uppercase** SAP system ID.

```
# SAPHanaSR-ScaleUp entries for writing srHook cluster attribute
Cmnd_Alias SOK_SITEA    = /usr/sbin/crm_attribute -n hana_<sid>_site_srHook_<siteA> -v
 SOK   -t crm_config -s SAPHanaSR
Cmnd_Alias SFAIL_SITEA  = /usr/sbin/crm_attribute -n hana_<sid>_site_srHook_<siteA> -v
 SFAIL -t crm_config -s SAPHanaSR
Cmnd_Alias SOK_SITEB    = /usr/sbin/crm_attribute -n hana_<sid>_site_srHook_<siteB> -v
 SOK   -t crm_config -s SAPHanaSR
Cmnd_Alias SFAIL_SITEB  = /usr/sbin/crm_attribute -n hana_<sid>_site_srHook_<siteB> -v
 SFAIL -t crm_config -s SAPHanaSR
Cmnd_Alias HOOK_HELPER  = /usr/bin/SAPHanaSR-hookHelper --sid=<SID> --case=checkTakeover

<sid>adm ALL=(ALL) NOPASSWD: SOK_SITEA, SFAIL_SITEA, SOK_SITEB, SFAIL_SITEB, HOOK_HELPER
```

8.6 Start SAP HANA

Start SAP HANA as user <sid>adm.

```
~> HDB start
```

8.7 Testing the hook integration

8.7.1 Check loading of the hook scripts

As user <sid>adm check the SAP HANA tracefiles, if the HA/DR provider scripts are loaded and called successfully during an SAP HANA event.

Check if script *susHanaSR.py* is loaded and initialized correctly:

```
~> cdtrace
~> grep HADR.*load.*susHanaSR nameserver_*.trc
~> grep susHanaSR.init nameserver_*.trc
```

Check if script *susTkOver.py* is loaded and initialized correctly:

```
~> cdtrace
~> grep HADR.*load.*susTkOver nameserver_*.trc
~> grep susTkOver.init nameserver_*.trc
```

Check if script *susChkSrv.py* is loaded and initialized correctly:

```
~> cdtrace
~> grep HADR.*load.*susChkSrv nameserver_*.trc
~> grep susChkSrv.init nameserver_*.trc
~> egrep '(LOST:|STOP:|START:|DOWN:|init|load|fail)' nameserver_suschksrv.trc
```

8.7.2 Check an **srConnectionChanged** event

After an event has been processed by the HA/DR provider script, check for the correct behavior.

As user <sid>adm check the SAP HANA tracefiles and verify if *susHanaSR.py* did successfully interact with the cluster or created a fallback file.

```
~> cdtrace
~> grep susHanaSR.srConnection.*CRM nameserver_*.trc
~> grep susHanaSR.srConnection.*fallback nameserver_*.trc
```

8.7.3 Check a **preTakeover** event

After an event has been processed by the HA/DR provider script, check for the correct behavior.

To test script *susTkOver.py* with stopped cluster, procedure *Section 7.4, "Manually testing the SAP HANA SR takeover"* (page 24) can be used here again. While the cluster is not set up at this moment the takeover will not be blocked.

As user root check the system messages and verify if the sudo permissions for *susTkOver.py* calling SAPHanaSR-hookHelper are set successfully.

```
# grep "sudo.*SAPHanaSR-hookHelper"  /var/log/messages
```

As user <sid>adm check the SAP HANA tracefiles, to verify if *susTkOver.py* did successfully block or permit the takeover request. First check for permitted takeover requests, then check for blocked takeover requests.

```
~> cdtrace
~> grep susTkOver.preTakeover.*permit nameserver_*.trc
~> grep susTkOver.preTakeover.*failed.*50277 nameserver_*.trc
```

9 Configuring the cluster

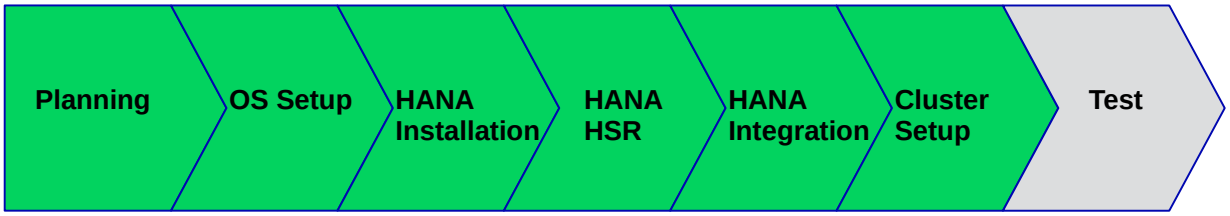

This chapter describes the configuration of the cluster software SUSE Linux Enterprise High Availability, which is part of SUSE Linux Enterprise Server for SAP applications, and the SAP HANA database integration.

ACTIONS

1. Basic cluster configuration
2. Configuration of cluster properties and resources
3. Testing the HA/DR provider hook integration

9.1 Configuring the basic cluster

The first step is to set up the basic cluster framework. For convenience, use the crmsh command `crm cluster init`.

9.1.1 Setting up watchdog for SBD fencing

If you use the storage-based fencing (SBD) mechanism (diskless or disk-based), you must also configure a watchdog. The watchdog is needed to reset a node if the system cannot longer access the SBD (diskless or disk-based). It is mandatory to configure the Linux system for loading a watchdog driver. It is strongly recommended to use a watchdog with hardware assistance (as is available on most modern systems), such as hpwdt, iTCO_wdt, or others. As fallback, you can use the softdog module.

Important: Setup for Watchdog

Access to the watchdog timer: No other software must access the watchdog timer; it can only be accessed by one process at any time. Some hardware vendors ship systems management software that use the watchdog for system resets. Such software must be disabled if the watchdog is to be used by SBD.

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|ibm)"
```

If you get a result, the system has already a loaded watchdog. If the watchdog does not match your watchdog device, you need to unload the module.

To safely unload the module, check first if an application is using the watchdog device.

```
# lsof /dev/watchdog
# rmmod <wrong_module>
```

Enable your watchdog module and make it persistent. For the example below, *softdog* has been used. However, *softdog* has some restrictions and should not be used as first option.

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

Check if the watchdog module is loaded correctly.

```
# lsmod | grep dog
# ls -l /dev/watchdog
```

```
# systool -vc watchdog
```

Testing the watchdog can be done with a simple action. Make sure to stop your SAP HANA first, because the watchdog will force an unclean reset or shutdown of your system.

Important

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

The following action will activate the watchdog without feeding it and should cause a reset/switch off:

```
# sync; cat /dev/watchdog & while date; do sleep 10; done
```

After your test was successful, you must implement the watchdog on all cluster members.

9.1.2 Setting up the initial cluster using `crm cluster init`

For more detailed information about setting up a cluster, refer to the sections *Setting Up the First Node* and *Adding the Second Node* of the Installation and Setup Quick Start for SUSE Linux Enterprise High Availability 16 0 at https://documentation.suse.com/sle-ha/16.0/html/HA-installing-two-node-cluster/.

This setup uses knet for corosync communication. Refer to the https://documentation.suse.com/sle-ha/16.0/ on detailed explanations of the term knet.

Create an initial setup, using the `crm cluster init` command, and follow the dialogs. Do this only on the first cluster node. Answer "no" to "Do you wish to configure a virtual IP address" and to "Do you want to configure QDevice".

```
suse01:~ # crm cluster init -s /dev/disk/by-id/SBDA
```

This command configures the basic cluster framework including:

- SSH keys
- csync2 to transfer configuration files
- SBD (one in this guide, better three for production)
- corosync (at least one ring, two for production)

Important

As requested by `crm cluster init`, change the password of the user hacluster.

9.1.3 Checking and adapting the corosync and SBD configuration

9.1.3.1 Checking the corosync configuration

Check the following blocks in the file */etc/corosync/corosync.conf*. The important parts are *knet* and the correct ring/IP configuration.

See also the example at the end of this document and refer to the manual pages corosync.conf(5), votequorum(5) and corosync_overview(8) for details on parameters and features.

```
totem {
    ...
    transport: knet
    ...
}

nodelist {
    node {
            ring0_addr: 192.168.1.11
            name: suse01
            nodeid: 1
    }

    node {
            ring0_addr: 192.168.1.12
            name: suse02
            nodeid: 2
    }
}
    ...
```

9.1.3.2 Adapting the SBD configuration

You can skip this section if you do not have any SBD devices, but be sure to implement another supported fencing mechanism.

See the manual pages sbd(8) and stonith_sbd(7) for details.

TABLE 3: SBD OPTIONS IN FILE /ETC/SYSCONFIG/SBD

Parameter	Description
SBD_WATCHDOG_DEV	Define the watchdog device. It is mandatory to use a watchdog. SBD does not work reliable without watchdog. Refer to the SLES manual and SUSE TID 7016880 for setting up a watchdog.
SBD_WATCHDOG_TIMEOUT	This parameter is used with diskless SBD. It defines the timeout, in seconds, the watchdog will wait before panicking the node if noone tickles it. The value of fencing-watchdog-timeout in CIB should be more than twice the value of SBD_WATCHDOG_TIMEOUT.
SBD_STARTMODE	Start mode. If set to `clean`, sbd will only start if the node was previously shut down cleanly or if the slot is empty.
SBD_DELAY_START	Whether to delay after starting sbd on boot. Should be `no` for SBD_STSRTMODE `clean`.
SBD_PACEMAKER	Check Pacemaker quorum and node health.

In the following, replace /dev/disk/by-id/SBDA by your real sbd device names. As an example, the SBD_WATCHDOG_TIMEOUT is set to 15s.

```
# egrep -v "(^#|^$)" /etc/sysconfig/sbd
SBD_PACEMAKER="yes"
SBD_STARTMODE="clean"
SBD_DELAY_START="no"
SBD_WATCHDOG_DEV="/dev/watchdog"
SBD_WATCHDOG_TIMEOUT="15"
SBD_TIMEOUT_ACTION="flush,reboot"
SBD_MOVE_TO_ROOT_CGROUP="auto"
SBD_OPTS=""
SBD_DEVICE="/dev/disk/by-id/SBDA"
```

Important

Also read the SUSE product documentation about calculation of timeouts for more details: https://documentation.suse.com/sle-ha/16.0/html/HA-sbd-configuring-diskbased/index.html#ha-sbd-setting-up-watchdog

9.1.3.3 Verifying the SBD device

You can skip this section if you do not have any SBD devices, but make sure to implement a supported fencing mechanism.

It is a good practice to check if the SBD device can be accessed from both nodes and does contain valid records. Check this for all devices configured in */etc/sysconfig/sbd*. You can do so, for example, by calling `cs_show_sbd_devices`.

