

SAP S/4 HANA - Enqueue Replication 2 High Availability Cluster

Setup Guide

SUSE Linux Enterprise Server for SAP Applications 15
SAP HANA

Fabian Herschel, Distinguished Architect SAP (SUSE)
Bernd Schubert, SAP Solution Architect (SUSE)

SAP S/4 HANA - Enqueue Replication 2 High Availability Cluster

Setup Guide

Date: 2024-11-14

SUSE® Linux Enterprise Server for SAP Applications is optimized in various ways for SAP* applications. This document explains how to deploy an S/4 HANA Enqueue Replication 2 High Availability Cluster solution. It is based on SUSE Linux Enterprise Server for SAP Applications 15. The concept however can also be used with newer service packs of SUSE Linux Enterprise Server for SAP Applications.

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

Contents

- 1 About this guide 4
- 2 Scope of this document 6
- 3 Overview 6
- 4 SAP installation 16
- 5 Implementing the cluster 24
- 6 Administration 30
- 7 Multi-node cluster setups for an SAP S/4HANA 36
- 8 References 44
- 9 Appendix 47
- 10 Legal notice 52
- 11 GNU Free Documentation License 53

1 About this guide

The following sections focus on background information and the purpose of the document at hand.

1.1 Introduction

SUSE® Linux Enterprise Server for SAP Applications is the optimal platform to run SAP* applications with high availability. Together with a redundant layout of the technical infrastructure, single points of failure can be eliminated.

SAP* Business Suite is a sophisticated application platform for large enterprises and mid-size companies. Many critical business environments require the highest possible SAP* application availability.

The described cluster solution can be used for SAP* SAP S/4HANA ABAP Platform.

SAP S/4HANA ABAP Platform is a common stack of middleware functionality used to support SAP business applications. The SAP Enqueue Replication Server 2 constitutes application level redundancy for one of the most crucial components of the SAP S/4HANA ABAP Platform stack, the enqueue service. An optimal effect of the enqueue replication mechanism can be achieved when combining the application level redundancy with a high availability cluster solution, as provided for example by SUSE Linux Enterprise Server for SAP Applications. Over several years of productive operations, the components mentioned have proven their maturity for customers of different sizes and industries.

The here described HA setup is based on cluster-controlled file systems for the SAP central services working directories. This setup with cluster-controlled file system resources has been used in former HA certifications. It still is supported. For deploying new HA clusters, we recommend using the so-called simple mount setup, described in "SAP S/4 HANA - Enqueue Replication 2 High Availability Cluster With Simple Mount - Setup Guide" (https://documentation.suse.com/sbp/sap/html/SAP_S4HA10_SetupGuide_SimpleMount-SLE15).

1.2 Additional documentation and resources

Several chapters in this document contain links to additional documentation resources that are either available on the system or on the Internet.

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

More whitepapers, guides and best practices documents referring to SUSE Linux Enterprise Server and SAP can be found and downloaded at the SUSE Best Practices Web page: <https://documentation.suse.com/sbp/sap/>

Here you can access guides for SAP HANA system replication automation and High Availability (HA) scenarios for SAP NetWeaver and SAP S/4HANA.

Additional resources, such as customer references, brochures or flyers, can be found at the SUSE Linux Enterprise Server for SAP Applications resource library: <https://www.suse.com/products/sles-for-sap/resource-library/>.

Supported high availability solutions by SUSE Linux Enterprise Server for SAP Applications overview: <https://documentation.suse.com/sles-sap/sap-ha-support/html/sap-ha-support/article-sap-ha-support.html>

Lastly, there are manual pages shipped with the product.

1.3 Errata

To deliver urgent smaller fixes and important information in a timely manner, the Technical Information Document (TID) for this document will be updated, maintained and published at a higher frequency:

- SAP S/4 HANA - Enqueue Replication 2 High Availability Cluster - Setup Guide - Errata (<https://www.suse.com/support/kb/doc/?id=7023714>)
- In addition to this guide, check the SUSE SAP Best Practice Guide Errata for other solutions (<https://www.suse.com/support/kb/doc/?id=7023713>).