```
suse01:~ # cs_show_sbd_devices
==Dumping header on disk /dev/disk/by-id/SBDA
Header version     : 2.1
UUID               : c2d5b54e-aa7e-421d-8f08-ed929499cd84
Number of slots    : 255
Sector size        : 512
Timeout (watchdog) : 15
Timeout (allocate) : 2
Timeout (loop)     : 1
Timeout (msgwait)  : 30
==Header on disk /dev/disk/by-id/SBDA is dumped
0       suse01 clear
```

Important

The timeout values in our example are only start values. It is a requirement that they are tuned to your environment. Refer to the TIDs 7011346 and 7023689 for more information.

For more information on SBD configuration parameters, consult the respective sections of the SUSE Linux Enterprise High Availability Administration Guide and the TIDs 7016880 and 7008216.

9.1.4 Configuring the cluster on the second node

The second node of the two nodes cluster can be integrated by starting the command `crm cluster join`. This command asks for the IP address or name of the first cluster node. With this command, all needed configuration files are copied over. As a result, the cluster is started on both nodes.

```
# crm cluster join -c <host1>
```

Press *RETURN* to acknowledge the IP address.

9.1.5 Checking the cluster for the first time

Now it is time to check and optionally start the cluster for the first time on both nodes.

```
suse01:~ # systemctl status pacemaker
suse01:~ # systemctl status sbd
suse02:~ # systemctl status pacemaker
suse02:~ # systemctl status sbd
suse01:~ # crm cluster stop --all
suse01:~ # crm cluster start --all
```

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

```
# crm_mon -r1
```

The command will show the "empty" cluster and will print something similar to the screen output below. The most interesting pieces of information for now are that there are two nodes in the status "online", the message "partition with quorum", and a running SBD resource.

```
Cluster Summary:
  * Stack: corosync (Pacemaker is running)
  * Current DC: suse01 (version
 3.0.2+20260429.40d19b75-160099.1.8-3.0.2+20260429.40d19b75) - partition with quorum
  * Last updated: Fri Jul 24 15:05:26 2026 on suse01
  * Last change:  Tue Jul 14 13:36:53 2026 by hacluster via crmd on suse01
  * 2 nodes configured
  * 1 resource instance configured

Node List:
  * Online: [ suse01 suse02  ]

Full List of Resources:
  * fencing-sbd (stonith:fence_sbd):  Started suse01
```

9.2 Configuring cluster properties and resources

This section describes how to configure constraints, resources, bootstrap, and fencing, using the `crm configure` shell command as described in the SUSE Linux Enterprise High Availability Administration Guide.

Use the command `crm` to add the objects to the cluster information base (CIB). Copy the following examples to a local file, edit the file and then load the configuration to the CIB:

```
suse01:~ # vi crm-fileXX
suse01:~ # crm configure load update crm-fileXX
```

9.2.1 Cluster bootstrap and more

The first example defines the cluster bootstrap options, the resource and operation defaults. The fencing-timeout should be greater than 1.2 times the SBD on-disk msgwait timeout. The priority-fencing-delay should be at least 2 times the SBD CIB pcmk_delay_max.

```
suse01:~ # vi crm-bs.txt
# enter the following to crm-bs.txt
property cib-bootstrap-options: \
    fencing-enabled="true" \
    fencing-action="reboot" \
    fencing-timeout="150" \
    priority-fencing-delay="30"
rsc_defaults build-resource-defaults: \
    resource-stickiness="1000" \
    migration-threshold="5000"
op_defaults op-options: \
    timeout="600" \
    record-pending=true
```

Now add the configuration to the cluster.

```
suse01:~ # crm configure load update crm-bs.txt
```

9.2.2 Fencing device

Skip this section if you are using diskless SBD.

The next configuration part defines an SBD disk fencing resource.

```
suse01:~ # vi crm-sbd.txt
# enter the following to crm-sbd.txt
```

```
primitive fencing-sbd stonith:fence_sbd \
        params pcmk_delay_max=15
```

Again we add the configuration to the cluster.

```
suse01:~ # crm configure load update crm-sbd.txt
```

9.2.3 Using other fencing mechanisms

This section is only relevant if the recommended disk-based or diskless SBD fencing is not used.

We recommend to use SBD (best practice) or IPMI (second choice) as fencing mechanism. The SUSE Linux Enterprise High Availability product also supports additional fencing mechanism not covered here.

For further information about fencing, read the SUSE Linux Enterprise High Availability Administration Guide . For public cloud environments, refer to your cloud provider's documentation on supported fencing mechanisms.

9.2.4 SAPHanaTopology

This step is to define the resources needed, to analyze the SAP HANA topology for the replicated pair. Prepare the changes in a text file, for example *crm-saphanatop.txt*, and load it with the command:

`crm configure load update crm-saphanatop.txt`

```
suse01:~ # vi crm-saphanatop.txt
# enter the following to crm-saphanatop.txt
primitive rsc_SAPHanaTop_HA1_HDB10 ocf:suse:SAPHanaTopology \
    op start interval=0 timeout=600 \
    op stop interval=0 timeout=300 \
    op monitor interval=50 timeout=600 \
    params SID=HA1 InstanceNumber=10
clone cln_SAPHanaTop_HA1_HDB10 rsc_SAPHanaTop_HA1_HDB10 \
    meta clone-node-max=1 interleave=true
```

Additional information about all parameters can be found in manual page ocf_suse_SAPHanaTopology(7).

Again, add the configuration to the cluster.

```
suse01:~ # crm configure load update crm-saphanatop.txt
```

The most important parameters here are SID and InstanceNumber, which are quite self explaining in the SAP context. Beside these parameters, typical tunables are the timeout values for the operations (start, monitor, stop).

9.2.5 SAPHanaFilesystem

This step is to define the resources to monitor the file system used by HANA, for example /hana/shared/<SID>. The RA just monitors the file system, but neither does mount nor umount it. Mounting and umounting is done by the OS thru /etc/fstab. Prepare the changes in a text file, for example *crm-saphanafil.txt*, and load it with the command:

`crm configure load update crm-saphanafil.txt`

```
suse01:~ # vi crm-saphanafil.txt
# enter the following to crm-saphanafil.txt
primitive rsc_SAPHanaFil_HA1_HDB10 ocf:suse:SAPHanaFilesystem \
    op start interval=0 timeout=10 \
    op stop interval=0 timeout=20 \
    op monitor interval=120 timeout=120 \
    params SID=HA1 InstanceNumber=10 ON_FAIL_ACTION=fence
clone cln_SAPHanaFil_HA1_HDB10 rsc_SAPHanaFil_HA1_HDB10 \
    meta clone-node-max=1 interleave=true
```

Additional information about all parameters can be found in manual page ocf_suse_SAPHanaFilesystem(7).

Again, add the configuration to the cluster.

```
suse01:~ # crm configure load update crm-saphanafil.txt
```

The most important parameters here are SID and InstanceNumber, which are quite self explaining in the SAP context. ON_FAIL_ACTION defines how the RA should react on monitor failures. Beside these parameters, typical tunables are the timeout values for the operations (start, monitor, stop).

9.2.6 SAPHanaController

This step is to define the resource needed, to control the replicated SAP HANA pair. Edit the changes in a text file, for example *crm-saphanacon.txt*, and load it with the following command:

`crm configure load update crm-saphanacon.txt`

TABLE 4: DESCRIPTION OF IMPORTANT RESOURCE AGENT PARAMETERS

Parameter	Description
PREFER_SITE_TAKEOVER	Defines whether RA should prefer to take over to the secondary instance instead of restarting the failed primary locally.

Parameter	Description
AUTOMATED_REGISTER	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. 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.
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 instances 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, the former primary will be registered immediately. After this registration to the new primary, all data will be overwritten by the system replication.
ON_FAIL_ACTION	Defines how the RA escalates monitor failures on an HANA primary node. If srHook=SOK, in case of monitor failure, a node fencing could be triggered. For srHook=SFAIL, the restart will be proceeded as usual. ON_FAIL_ACTION=**fence** may speed up takeover, depending on how long HANA needs for stopping.

Additional information about all parameters of the SAPHanaController RA can be found in manual page ocf_suse_SAPHanaController(7).

```
suse01:~ # vi crm-saphanacon.txt
# enter the following to crm-saphanacon.txt
primitive rsc_SAPHanaCon_HA1_HDB10 ocf:suse:SAPHanaController \
    op start interval=0 timeout=3600 \
    op stop interval=0 timeout=600 \
    op promote interval=0 timeout=900 \
    op demote interval=0 timeout=320 \
```

```
    op monitor interval=60 role=Promoted timeout=700 \
    op monitor interval=61 role=Unpromoted timeout=700 \
    params SID=HA1 InstanceNumber=10 PREFER_SITE_TAKEOVER=true \
        DUPLICATE_PRIMARY_TIMEOUT=7200 AUTOMATED_REGISTER=false \
        ON_FAIL_ACTION=fence \
    meta priority=100
clone mst_SAPHanaCon_HA1_HDB10 rsc_SAPHanaCon_HA1_HDB10 \
    meta clone-node-max=1 promotable=true interleave=true maintenance=true
```

Now add the configuration to the cluster.

```
suse01:~ # crm configure load update crm-saphanacon.txt
```

The most important parameters here are again SID and InstanceNumber. Beside these parameters the timeout values for the operations (start, promote, monitors, stop) are typical tunables.

9.2.7 Adding a virtual IP address for the primary site

The last resource to be added is covering the virtual IP address. For details, see manual page ocf_heartbeat_IPaddr2(7).

```
# vi crm-vip.txt
# enter the following to crm-vip.txt
primitive rsc_ip_HA1_HDB10 ocf:heartbeat:IPaddr2 \
    op monitor interval=10 timeout=20 \
    params ip=192.168.1.20
```

Load the file to the cluster.

```
suse01:~ # crm configure load update crm-vip.txt
```

In most on-premise installations, only the parameter *ip* needs to be set to the virtual IP address to be presented to the client systems. Public cloud environments often need specific settings.

9.2.8 Constraints for SAPHanaSR-angi

Two constraints are organizing the correct placement of the virtual IP address for the client database access and the start order between the resource agents SAPHanaController and SAPHanaTopology.

```
suse01:~ # vi crm-cs.txt
# enter the following to crm-cs.txt
location loc_saphana_ip_HA1_HDB10 rsc_ip_HA1_HDB10 \
    rule 2000: master-rsc_SAPHanaCon_HA1_HDB10 eq 150
order ord_saphana_HA1_HDB10 Optional: cln_SAPHanaTop_HA1_HDB10 \
```

```
    mst_SAPHanaCon_HA1_HDB10
```

Load the file to the cluster.

```
suse01:~ # crm configure load update crm-cs.txt
```

9.2.9 Activating multi-state resource for cluster operation

The multi state resource was added to the cluster with in maintenance mode. To get the resource operated by the cluster the maintenance must be ended by the command:

```
suse01:~ # crm resource refresh mst_SAPHanaCon_HA1_HDB10
suse01:~ # cs_wait_for_idle -s 5
suse01:~ # crm resource maintenance mst_SAPHanaCon_HA1_HDB10 off
```

The command `cs_wait_for_idle` is part of the package `ClusterTools2`. For more details, see manual pages cs_wait_for_idle(8), crm(8), SAPHanaSR_maintenance_examples(7).

9.2.10 Active/active read-enabled scenario

This step is optional. If you have an active/active SAP HANA system replication with a read-enabled secondary, it is possible to integrate the needed second virtual IP address into the cluster. This is done by adding a second virtual IP address resource and a location constraint binding the address to the secondary site.

```
suse01:~ # vi crm-re.txt
# enter the following to crm-re.txt
primitive rsc_ip_HA1_HDB10_readenabled ocf:heartbeat:IPaddr2 \
    op monitor interval=10 timeout=20 \
    params ip=192.168.1.21
location loc_saphana_ip_readenabled_HA1_HDB10 rsc_ip_HA1_HDB10_readenabled \
    rule 2000: master-rsc_SAPHanaCon_HA1_HDB10 eq 100
```

Load the file to the cluster.

```
suse01:~ # crm configure load update crm-re.txt
```

10 Testing the cluster

The chapter "Tests" refers to our two-node lab cluster *suse12*. The two nodes are named *suse01* and *suse02*. SAP HANA has SID *HA1* and instance number *10*.

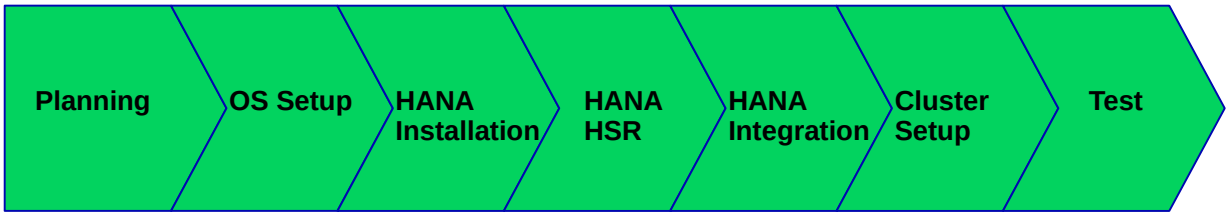

The lists of tests will be further enhanced in one of the next updates of this document.

As with any cluster testing is crucial. Make sure that all test cases derived from customer expectations are implemented and passed fully. Otherwise the project is likely to fail in production.

The test prerequisite, if not described differently, is always that both nodes are booted, normal members of the cluster and the HANA RDBMS is running. The system replication is in sync (SOK).

In the following test descriptions, we assume resource agents and HA/DR hooks are configured for the progressive variant. See manual page SAPHanaSR-angi-scenarios(7) for details. The SAPHanaController uses PREFER_SITE_TAKEOVER=true, AUTOMATED_REGISTER=true and ON_FAIL_ACTION=fence. The susChkSrv uses action_on_lost=fence. The SBD uses SBD_STARTMODE=clean and SBD_DELAY_START=0. The pacemaker service is disabled.

Note

The following tests are designed to be run in sequence and depend on the exit state of the preceding tests.

10.1 Test: Stop primary database on site A (node 1)

EXAMPLE 3: TEST STOP PRIMARY SITE A

COMPONENT:

- Primary Database

DESCRIPTION:

- The primary SAP HANA database is stopped during normal cluster operation.

TEST PROCEDURE:

1. Stop the primary SAP HANA database gracefully as _<sid>adm_.

```
node 1:~> HDB stop
```

RECOVERY PROCEDURE:

1. Clean up the sbd device (or devices) - to be called on any cluster node

```
# cs_clear_sbd_devices --all
```

2. Start the cluster on failed node - to be called on any cluster node

```
# crm cluster start --all
```

EXPECTED:

1. The cluster detects the stopped primary SAP HANA database (on node 1) and marks the resource failed.
2. The node 1 gets fenced.
3. The cluster promotes the secondary SAP HANA database (on node 2) to take over as primary.
4. The cluster migrates the IP address to the new primary (on node 2).
5. The failed node 1 does not join the cluster - **admin to process the recovery procedure**
6. When node 1 joins back to the cluster, SAP HANA is registered as new secondary and started by the cluster.
7. After some time, the new secondary is marked as in sync (SOK).
8. The cluster "failed actions" are cleaned up after following the recovery procedure.

10.2 Test: Stop primary database on site B (node 2)

EXAMPLE 4: TEST STOP PRIMARY SITE B

COMPONENT:

- Primary Database

DESCRIPTION:

- The primary SAP HANA database is stopped during normal cluster operation.

TEST PROCEDURE:

1. Stop the primary SAP HANA database gracefully as _<sid>adm_.

```
node 2:~> HDB stop
```

RECOVERY PROCEDURE:

1. Clean up the sbd device (or devices) - to be called on any cluster node

```
# cs_clear_sbd_devices --all
```

2. Start the cluster on failed node - to be called on any cluster node

```
# crm cluster start --all
```

EXPECTED:

1. The cluster detects the stopped primary SAP HANA database (on node 2) and marks the resource failed.
2. The node 2 gets fenced.
3. The cluster promotes the secondary SAP HANA database (on node 1) to take over as primary.
4. The cluster migrates the IP address to the new primary (on node 1).
5. The failed node 2 does not join the cluster - **admin to process the recovery procedure**
6. When node 2 joins back to the cluster, SAP HANA is registered as new secondary and started by the cluster.
7. After some time, the new secondary is marked as in sync (SOK).
8. The cluster "failed actions" are cleaned up after following the recovery procedure.

10.3 Test: Crash primary database on site A (node 1)

EXAMPLE 5: TEST CRASH PRIMARY DATABASE SITE A

COMPONENT:

- Primary Database

DESCRIPTION:

- Simulate a complete break-down of the primary database system.

TEST PROCEDURE:

1. Kill the primary database system using signals as _<sid>adm_.

```
node 1:~> HDB kill-9
```

RECOVERY PROCEDURE:

1. Clean up the sbd device (or devices) - to be called on any cluster node

```
# cs_clear_sbd_devices --all
```

2. Start the cluster on failed node - to be called on any cluster node

```
# crm cluster start --all
```

EXPECTED:

1. The cluster detects the stopped primary SAP HANA database (on node 1) and marks the resource failed.
2. The node 1 gets fenced.
3. The cluster promotes the secondary SAP HANA database (on node 2) to take over as primary.
4. The cluster migrates the IP address to the new primary (on node 2).
5. The failed node 1 does not join the cluster - **admin to process the recovery procedure**
6. When node 1 joins back to the cluster, SAP HANA is registered as new secondary and started by the cluster.
7. After some time, the new secondary is marked as in sync (SOK).
8. The cluster "failed actions" are cleaned up after following the recovery procedure.

10.4 Test: Crash primary database on site B (node 2)

EXAMPLE 6: TEST CRASH PRIMARY DATABASE SITE B

COMPONENT:

- Primary Database

DESCRIPTION:

- Simulate a complete break-down of the primary database system.

TEST PROCEDURE:

1. Kill the primary database system using signals as _<sid>adm_.

```
node 2:~> HDB kill-9
```

RECOVERY PROCEDURE:

1. Clean up the sbd device (or devices) - to be called on any cluster node

```
# cs_clear_sbd_devices --all
```

2. Start the cluster on failed node - to be called on any cluster node

```
# crm cluster start --all
```

EXPECTED:

1. The cluster detects the stopped primary SAP HANA database (on node 2) and marks the resource failed.
2. The node 2 gets fenced.
3. The cluster promotes the secondary SAP HANA database (on node 1) to take over as primary.
4. The cluster migrates the IP address to the new primary (on node 1).
5. The failed node 2 does not join the cluster - **admin to process the recovery procedure**
6. When node 2 joins back to the cluster, SAP HANA is registered as new secondary and started by the cluster.
7. After some time, the new secondary is marked as in sync (SOK).
8. The cluster "failed actions" are cleaned up after following the recovery procedure.

10.5 Test: Crash primary node on site A (node 1)

EXAMPLE 7: TEST CRASH PRIMARY NODE SITE A

COMPONENT:

- Cluster node of primary site

DESCRIPTION:

- Simulate a crash of the primary site node running the primary SAP HANA database.

TEST PROCEDURE:

1. Crash the primary node by sending a 'fast-reboot' system request.

```
node 1:~ # sync; echo b > /proc/sysrq-trigger
```

RECOVERY PROCEDURE:

1. Clean up the sbd device (or devices) - to be called on any cluster node

```
# cs_clear_sbd_devices --all
```

2. Start the cluster on failed node - to be called on any cluster node

```
# crm cluster start --all
```

EXPECTED:

1. The cluster detects missing node 1
2. The node 1 gets fenced.
3. The cluster promotes the secondary SAP HANA database (on node 2) to take over as primary.
4. The cluster migrates the IP address to the new primary (on node 2).
5. The failed node 1 does not join the cluster - **admin to process the recovery procedure**
6. When node 1 joins back to the cluster, SAP HANA is registered as new secondary and started by the cluster.
7. After some time, the new secondary is marked as in sync (SOK).
8. The cluster "failed actions" are cleaned up after following the recovery procedure.

10.6 Test: Crash primary node on site B (node 2)

EXAMPLE 8: TEST CRASH PRIMARY NODE SITE B

COMPONENT:

- Cluster node of primary site

DESCRIPTION:

- Simulate a crash of the primary site node running the primary SAP HANA database.

TEST PROCEDURE:

1. Crash the primary node by sending a 'fast-reboot' system request.

```
node 2:~ # sync; echo b > /proc/sysrq-trigger
```

RECOVERY PROCEDURE:

1. Clean up the sbd device (or devices) - to be called on any cluster node

```
# cs_clear_sbd_devices --all
```

2. Start the cluster on failed node - to be called on any cluster node

```
# crm cluster start --all
```

EXPECTED:

1. The cluster detects missing node 2
2. The node 2 gets fenced.
3. The cluster promotes the secondary SAP HANA database (on node 1) to take over as primary.
4. The cluster migrates the IP address to the new primary (on node 1).
5. The failed node 2 does not join the cluster - **admin to process the recovery procedure**
6. When node 2 joins back to the cluster, SAP HANA is registered as new secondary and started by the cluster.
7. After some time, the new secondary is marked as in sync (SOK).
8. The cluster "failed actions" are cleaned up after following the recovery procedure.

10.7 Test: Crash the primary database indexserver on site A (node 1)

EXAMPLE 9: TEST CRASH PRIMARY INDEXSERVER SITE A

COMPONENT:

- Primary database

DESCRIPTION:

- Simulate a crash of the primary SAP HANA database indexserver.

TEST PROCEDURE:

1. Crash the primary indexserver by killing the hdbindexserver process.

```
node 1:~ # killall -9 hdbindexserver
```

RECOVERY PROCEDURE:

1. Clean up the sbd device (or devices) - to be called on any cluster node

```
# cs_clear_sbd_devices --all
```

2. Start the cluster on failed node - to be called on any cluster node

```
# crm cluster start --all
```

EXPECTED:

1. The HA/DR hook susChkSrv fences node 1
2. The cluster promotes the secondary SAP HANA database (on node 2) to take over as primary.
3. The cluster migrates the IP address to the new primary (on node 2).
4. The failed node 1 does not join the cluster - **admin to process the recovery procedure**
5. When node 1 joins back to the cluster, SAP HANA is registered as new secondary and started by the cluster.
6. After some time, the new secondary is marked as in sync (SOK).
7. The cluster "failed actions" are cleaned up after following the recovery procedure.

10.8 Test: Crash the primary database indexserver on site B (node 2)

EXAMPLE 10: TEST CRASH PRIMARY INDEXSERVER SITE B

COMPONENT:

- Primary database

DESCRIPTION:

- Simulate a crash of the primary SAP HANA database indexserver.

TEST PROCEDURE:

1. Crash the primary indexserver by killing the hdbindexserver process.