1.4 Feedback


Several feedback channels are available:

Bugs and Enhancement Requests

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

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

Mail

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

2 Scope of this document

The document at hand explains how to:

- Plan a SUSE Linux Enterprise High Availability platform for SAP S/4HANA ABAP Platform, including SAP Enqueue Replication Server 2.
- Set up a Linux high availability platform and perform a basic SAP S/4HANA ABAP Platform installation including SAP Enqueue Replication Server 2 on SUSE Linux Enterprise.
- Integrate the high availability cluster with the SAP control framework via **sap-suse-cluster-connector** version 3, as certified by SAP.



Note

This guide implements the cluster architecture for enqueue replication version 2. For SAP S/4HANA ABAP Platform versions 1809 or newer enqueue replication version 2 is the default.

This guide focuses on the high availability of the central services. For SAP HANA system replication consult the guides for the performance-optimized or cost-optimized scenario (see [Section 1.2, "Additional documentation and resources"](#)).

3 Overview

This document describes setting up a pacemaker cluster using SUSE Linux Enterprise Server for SAP Applications 15 for the Enqueue Replication scenario. The focus is on matching the SAP S/4-HA-CLU 1.0 certification specifications and goals. For the setup described in this document, two nodes are used for the ASCS central services instance and ERS replicated enqueue instance. These

two nodes are controlled by the SUSE Linux Enterprise High Availability cluster. A third node is used for running the database and PAS and AAS application server instances. An additional fourth node is used as NFS server.

The goals for the setup include:

- Integration of the cluster with the native systemd based SAP start framework *sapstartsrv* to ensure that maintenance procedures do not break the cluster stability
- Rolling Kernel Switch (RKS) awareness
- Standard SAP installation to improve support processes
- Support of automated HA maintenance mode for SAP resources by implementing support of SAP HACheckMaintenanceMode and HASetMaintenanceMode
- Support of more than two cluster nodes for ASCS and ERS instances allowed

The updated certification SAP S/4-HA-CLU 1.0 redefines some of the test procedures and describes new expectations how the cluster should behave in special conditions. These changes allowed to improve the cluster architecture and to design it for easier usage and setup.

Shared SAP resources are located on a central NFS server.

The SAP instances themselves are installed on a shared disk to allow for switching over the file systems for proper functionality. A shared disk setup also allows using the SBD for the cluster fencing mechanism STONITH.

3.1 Differences to previous cluster architectures

- The described concept differs from the old stack that used a multi-state architecture. With the new certification there is a switch to a simpler model with primitives. This means: on one machine there is the ASCS instance with its own resources and on the other machine there is the ERS instance with its own resources.
- For SAP S/4HANA the new concept implies that, after a resource failure, the ASCS does not need to be started at the ERS side. The new enqueue architecture is also named ENSA2.
- Use of native systemd integration for SAP hostagent and instance's *sapstartsrv*. Refer to the SAP documentation for the necessary product version. See also SAP Note 3139184. SUSE **systemd** version 234 at least is needed. For details refer to the SUSE Linux Enterprise Server for SAP Applications product documentation. SUSE resource agents are needed; at least **sapstartsrv-resource-agents 0.9.1** and **resource-agents 4.x** from November 2021.

3.2 Three systems or more for ASCS, ERS, database and additional SAP instances

This document describes the installation of a distributed SAP system on three and more systems. In this setup, only two systems reside inside the cluster. The database and SAP dialog instances can be added to the cluster either by adding more nodes to the cluster or by installing the database on either of the existing nodes. However, it is recommended to install the database on a separate cluster. The cluster configuration for three and more nodes is described at the end of this document. The number of nodes within one cluster should be either two or an odd number. NOTE: Because the setup at hand focuses on the SAP S/4-HA-CLU 1.0 certification, the cluster detailed in this guide only manages the SAP instances ASCS and ERS.

If your database is SAP HANA, we recommend setting up the performance-optimized system replication scenario using the automation solution SAPHanaSR. The SAPHanaSR automation should be set up in an own two-node cluster. The setup is described in a separate best practices document available from the SUSE Best Practices Web page.




FIGURE 1: THREE SYSTEMS FOR THE CERTIFICATION SETUP

CLUSTERED MACHINES TWO-NODE SCENARIO

- one machine (valuga11) for ASCS
 - Hostname: sapen2as
- one machine (valuga12) for ERS

- Hostname: sapen2er

NON-CLUSTERED MACHINE

- one machine (valuga01) for DB and as application server
- Hostname: sapen2db
- Hostname: sapen2d1
- Hostname: sapen2d2

3.3 High availability for the database

Depending on your needs, you can increase the availability of the database, if your database is not already highly available by design.