```
node 2:~ # killall -9 hdbindexserver
```

RECOVERY PROCEDURE:

1. Clean up the sbd device (or devices) - to be called on any cluster node

```
# cs_clear_sbd_devices --all
```

2. Start the cluster on failed node - to be called on any cluster node

```
# crm cluster start --all
```

EXPECTED:

1. The HA/DR hook susChkSrv fences node 2
2. The cluster promotes the secondary SAP HANA database (on node 1) to take over as primary.
3. The cluster migrates the IP address to the new primary (on node 1).
4. The failed node 2 does not join the cluster - **admin to process the recovery procedure**
5. When node 2 joins back to the cluster, SAP HANA is registered as new secondary and started by the cluster.
6. After some time, the new secondary is marked as in sync (SOK).
7. The cluster "failed actions" are cleaned up after following the recovery procedure.

10.9 Test: Stop the secondary database on site B (node 2)

EXAMPLE 11: TEST STOP SECONDARY DATABASE SITE B

COMPONENT:

- Secondary database

DESCRIPTION:

- The secondary SAP HANA database is stopped during normal cluster operation

TEST PROCEDURE:

1. Stop the secondary SAP HANA database gracefully as _<sid>adm_.

```
node 2:~> HDB stop
```

RECOVERY PROCEDURE:

1. Refresh the failed resource status of the secondary SAP HANA database (on node 2) as root.

```
# crm resource refresh rsc_SAPHanaCon_HA1_HDB10 suse02
```

EXPECTED:

1. The cluster detects the stopped secondary SAP HANA database (on node 2) and marks the resource failed.
2. The cluster detects the broken system replication and marks it as failed (SFAIL).
3. The cluster restarts the secondary SAP HANA database on the same node (node 2).
4. The cluster detects that the system replication is in sync again and marks it as ok (SOK).
5. The cluster "failed actions" are cleaned up after following the recovery procedure.

10.10 Test: Crash the secondary database on site B (node 2)

EXAMPLE 12: TEST STOP SECONDARY DATABASE SITE B

COMPONENT:

- Secondary database

DESCRIPTION:

- Simulate a complete breakdown of the secondary database system.

TEST PROCEDURE:

1. Kill the secondary database system using signals as _<sid>adm_.

```
node 2:~> HDB kill-9
```

RECOVERY PROCEDURE:

1. Refresh the failed resource status of the secondary SAP HANA database (on node 2) as root.

```
# crm resource refresh rsc_SAPHanaCon_HA1_HDB10 suse02
```

EXPECTED:

1. The cluster detects the stopped secondary SAP HANA database (on node 2) and marks the resource failed.
2. The cluster detects the broken system replication and marks it as failed (SFAIL).
3. The cluster restarts the secondary SAP HANA database on the same node (node 2).

4. The cluster detects that the system replication is in sync again and marks it as ok (SOK).
5. The cluster "failed actions" are cleaned up after following the recovery procedure.

10.11 Test: Crash secondary node on site B (node 2)

EXAMPLE 13: TEST CRASH SECONDARY NODE SITE B

COMPONENT:

- Cluster node of secondary site

DESCRIPTION:

- Simulate a crash of the secondary site node running the secondary SAP HANA database.

TEST PROCEDURE:

1. Crash the primary node by sending a 'fast-reboot' system request.

```
node 2:~ # sync; echo b > /proc/sysrq-trigger
```

RECOVERY PROCEDURE:

1. Clean up the sbd device (or devices) - to be called on any cluster node

```
# cs_clear_sbd_devices --all
```

2. Start the cluster on failed node - to be called on any cluster node

```
# crm cluster start --all
```

EXPECTED:

1. The cluster detects missing node 2
2. The node 2 gets fenced.
3. After some time, the cluster shows the `srHook` attribute on the missing secondary site B (node 2) as `SFAIL`.
4. The failed node 2 does not join the cluster - **admin to process the recovery procedure**
5. When node 2 joins back to the cluster, SAP HANA is started by the cluster.
6. After some time, the new secondary is marked as in sync (SOK).
7. The cluster "failed actions" are cleaned up after following the recovery procedure.

10.12 Test: Failure of replication LAN

EXAMPLE 14: TEST FAIL_NETWORK_SR

COMPONENT:

- Replication LAN

DESCRIPTION:

- Loss of replication LAN connectivity between the primary and secondary node.

TEST PROCEDURE:

1. Block HANA system replication on the secondary node (node 2). The below examples assumes HANA instance number 00 for the ports to be blocked.

```
# ss -tulpan | grep hdb.*server | grep -v 127.0.0.1 |\
  grep -v "\*" | awk '{print $6}' | awk -F: '{print $2}'
...
# nft add table inet t_test
# nft add chain inet t_test c_test \
    "{ type filter hook input priority filter; }"
# nft add rule inet t_test c_test tcp sport \
    { 40001, 40002, 40003, 40007 } drop
# nft list ruleset
table inet t_test {
    type filter hook input priority filter; policy accept;
    chain c_test {
        tcp sport { 40001, 40002, 40003, 40007 } drop
    }
}
# SAPHanaSR-showStatus -w
...
```

2. Wait for srHook going to SFAIL for the secondary node.

RECOVERY PROCEDURE:

1. Re-establish the connection between the cluster nodes on the replication LAN.

```
# nft delete table inet t_test
```

EXPECTED:

1. After some time, the cluster shows the srHook of the secondary (on node 2) as SFAIL.

2. The primary SAP HANA database (node 1) "cdpy; python3 systemReplicationStatus.py" shows "CONNECTION TIMEOUT" and the secondary SAP HANA database (node 2) is not able to reach the primary database (node 1).
3. The primary SAP HANA database continues to operate, but no system replication takes place and there is no takeover destination.
4. When the LAN connection is re-established, SAP HANA automatically detects connectivity between the databases and restarts the system replication.
5. After some time, the secondary is marked as in sync (SOK).

10.13 Test: Failure of cluster communication LAN

EXAMPLE 15: TEST FAIL_NETWORK_CL

COMPONENT:

- Cluster communication LAN

DESCRIPTION:

- Loss of cluster connectivity between the primary and secondary node.

TEST PROCEDURE:

1. Block cluster communication on the secondary node.

```
# nft add table inet t_test
# nft add chain inet t_test c_test \
    "{ type filter hook input priority filter; }"
# nft add rule inet t_test c_test udp dport \
    { 5404, 5405, 5406, 5407, 5408, 5409, 5410, 5411, 5412 } drop
# exit
```

2. Wait for secondary node to get fenced.

RECOVERY PROCEDURE:

1. Clean up the sbd device (or devices) - to be called on any cluster node

```
# cs_clear_sbd_devices --all
```

2. Start the cluster on failed node - to be called on any cluster node

```
# crm cluster start --all
```

EXPECTED:

1. The secondary node (node 2) gets fenced.
2. The fenced node 2 does not join the cluster - **admin to process the recovery procedure**
3. When node 2 joins back to the cluster, SAP HANA is started by the cluster.
4. After some time, the secondary is marked as in sync (SOK).

11 Administration

The chapter "Administration" refers to our two-node lab cluster *suse12*. The two nodes are named *suse01* and *suse02*. SAP HANA has the SID *HA1* and instance number *10*.

11.1 Dos and don’ts

In your project, you should:

- Define fencing before adding other resources to the cluster.
- Tune the timeouts of operations of SAPHanaController, SAPHanaTopology and SAPHanaFilesystem.
- Use SUSE Trento to check your cluster.
- Do intensive testing.
- Always check HANA and cluster status before doing anything and after having done something, see manual page SAPHanaSR-showStatus(8).
- Always wait for pending cluster actions to finish before doing something, see manual page cs_wait_for_idle(8).
- Set up a test cluster for testing configuration changes and administrative procedure before applying them on the production cluster.

In your project, avoid:

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

- Using site names other than the ones already known by the cluster when manually re-registering a site.
- Adding location rules for the clone, multi-state or IP resource. Only location rules mentioned in this setup guide are allowed. For public cloud refer to the cloud specific documentation.
- Using SAP tools for attempting start/stop/takeover actions on a database while the cluster is in charge of managing that database. Same for unregistering/disabling system replication.

Important

As "migrating" or "moving" resources in crm-shell, HAWK or other tools would add client-prefer location rules, support is limited to maintenance procedures described in this document. See *Section 10, "Testing the cluster"* and *Section 11.3, "Maintenance"* (page 63) for proven procedures.

11.2 Monitoring and tools

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

11.2.1 HAWK – cluster status and more

You can use a Web browser to check the cluster status.

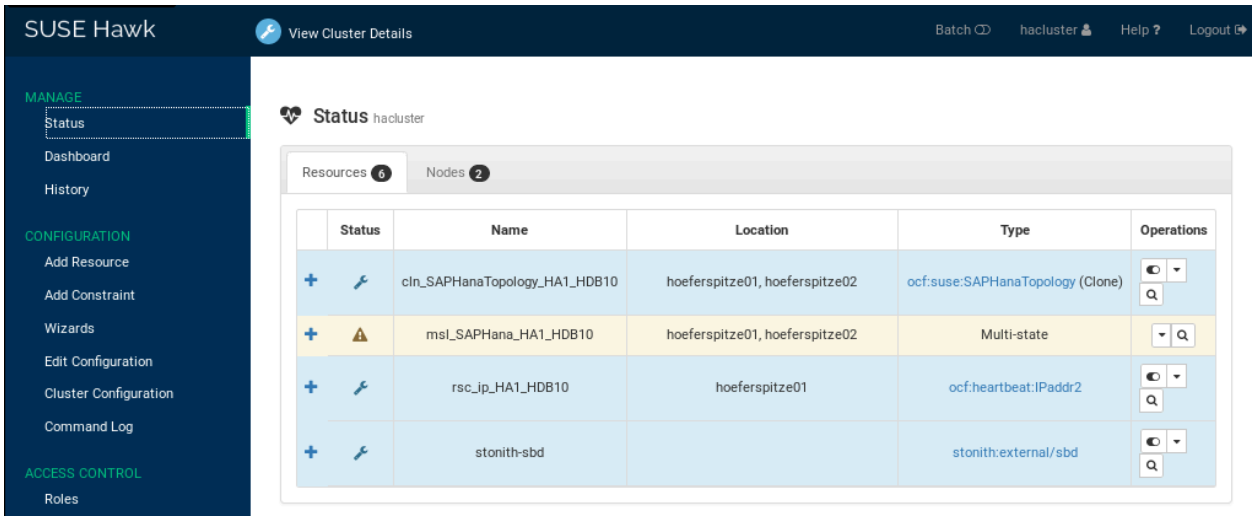

FIGURE 3: CLUSTER STATUS IN HAWK

If you set up the cluster using crm 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, perform administrative tasks or configure resources and cluster bootstrap parameters. Read the product manuals for a complete documentation of this user interface. For the SAP HANA system replication scale-up performance optimized scenario the use of HAWK should follow the guidance given in this guide.

11.2.2 SAP HANA Cockpit

Database-specific administration and checks can be done with SAP HANA Cockpit. Before trying start/stop/takeover for the database, make sure the cluster is not in charge of managing the respective resource. See also *Section 11.3, "Maintenance"* (page 63).

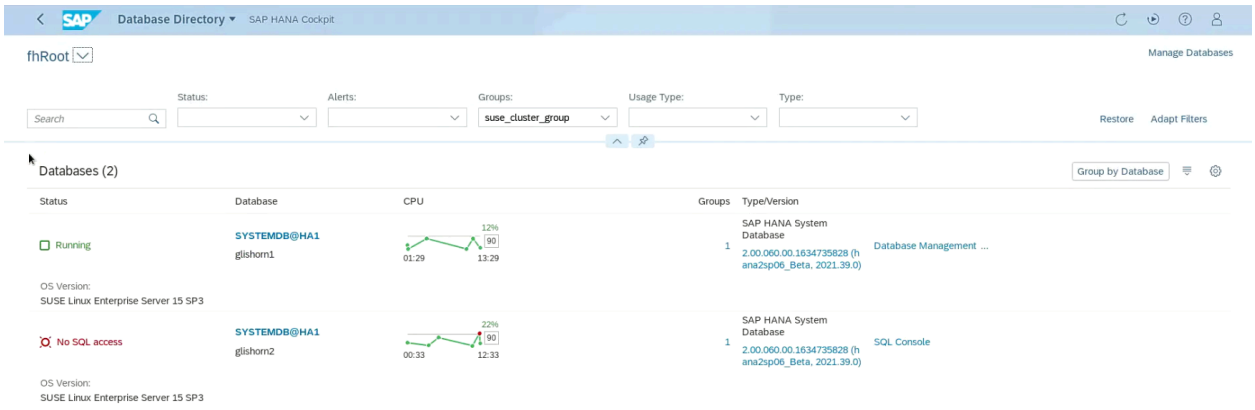

FIGURE 4: SAP HANA COCKPIT – DATABASE DIRECTORY

11.2.3 Cluster command line tools

A simple overview can be obtained by calling `crm_mon`. Using option `-r` shows also stopped but already configured resources. Option `-1` tells `crm_mon` to output the status once instead of periodically.

```
besler3:~ # crm_mon -1r
Cluster Summary:
  * Stack: corosync (Pacemaker is running)
  * Current DC: besler4 (version
 3.0.2+20260601.c75e25851c-160099.1.5-3.0.2+20260601.c75e25851c) - partition with quorum
  * Last updated: Mon Aug  3 16:49:12 2026 on besler3
  * Last change:  Mon Aug  3 16:48:56 2026 by root via crm_attribute on besler3
```

```
  * 2 nodes configured
  * 8 resource instances configured

Node List:
  * Online: [ besler3 besler4 ]

Full List of Resources:
  * fencing-sbd (stonith:fence_sbd):     Started besler3
  * Clone Set: cln_SAPHanaTop_HA1_HDB00 [rsc_SAPHanaTop_HA1_HDB00]:
    * Started: [ besler3 besler4 ]
  * Clone Set: cln_SAPHanaFil_HA1_HDB00 [rsc_SAPHanaFil_HA1_HDB00]:
    * Started: [ besler3 besler4 ]
  * Clone Set: mst_SAPHanaCon_HA1_HDB00 [rsc_SAPHanaCon_HA1_HDB00] (promotable):
    * Promoted: [ besler3 ]
    * Unpromoted: [ besler4 ]
  * rsc_ip_HA1_HDB00    (ocf:heartbeat:IPaddr2):         Started besler3
```

See the manual page crm_mon(8) for details.

11.2.4 SAPHanaSR command line tools

To show SAPHanaController and SAPHanaTopology resource agent internal values, you can call the program `SAPHanaSR-showAttr`. The internal values, the attribute location and their parameter names may change. The command `SAPHanaSR-showAttr` will always fetch the values from the correct attribute location. Do not use cluster commands like `crm_attribute` to fetch the values directly from the cluster. `SAPHanaSR-showAttr` should not be used for automated system monitoring.

```
besler3:~ # SAPHanaSR-showAttr
Global cib-update dcid prim sec sid topology
--------------------------------------------
global 0.165.0    2    S1   S2  HA1 ScaleUp

Resource                 maintenance promotable
-----------------------------------------------
mst_SAPHanaCon_HA1_HDB00 false       true
cln_SAPHanaTop_HA1_HDB00

Site lpt        lss mns     opMode    srHook srMode  srPoll srr
-------------------------------------------------------------
S2   30         4   besler4 logreplay SOK    syncmem SOK    S
S1   1785768660 4   besler3 logreplay PRIM   syncmem PRIM   P

Host    clone_state roles                        score site srah version     vhost
-----------------------------------------------------------------------------------
```

```
besler3 PROMOTED    master1:master:worker:master 150   S1   -    2.00.089.00 besler3
besler4 DEMOTED     master1:master:worker:master 100   S2   -    2.00.089.00 besler4
```

`SAPHanaSR-showAttr` also supports other output formats such as **script**. The script format is intended to allow running filters. Based on `SAPHanaSR-showAttr` output format script you can define effective queries:

```
# SAPHanaSR-showAttr --format=script --experimental_attributes host:vhost
0 Host/suse01/vhost="suse01"
0 Host/suse02/vhost="suse02"
```

See the manual page SAPHanaSR-showAttr(8) and SAPHanaSR-showStatus(8) for details.

`SAPHanaSR-replay-archive` can help to analyze the SAPHanaSR attribute values from `crm_report` archives. This allows post mortem analysis.

In our example, the administrator killed the primary SAP HANA instance using the command `HDB kill-9`. This happened around 9:10 pm.

```
suse01:~ # crm_report -f 17:00
suse01: INFO: The report is saved in ./crm_report-Mon-03-Aug-2026.tar.gz
suse01: INFO: Report timespan: 2026-08-03 17:00:00 - 2026-08-03 17:03:29
suse01: INFO: Including nodes: suse01 suse02
suse01: INFO: Thank you for taking time to create this report

suse01:~ # SAPHanaSR-replay-archive --format=script --experimental_attributes host:score
 site:lss --dir=crm_report-Mon-03-Aug-2026 --sid=HA1

2026-08-03T15:01:20 Site/S2/srHook="SOK"
2026-08-03T15:01:20 Site/S2/lss="4"
2026-08-03T15:01:20 Host/suse02/score="100"
2026-08-03T15:02:04 Site/S2/srHook="SFAIL"
2026-08-03T15:02:04 Site/S2/lss="4"
2026-08-03T15:02:04 Host/suse02/score="100"
2026-08-03T15:02:11 Site/S2/srHook="SFAIL"
2026-08-03T15:02:11 Site/S2/lss="4"
2026-08-03T15:02:11 Host/suse02/score="-INFINITY"
...
2026-08-03T15:02:14 Site/S2/srHook="SFAIL"
2026-08-03T15:02:14 Site/S2/lss="1"
2026-08-03T15:02:14 Host/suse02/score="-1"
2026-08-03T15:03:09 Site/S2/srHook="SOK"
2026-08-03T15:03:09 Site/S2/lss="4"
2026-08-03T15:03:09 Host/suse02/score="0"
2026-08-03T15:03:11 Site/S2/srHook="SOK"
2026-08-03T15:03:11 Site/S2/lss="4"
2026-08-03T15:03:11 Host/suse02/score="100"
```

In the above example the attributes indicate that at the beginning suse02 was running secondary (lss "4", score "100").

At 15:02:04 (UTC) suddenly the system replication is maked as srHook="SFAIL".

At 15:02:11 (UTC) The secondary score is degraded: score="-INFINITY".

At 15:02:14 (UTC) The secondary is also marked as in error (lss="1"). The cluster initiated a restart for the failed secondary.

At 15:03:09 (UTC) The seconary is marked as back (lss="4") and the system replication is up again (srHook="SOK").

At 15:03:10 (UTC) The seconary score is set back to "100".

11.2.5 SAP HANA LandscapeHostConfiguration

To check the status of an HANA database and to find out if the cluster should react, you can use the script `landscapeHostConfiguration.py` to be called as Linux user *<sid>adm*.

```
suse01:~> cdpy; python3 landscapeHostConfiguration.py
| Host    | Host   | ... NameServer   | NameServer  | IndexServer | IndexServer |
|         | Active | ... Config Role  | Actual Role | Config Role | Actual Role |
| ------  | ------ | ... ------------ | ----------- | ----------- | ----------- |
| suse01 | yes    | ... master 1     | master      | worker      | master      |

overall host status: ok
```

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

TABLE 5: INTERPRETATION OF RETURN CODES

Return Code	Description
4	SAP HANA database is up and OK. The cluster interprets this as a correctly running database.
3	SAP HANA database is up and in status info. The cluster interprets this as a correctly running database.
2	SAP HANA database is up and in status warning. The cluster interprets this as a correctly running database.

Return Code	Description
1	SAP HANA database is down. If the database should be up and is not down by intention, this could trigger a takeover.
0	Internal Script Error - Database status unknown. Cluster action depends on RA parameter ON_FAIL_ACTION.

11.3 Maintenance

To receive updates for the operating system or SUSE Linux Enterprise High Availability, it is recommended to register your systems to either a local SUSE Mulit-Linux Manager, to Repository Mirroring Tool (RMT), or remotely with SUSE Customer Center. For more information, visit the respective Web pages: https://www.suse.com/products/multi-linux-manager/ TBD https://scc.suse.com/docs/help Examples for maintenance tasks are also given in manual page SAPHanaSR_maintenance_examples(7).

11.3.1 Updating the operating system and cluster

For an update of SUSE Linux Enterprise Server for SAP applications packages including cluster software, follow the rolling update procedure defined in the product documentation of the SUSE Linux Enterprise High Availability Administration Guide.

11.3.2 Updating SAP HANA - seamless SAP HANA maintenance

For updating SAP HANA database systems in system replication, you need to follow the defined SAP processes. This section describes the steps required before and after the update procedure to get the system replication automated again.

SUSE has optimized the SAP HANA maintenance process in the cluster. The improved procedure only sets the multi-state resource to maintenance and keeps the rest of the cluster (SAPHanaTopology clones and IPaddr2 vIP resource) still active. Using the updated procedure allows a seamless SAP HANA maintenance in the cluster, as the virtual IP address can automatically follow the running primary.

Prepare the cluster not to react on the maintenance work to be done on the SAP HANA database systems. Set the multi-state resource to `maintenance`.

EXAMPLE 16: MAIN SAP HANA UPDATE PROCEDURE

Pre-Update Tasks

For the multi-state-resource set the maintenance mode as follows:

```
# crm resource maintenance <multi-state-resource>
```

The <multi-state-resource> in the guide at hand is `mst_SAPHanaCon_HA1_HDB10`.

Update

Process the SAP Update for both SAP HANA database systems. This procedure is described by SAP.

During this SAP HANA update, the active primary the moves from one site to an other. After each move of primary and register the new secondary the cluster should to be informed to score the new SAP HANA sites.

```
# crm resource refresh <multi-state-resource>
```

Using this resource refresh will set the new scores and allows the seamless maintenance. The IP adress for the primary will follow the score 150 (running primary), while an optional IP address for the secondary will follow the score 100 (running and in-sync secondary).

Post-Update Tasks

Expect the primary/secondary roles to be exchanged after the maintenance. Therefore, tell the cluster to forget about these states and to reprobe the updated SAP HANA database systems.

```
# crm resource refresh <multi-state-resource>
```

After the SAP HANA update is complete on both sites, tell the cluster about the end of the maintenance process. This allows the cluster to actively control and monitor the SAP again.

```
# crm resource maintenance <multi-state-resource> off
```

Optionally, you could completety remove the maintenance attribute from the resource. For more details see manual page SAPHanaSR_maintenance_examples(7) and crm(8).

11.3.3 Migrating an SAP HANA primary using handshake takeover

In the following procedures, we assume the primary runs on node 1 and the secondary on node 2. The goal is to "exchange" the roles of the nodes: the primary should then run on node 2 and the secondary should run on node 1.

There are different methods to get the exchange of the roles done. The following procedure shows how to tell the cluster to "accept" a role change via native SAP HANA commands.

EXAMPLE 17: MIGRATING AN SAP HANA PRIMARY USING SAP TOOLSET - HANDSHAKE TAKEOVER

Pre-Migration Tasks

Set the multi-state resource to `maintenance`. This can be done on any cluster node.

```
# crm resource maintenance <multi-state-resource>
```

Manual Takeover Process

- On the secondary SAP HANA database system trigger the takeover with suspend of the primary

  ```
  # su - ha1adm -c "hdbnsutil -sr_takeover --suspendPrimary"
  ```

  The former primary keeps started, but suspended. The former secondary will get the new primary role.

- Wait until the secondary has completely taken over to be the new primary role. You see this using the command line tool `SAPHanaSR-showAttr`. Check for the site attributes "srr" ("P") and "lss" ("4") for the new primary site.

  ```
  suse01:~ # SAPHanaSR-showAttr --format script --experimental_attributes
   site:lss,srr,mns
  0 Site/S2/lss="4"
  0 Site/S2/srr="P"
  0 Site/S2/mns="suse02"
  0 Site/S1/lss="4"
  0 Site/S1/srr="P"
  0 Site/S1/mns="suse01"
  ```

- On the former primary run the online register command

  ```
  ~> hdbnsutil -sr_register --remoteHost=suse02 --remoteInstance=10 \
   --replicationMode=sync --name=WDF \
   --operationMode=logreplay --online
  ```

- Wait until the new secondary has been registered and started. You see this using the command line tool `SAPHanaSR-showAttr` and check for the site attributes "lss" ("4") and "srr" ("S") for the new secondary.

```
suse01:~ # SAPHanaSR-showAttr --format script --experimental_attributes
 site:lss,srr,mns
0 Site/S2/lss="4"
0 Site/S2/srr="P"
0 Site/S2/mns="suse02"
0 Site/S1/lss="4"
0 Site/S1/srr="S"
0 Site/S1/mns="suse01"
```

Post-Migration Tasks

- Tell the cluster to forget about the former multi-state roles and to re-monitor the failed master. The command can be submitted on any cluster node as user root.

```
# crm resource refresh <multi-state-resource>
```

- Set the multi-state resource to the status managed again. The command can be submitted on any cluster node as user root.

```
# crm resource maintenance <multi-state-resource> off
```

11.3.4 Migrating an SAP HANA primary using standard takeover

EXAMPLE 18: MIGRATING AN SAP HANA PRIMARY USING SAP TOOLSET (STANDARD TAKEOVER)

Pre-Migration Tasks

Set the multi-state resource to `maintenance`. This can be done on any cluster node.

```
# crm resource maintenance <multi-state-resource>
```

Manual Takeover Process

- Stop the primary SAP HANA database system. Enter the command in our example on node1 as user *<sid>adm*.

```
~> HDB stop
```

- Before proceeding, make sure the primary SAP HANA database is stopped.
- Start the takeover process on the secondary SAP HANA database system. Enter the command in our example on node 2 as user *<sid>adm*.

```
~> hdbnsutil -sr_takeover
```

- Register the former primary to become the new secondary. Enter the command in our example on node1 as user <*sid*>*adm*.

```
~> hdbnsutil -sr_register --remoteHost=suse02 --remoteInstance=10 \
 --replicationMode=sync --name=WDF \
 --operationMode=logreplay
```

- Start the new secondary SAP HANA database system. Enter the command in our example on node1 as user <*sid*>*adm*.

```
~> HDB start
```

Post-Migration Tasks

- Wait some time until `SAPHanaSR-showAttr` shows both SAP HANA database systems to be up again (field roles must start with the digit 4). The new secondary should have role "S" (for secondary).
- Tell the cluster to forget about the former multi-state roles and to re-monitor the failed master. The command can be submitted on any cluster node as user root.

```
# crm resource refresh <multi-state-resource>
```

- Set the multi-state resource to the status managed again. The command can be submitted on any cluster node as user root.

```
# crm resource maintenance <multi-state-resource> off
```

The following paragraphs explain how to use the cluster to partially automate the migration.

EXAMPLE 19: MOVING AN SAP HANA PRIMARY USING THE CLUSTER TOOLSET

- Create a "move away" from this node rule by using the **force** option.

```
# crm resource move <multi-state-resource> force
```

Because of the "move away" (**force**) rule, the cluster will **stop** the current primary. After that, run a **promote** on the secondary site if the system replication was in sync before. You should not migrate the primary if the status of the system replication is not in sync (SFAIL).

Important

The **force** option of the move comamnd shall be used. Only the migration with **force** option is supported. Migration without the **force** option will cause a takeover without the former primary to be stopped. This would result in a dual active primary.

- Wait until the secondary has completely taken over to be the new primary role. You see this using the command line tool `SAPHanaSR-showAttr`. Check for the site attributes "srr" ("P") and "lss" ("4") for the new primary site.

```
suse01:~ # SAPHanaSR-showAttr --format script --experimental_attributes
 site:lss,srr,mns
0 Site/S2/lss="4"
0 Site/S2/srr="P"
0 Site/S2/mns="suse02"
0 Site/S1/lss="1"
0 Site/S1/srr="P"
0 Site/S1/mns="suse01"
```

- If you have set up the parameter value `AUTOMATED_REGISTER="true"`, you can skip this step. In other cases you now need to register the old primary. Enter the command in our example on node1 as user <*sid*>*adm*.

```
~> hdbnsutil -sr_register --remoteHost=suse02 --remoteInstance=10 \
    --replicationMode=sync --operationMode=logreplay \
    --name=WDF
```

- Clear the ban rules of the resource to allow the cluster to start the new secondary.

```
# crm resource clear <multi-state-resource>
```

- Wait until the new secondary has started. You see this using the command line tool `SAPHanaSR-showAttr` and check for the site attributes "lss" ("4") and "srr" ("S") for the new secondary.

```
suse01:~ # SAPHanaSR-showAttr --format script --experimental_attributes
 site:lss,srr,mns
0 Site/S2/lss="4"
0 Site/S2/srr="P"
0 Site/S2/mns="suse02"
0 Site/S1/lss="4"
0 Site/S1/srr="S"
```

```
0 Site/S1/mns="suse01"
```

12 Examples

The chapter "Examples" refers to our two-node lab cluster *besler34*. The two nodes are named *besler3* and *besler4*. SAP HANA has the SID *HA1* and instance number *00*.

12.1 Example `crm cluster init` and join

1. Example for `crm cluster init` dialog

```
INFO: Loading "default" profile from /etc/crm/profiles.yml
INFO: Loading "knet-default" profile from /etc/crm/profiles.yml
INFO: A public key is added to authorized_keys for user root: 2048
 SHA256:BpU1NYKSI1pLWQVkZCotf1jdiThQlMrCd77Pz2XTGf0 root@glishorn01 (RSA)
INFO: The user 'hacluster' will have the login shell configuration changed to /bin/bash
Continue (y/n)? y
INFO: A new ssh keypair is generated for user hacluster.
INFO: A public key is added to authorized_keys for user hacluster: 3072
 SHA256:sOI39ogMRs3WOJqFzFO12mTy+gAf9jS7Bu7Y334W5g4 Cluster internal on besler3 (RSA)
INFO: Added firewalld service high-availability on besler3
Address for ring0 [192.168.2.62]
INFO: Configuring disk-based SBD
INFO: Initializing SBD device /dev/disk/by-id/scsi-0QEMU_QEMU_HARDDISK_drive-scsi0-0-0
INFO: Update SBD_DEVICE in /etc/sysconfig/sbd: /dev/disk/by-id/
scsi-0QEMU_QEMU_HARDDISK_drive-scsi0-0-0
INFO: Update SBD_WATCHDOG_DEV in /etc/sysconfig/sbd: /dev/watchdog0
INFO: Enable sbd.service on node besler3
INFO: Configure Qdevice/Qnetd:
  QDevice participates in quorum decisions. With the assistance of
  a third-party arbitrator Qnetd, it provides votes so that a cluster
  is able to sustain more node failures than standard quorum rules
  allow. It is recommended for clusters with an even number of nodes
  and highly recommended for 2 node clusters.

Do you want to configure QDevice (y/n)? n
INFO: Hawk cluster interface is now running. To see cluster status, open:
INFO:   https://besler:37630/
INFO: Log in with username 'hacluster', password 'linux'
WARNING: You should change the hacluster password to something more secure!
INFO: Starting and enable pacemaker.service on besler3
INFO: BEGIN Waiting for cluster
............
```

```
INFO: END Waiting for cluster
INFO: Loading initial cluster configuration
INFO: Configuring fence agent stonith:fence_sbd
INFO: Update SBD_DELAY_START in /etc/sysconfig/sbd: 41
INFO: Adjusting stonith-timeout to 71
WARNING: "stonith-timeout" in crm_config is set to 71, it was 60s
INFO: Setting stonith-enabled to true
WARNING: "stonith-enabled" in crm_config is set to true, it was false
INFO: Createing runtime drop-in file /run/systemd/system/sbd.service.d/
sbd_delay_start_disabled.conf to unset SBD_DELAY_START
WARNING: "stonith-timeout" is deprecated, please consider using "fencing-timeout"
WARNING: "stonith-enabled" is deprecated, please consider using "fencing-enabled"
INFO: Configure Administration IP Address:
  Optionally configure an administration virtual IP
  address. The purpose of this IP address is to
  provide a single IP that can be used to interact
  with the cluster, rather than using the IP address
  of any specific cluster node.

Do you wish to configure a virtual IP address (y/n)? n
INFO: Done (log saved to /var/log/crmsh/crmsh.log on besler3)
```

1. Example for `crm cluster join` dialog

```
besler4:~ # crm cluster join -c besler3
INFO: The user 'hacluster' will have the login shell configuration changed to /bin/bash
Continue (y/n)? y
INFO: A new ssh keypair is generated for user hacluster.
INFO: A public key is added to authorized_keys for user hacluster: 3072
 SHA256:XU8AHmSS8Zr1UT5FSBuoOoDP7ljFFHdyjHZY7w55kXs Cluster internal on besler4 (RSA)
INFO: Added firewalld service high-availability on besler4
INFO: Merging known_hosts
INFO: BEGIN Probing for new partitions
INFO: END Probing for new partitions
INFO: BEGIN Retrieving all configuration files from besler3
INFO: END Retrieving all configuration files from besler3
Address for ring0 [192.168.2.64]

INFO: Sync file /etc/corosync/corosync.conf to besler3
INFO: Got SBD configuration
INFO: Enable sbd.service on node besler4
INFO: Hawk cluster interface is now running. To see cluster status, open:
INFO:   https://192.168.2.64:7630/
INFO: Log in with username 'hacluster', password 'linux'
WARNING: You should change the hacluster password to something more secure!
INFO: Starting and enable pacemaker.service on besler4
INFO: BEGIN Waiting for cluster
```

```
..
INFO: END Waiting for cluster
INFO: Setting pcmk_delay_max parameter to 30s in resource fencing-sbd
INFO: Update SBD_DELAY_START in /etc/sysconfig/sbd: 71
INFO: Sync file /etc/sysconfig/sbd to besler3
WARNING: "stonith-timeout" is deprecated, please consider using "fencing-timeout"
WARNING: "stonith-enabled" is deprecated, please consider using "fencing-enabled"
WARNING: "priority" in rsc_defaults is set to 1, it was 0
WARNING: "priority-fencing-delay" in crm_config is set to 60, it was 0
INFO: Sync file /etc/corosync/corosync.conf to besler3
INFO: BEGIN Reloading cluster configuration
INFO: END Reloading cluster configuration
INFO: Done (log saved to /var/log/crmsh/crmsh.log on besler4)
```

12.2 Example for corosync configuration

1. Example */etc/corosync/corosync.conf* from lab cluster *besler34*

```
# Generated by crmsh
# For more details please see corosync.conf.5 man page
totem {
        version: 2
        cluster_name: besler34
        transport: knet
        token: 5000
        join: 60
        max_messages: 20
        token_retransmits_before_loss_const: 10
        crypto_hash: sha256
        crypto_cipher: aes256
}

quorum {
        provider: corosync_votequorum
        two_node: 1
}

logging {
        to_logfile: yes
        logfile: /var/log/cluster/corosync.log
        to_syslog: yes
        timestamp: hires
}

nodelist {
```

```
        node {
                ring0_addr: 192.168.2.62
                name: besler3
                nodeid: 1
        }

        node {
                ring0_addr: 192.168.2.64
                name: besler4
                nodeid: 2
        }

}
```

12.3 Example cluster configuration

The following complete progressive crm configuration is for a two-node cluster (besler3, besler4) and an SAP HANA database with SID HA1 and instance number 00. A failed primary will be automatically registered as secondary after reboot. Priority fencing prefers the SAP HANA primary in case of split-brain. The virtual IP address in the example is 192.168.2.66.

1. Example crm configuration from lab cluster *besler34*

```
node 1: besler3 \
        attributes hana_ha1_vhost=besler3 hana_ha1_site=S1
node 2: besler4 \
        attributes hana_ha1_vhost=besler4 hana_ha1_site=S2
#
primitive fencing-sbd stonith:fence_sbd \
        params pcmk_delay_max=15
#
primitive rsc_SAPHanaCon_HA1_HDB00 ocf:suse:SAPHanaController \
        op start interval=0 timeout=3600 \
        op stop interval=0 timeout=600 \
        op promote interval=0 timeout=900 \
        op demote interval=0 timeout=320 \
        op monitor interval=60 role=Promoted timeout=700 \
        op monitor interval=61 role=Unpromoted timeout=700 \
        params SID=HA1 InstanceNumber=00 PREFER_SITE_TAKEOVER=true \
        DUPLICATE_PRIMARY_TIMEOUT=7200 AUTOMATED_REGISTER=false \
        ON_FAIL_ACTION=fence \
        meta priority=100
#
primitive rsc_SAPHanaFil_HA1_HDB00 ocf:suse:SAPHanaFilesystem \
```

```
        op start interval=0 timeout=10 \
        op stop interval=0 timeout=20 \
        op monitor interval=120 timeout=120 \
        params SID=HA1 InstanceNumber=00 ON_FAIL_ACTION=fence
#
primitive rsc_SAPHanaTop_HA1_HDB00 ocf:suse:SAPHanaTopology \
        op start interval=0 timeout=600 \
        op stop interval=0 timeout=300 \
        op monitor interval=50 timeout=600 \
        params SID=HA1 InstanceNumber=00
#
primitive rsc_ip_HA1_HDB00 IPaddr2 \
        op monitor interval=10 timeout=20 \
        params ip=192.168.2.66
#
clone cln_SAPHanaFil_HA1_HDB00 rsc_SAPHanaFil_HA1_HDB00 \
        meta clone-node-max=1 interleave=true
#
clone cln_SAPHanaTop_HA1_HDB00 rsc_SAPHanaTop_HA1_HDB00 \
        meta clone-node-max=1 interleave=true
#
clone mst_SAPHanaCon_HA1_HDB00 rsc_SAPHanaCon_HA1_HDB00 \
        meta clone-node-max=1 promotable=true interleave=true
#
location loc_ip_with_HA1_HDB00 rsc_ip_HA1_HDB00 \
        rule 2000: master-rsc_SAPHanaCon_HA1_HDB00 eq 150
#
order ord_saphana_HA1_HDB00 Optional: cln_SAPHanaTop_HA1_HDB00 \
        mst_SAPHanaCon_HA1_HDB00
#
property SAPHanaSR: \
        hana_ha1_site_srHook_S2=SOK \
        hana_ha1_glob_topology=ScaleUp \
        hana_ha1_site_lss_S2=4 \
        hana_ha1_site_lss_S1=4 \
        hana_ha1_site_srr_S2=S \
        hana_ha1_site_srr_S1=P \
        hana_ha1_site_srMode_S2=syncmem \
        hana_ha1_site_srMode_S1=syncmem \
        hana_ha1_site_mns_S1=besler3 \
        hana_ha1_site_mns_S2=besler4 \
        hana_ha1_site_lpt_S2=30 \
        hana_ha1_glob_sec=S2 \
        hana_ha1_site_srHook_S1=PRIM \
        hana_ha1_site_lpt_S1=1785762330 \
        hana_ha1_site_opMode_S2=logreplay \
        hana_ha1_site_opMode_S1=logreplay \
```

```
        hana_ha1_glob_prim=S1 \
        hana_ha1_site_srPoll_S1=PRIM \
        hana_ha1_site_srPoll_S2=SOK
property cib-bootstrap-options: \
        dc-version="3.0.2+20260601.c75e25851c-160099.1.5-3.0.2+20260601.c75e25851c" \
        cluster-infrastructure=corosync \
        have-watchdog=true \
        cluster-name=hacluster \
        priority-fencing-delay=30 \
        fencing-enabled=true \
        fencing-action=reboot \
        fencing-timeout=150 \
        last-lrm-refresh=1785761575
rsc_defaults build-resource-defaults: \
        resource-stickiness=1000 \
        migration-threshold=5000 \
        priority=1
op_defaults op-options: \
        timeout=600 \
        record-pending=true
```

13 References

For more detailed information, have a look at the documents listed below.

13.1 SUSE Product Documentation

Best practices for SAP on SUSE Linux Enterprise

https://documentation.suse.com/sbp/sap-16/

Supported high availability solutions by SUSE Linux Enterprise Server for SAP applications overview

https://documentation.suse.com/sles-sap/sap-ha-support/

SUSE product manuals and documentation

https://documentation.suse.com/

Release notes

https://documentation.suse.com/release-notes/16.0/

Online documentation of SLES for SAP

https://documentation.suse.com/sles-sap/16.0/

Online documentation of SUSE Linux Enterprise High Availability

https://documentation.suse.com/sle-ha/16.0/

Deployment guide for SUSE Linux Enterprise Server

https://documentation.suse.com/sles/16.0/html/SLES-deployment-pxe-client/

Tuning guide for SUSE Linux Enterprise Server

https://documentation.suse.com/sles-sap/16.0/html/SAP-saptune/

XFS file system

https://www.suse.com/c/xfs-the-file-system-of-choice/

SUSE YES certified hardware database

https://www.suse.com/yessearch/

SUSE Mulit-Linux Manager Product Page

https://www.suse.com/products/multi-linux-manager/

SUSE Mulit-Linux Manager Documentation

https://documentation.suse.com/multi-linux-manager/5.1/

RMT = Repository Mirroring Tool documentation

TBD

SUSE Customer Center Frequently Asked Questions

https://scc.suse.com/docs/help

13.2 Related Manual Pages

- chronyc(8)
- corosync.conf(8)
- corosync_overview(8)
- crm(8)
- crm_mon(8)
- crm_simulate(8)
- cs_clusterstate(8)
- cs_man2pdf(8)
- cs_show_hana_info(8)
- cs_show_sbd_devices(8)
- cs_wait_for_idle(8)
- ha_related_sap_notes(7)
- ha_related_suse_tids(7)
- ocf_heartbeat_IPaddr2(7)
- ocf_heartbeat_SAPInstance(7)
- ocf_suse_SAPHanaController(7)
- ocf_suse_SAPHanaFilesystem(7)
- ocf_suse_SAPHanaTopology(7)
- SAPHanaSR(7)
- SAPHanaSR-alert-fencing(8)

- SAPHanaSR-angi(7)
- SAPHanaSR-angi-scenarios(7)
- SAPHanaSR_basic_cluster(7)
- SAPHanaSR-CostOpt_setup(7)
- SAPHanaSR-hookHelper(8)
- SAPHanaSR_maintenance_examples(7)
- SAPHanaSR-manageAttr(8)
- SAPHanaSR-manageProvider(8)
- SAPHanaSR-replay-archive(8)
- SAPHanaSR-ScaleOut(7)
- SAPHanaSR-ScaleOut_basic_cluster(7)
- SAPHanaSR-ScaleOut-XSA(7)
- SAPHanaSR-showAttr(8)
- SAPHanaSR-showStatus(8)
- SAPHanaSR_with_DR_HA(7)
- SAPHanaSR_with_DR_HA_setup(7)
- saptune(8)
- sbd(8)
- stonith_sbd(7)
- sudo(8)
- sudoers(5)
- supportconfig(8)
- susChkSrv.py(7)
- susCostOpt.py(7)
- susHanaSR.py(7)
- susTkOver.py(7)

- systemctl(8)
- systemd-cgls(8)
- votequorum(5)
- zypper(8)

13.3 Related SUSE TIDs

Estimate correct multipath timeout

https://www.suse.com/support/kb/doc/?id=7016305

Can’t open watchdog device: /dev/watchdog: Device or resource busy

https://www.suse.com/support/kb/doc/?id=7008216

Systemd-udev-settle timing out

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

How to load the correct watchdog kernel module

https://www.suse.com/support/kb/doc/?id=7016880

TID XFS metadata corruption and invalid checksum on SAP Hana servers

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

Overcommit Memory in SLES

https://www.suse.com/support/kb/doc/?id=7002775

Recommended SUSE SLES 4 SAP Settings

https://www.suse.com/support/kb/doc/?id=7024082

SAPHanaController running in timeout when starting SAP Hana

https://www.suse.com/support/kb/doc/?id=000019899

Entry "CALLING CRM: ... rc=256" in HANA trace after upgrading SAPHanaSR-ScaleOut

https://www.suse.com/support/kb/doc/?id=000020599

SAP HANA monitors timed out after 5 seconds

https://www.suse.com/support/kb/doc/?id=000020626

HA for SAP HANA: Resource agents are running into HANA_CALL timeout lssRc=124

https://www.suse.com/support/kb/doc/?id=000022345

HA cluster takeover takes too long on HANA indexserver failure

https://www.suse.com/support/kb/doc/?id=000020845

Cluster node fence as SAPHanaTopology fails with error code 1 (OCF_ERR_GENERIC) during a normal cluster stop

https://www.suse.com/support/kb/doc/?id=000020964

SUSE HA for HANA cluster node fenced at shutdown, despite of systemd integration

https://www.suse.com/support/kb/doc/?id=000021046

SAP HANA scale-out - pacemaker.service: "Cannot find sapstartsrv and sapcontrol executable, please set DIR_EXECUTABLE parameter!"

https://www.suse.com/support/kb/doc/?id=000021062

SAPHanaSR-showAttr fails with error "Error: NIECONN_REFUSED ..."

https://www.suse.com/support/kb/doc/?id=000020548

Protect HANA against manually caused dual-primary situation in SUSE HA cluster

https://www.suse.com/support/kb/doc/?id=000021044

Address space monitoring and HANA DB performance

https://www.suse.com/support/kb/doc/?id=000020746

HANA DB resource failed to start

https://www.suse.com/support/kb/doc/?id=000020948

SAPHanaController monitor timeout leads to database restart

https://www.suse.com/support/kb/doc/?id=000021249

HANA Database Planning Engine crashes in __strncmp_avx2_rtm+0x1b3

https://www.suse.com/support/kb/doc/?id=000021026

Basic health check for two-node SAP HANA performance based model

https://www.suse.com/support/kb/doc/?id=7022984

How to re-enable replication in a two-node SAP performance based model

https://www.suse.com/support/kb/doc/?id=7023127

Showing SOK Status in Cluster Monitoring Tools Workaround

https://www.suse.com/support/kb/doc/?id=7023526

HANA SystemReplication doesn’t provide SiteName to Corosync Cluster

https://www.suse.com/support/kb/doc/?id=000019754

SUSE Cluster Support for SAP HANA System Replication Active / Active Read Enabled Feature

https://www.suse.com/support/kb/doc/?id=7023884

SAP Generating 'Database host operating system is not supported' alerts

https://www.suse.com/support/kb/doc/?id=7023744

sapstartsrv does not respawn after a forceful kill of the master nameserver

https://www.suse.com/support/kb/doc/?id=7024291

SAPHanaSR HANA system replication automation without layer 2 network

https://www.suse.com/support/kb/doc/?id=000020333

The vIP cluster resource does not follow the SAP HANA master ...

https://www.suse.com/support/kb/doc/?id=000019769

SAP Instances failed stop on shutdown (PACEMAKER, SYSTEMD, SAP)

https://www.suse.com/support/kb/doc/?id=7022671

SAP on SLES shows Error: NIECONN_REFUSED in the logs

https://www.suse.com/support/kb/doc/?id=7023236

Indepth HANA Cluster Debug Data Collection (PACEMAKER, SAP)

https://www.suse.com/support/kb/doc/?id=7022702

How to prevent certain values in limits.conf from being changed by saptune

https://www.suse.com/support/kb/doc/?id=7023104

Disabling fstrim - under which conditions?

https://www.suse.com/support/kb/doc/?id=7023805

saptune: WARNING saptune.io.go:66: 'noop' is not a valid scheduler for device

https://www.suse.com/support/kb/doc/?id=000019572

How to patch a SAP Application Pacemaker Cluster

https://www.suse.com/support/kb/doc/?id=000020268

crm_ipc_prepare: Message exceeds the configured ipc limit (PACEMAKER COROSYNC)

https://www.suse.com/support/kb/doc/?id=7013215

Unexpected SAP HANA DB failover if Pacemaker sub-daemon is unresponsive to IPC

https://www.suse.com/support/kb/doc/?id=000022175

HA for SAP HANA: pacemaker sub-daemon is unresponsive to IPC

https://www.suse.com/support/kb/doc/?id=000022179

SAP HANA swap and system hangs with THP on multi-NUMA node systems

https://www.suse.com/support/kb/doc/?id=000022286

13.4 Related SUSE blogs

What are SAPHanaSR-angi configuration variants?

https://www.suse.com/c/what-are-saphanasr-angi-configuration-variants/

Emergency Braking for SAP HANA Dying Indexserver

https://www.suse.com/c/emergency-braking-for-sap-hana-dying-indexserver/

SAP HANA Cockpit with SUSE HA integration greatly improves data integrity

https://www.suse.com/c/sap-hana-cockpit-with-suse-ha-integration-greatly-improves-data-integrity/

SAPHanaSR-ScaleOut for Multi-Target Architecture and Principles

https://www.suse.com/c/saphanasr-scaleout-multi-target/

SAP HANA Scale-Out System Replication for large ERP Systems

https://www.suse.com/c/sap-hana-scale-out-system-replication-for-large-erp-systems/

SUSE HA for SAP HANA scale-up cost-optimized improved

https://www.suse.com/c/suse-ha-for-sap-hana-scale-up-cost-optimized-improved/

Entry to blog series #towardsZeroDowntime

https://www.suse.com/c/tag/towardszerodowntime/

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

13.5 Related SAP Documentation

SAP Product Availability Matrix

https://support.sap.com/en/release-upgrade-maintenance.html#section_1969201630

SAP HANA Installation and Update Guide

https://help.sap.com/doc/e9702d76c3284623b02de196c0e79e49/2.0.05/en-US/SAP_HANA_Server_Installation_Guide_en.pdf

SAP HANA Administration Guide

https://help.sap.com/doc/eb75509ab0fd1014a2c6ba9b6d252832/2.0.05/en-US/SAP_HANA_Administration_Guide_en.pdf

SAP HANA Documentation Entry Page

https://help.sap.com/viewer/product/SAP_HANA_PLATFORM/2.0.05/en-US

SAP HANA Tailored Data Center Integration - FAQ

https://help.sap.com/docs/SAP_HANA_PLATFORM/eb3777d5495d46c5b2fa773206bbf-b46/0b15a92d554c4941a452c9ca127f8c70.html

SAP HANA and Persistent Memory

https://blogs.sap.com/2020/01/30/sap-hana-and-persistent-memory/

SAP HANA HA/DR Provider Hook Methods

https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.05/en-US/5df2e766549a405e95de4c5d7f2efc2d.html

13.6 Related SAP Notes

611361 - Hostnames of SAP servers

https://me.sap.com/notes/611361

768727 - Automatic restart functions in sapstart for processes

https://me.sap.com/notes/768727

927637 - Web service authentication in sapstartsrv as of Release 7.00

https://me.sap.com/notes/927637

1092448 - IBM XL C/C++ runtime environment for Linux on system p

https://me.sap.com/notes/1092448

1514967 - SAP HANA: Central Note

https://me.sap.com/notes/1514967

1552925 - Linux: High Availability Cluster Solutions

https://me.sap.com/notes/1552925

1763512 - Support details for SUSE Linux Enterprise for SAP applications

https://me.sap.com/notes/1763512

1846872 - "No space left on device" error reported from HANA

https://me.sap.com/notes/1846872

1876398 - Network configuration for System Replication in HANA SP6

https://me.sap.com/notes/1876398

1888072 - SAP HANA DB: Indexserver crash in strcmp sse42

https://me.sap.com/notes/1888072

2021789 - SAP HANA Revison and Maintenance Strategy

https://me.sap.com/notes/2021789

2196941 - SAP HANA Software Replication Takeover Hook Changes

https://me.sap.com/notes/2196941

2235581 - SAP HANA: Supported Operating Systems

https://me.sap.com/notes/2235581

2369981 - Required configuration steps for authentication with HANA System Replication

https://me.sap.com/notes/2369981

2369910 - SAP Software on Linux: General information

https://me.sap.com/notes/2369910

2380229 - SAP HANA Platform 2.0 - Central Note

https://me.sap.com/notes/2380229

2434562 - System Replication Hanging in Status "SYNCING" or "ERROR" With Status Detail "Missing Log" or "Invalid backup size"

https://me.sap.com/notes/2434562

3565382 - SUSE Linux Enterprise Server 16: Installation Note

https://me.sap.com/notes/3565382

2647673 - HANA Installation Failure

https://me.sap.com/notes/2647673

3577842 - SAP HANA DB: Recommended OS settings for SLES 16 / SLES for SAP applications 16

https://me.sap.com/notes/3577842

2733483 - Host Auto-Failover Not Occur when Indexserver Crash on Worker Node

https://me.sap.com/notes/2733483

2750199 - Incorrect Alert Regarding Unsupported Operating System Version

https://me.sap.com/notes/2750199

3007062 - FAQ: SAP HANA & Third Party Cluster Solutions

https://me.sap.com/notes/3007062

3014176 - Applying System Size Dependent Resource Limits During Installation or Upgrade

https://me.sap.com/notes/3014176

3091152 - sapstartsrv - improved deregistration for UNIX/Linux

https://me.sap.com/notes/3091152

3093542 - Transition to Python 3 of the Python Distribution Delivered With SAP HANA 2.0 Server

https://me.sap.com/notes/3093542

3139184 - Linux: systemd integration for sapstartsrv and SAP Hostagent

https://me.sap.com/notes/3139184

3145200 - SAP Host Agent 7.22 PL57

https://me.sap.com/notes/3145200

3366673 - landscapeHostConfiguration.py fails with Python runtime error "can't start new thread"

https://me.sap.com/notes/3366673

3446532 - Pacemaker cluster resource unavailable due to inadequate memory garbage collection parameter settings

https://me.sap.com/notes/3446532

3723107 - landscapeHostConfiguration.py Reports "Unknown" Status During SAP HANA Startup Causing Unintended Cluster Fencing

https://me.sap.com/notes/3723107

13.7 Pacemaker

Pacemaker Project Documentation

https://clusterlabs.org/pacemaker/doc/

14 Legal notice

Copyright © 2006-2026 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.

15 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. Secondarily, 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 https://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.

16 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. Secondarily, 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 https://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.