3.3.1 SAP HANA system replication

A perfect enhancement of the three-node scenario described in this document is to implement an SAP HANA system replication (SR) automation.

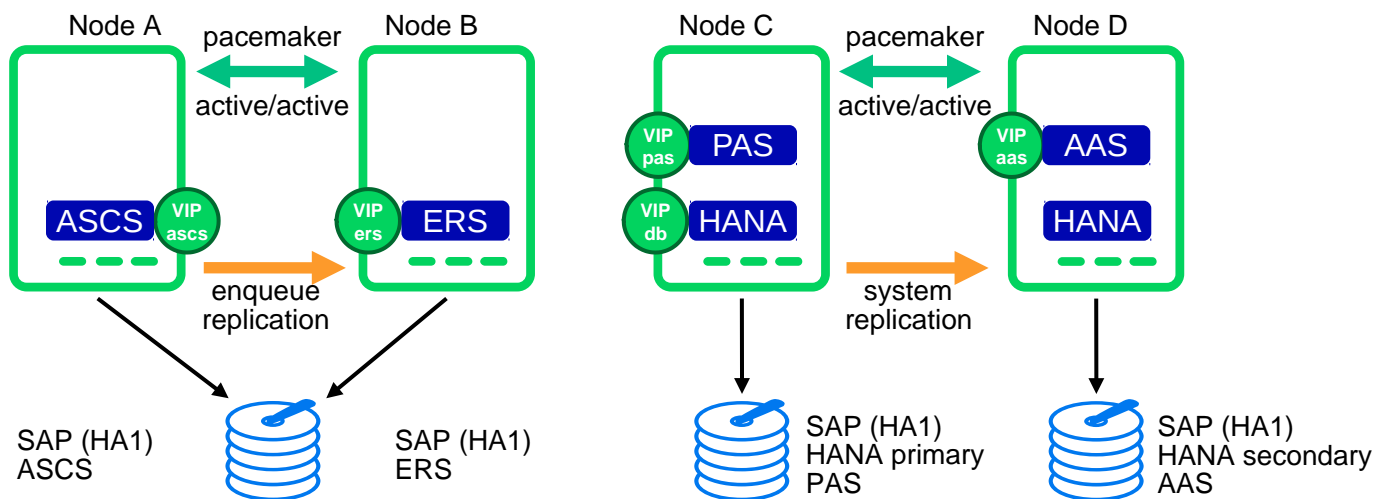


FIGURE 2: ONE CLUSTER FOR CENTRAL SERVICES, ONE FOR SAP HANA SR

TABLE 1: THE FOLLOWING OS/DATABASE COMBINATIONS ARE EXAMPLES FOR THIS SCENARIO

SUSE Linux Enterprise Server for SAP Applications 15	
Intel X86_64	SAP HANA DATABASE 2.0
IBM PowerLE	SAP HANA DATABASE 2.0



Note

Version for SAP S/4HANA ABAP Platform on Linux on AMD64/Intel 64 and PowerLE. More information about the supported combinations of OS and databases for SAP S/4HANA Server 1809 can be found at the SAP Product Availability Matrix at [SAP PAM \(https://apps.support.sap.com/sap/support/pam\)](https://apps.support.sap.com/sap/support/pam).

3.4 Integration of SAP S/4HANA into the cluster using the Cluster Connector

The integration of the HA cluster through the SAP control framework using the **sap_suse_cluster_connector** is of special interest. The service **sapstartsrv** controls SAP instances since SAP Kernel versions 6.40. One of the classic problems running SAP instances in a highly available environment is the following: If an SAP administrator changes the status (start/stop) of an SAP instance without using the interfaces provided by the cluster software, the cluster framework will detect that as an error status and will bring the SAP instance into the old status by either starting or stopping the SAP instance. This can result in very dangerous situations, if the cluster changes the status of an SAP instance during some SAP maintenance tasks. The new updated solution enables the central component **sapstartsrv** to report state changes to the cluster software. This avoids dangerous situations as previously described. More details can be found in the blog article "Using sap_vendor_cluster_connector for interaction between cluster framework and sapstartsrv" at <https://blogs.sap.com/2014/05/08/using-sapvendorclusterconnector-for-interaction-between-cluster-framework-and-sapstartsrv/comment-page-1/>.



Note

If you update from an SAP S/4HANA ABAP Platform version less than 1809, read SAP Note 2641019 carefully to adapt your cluster.

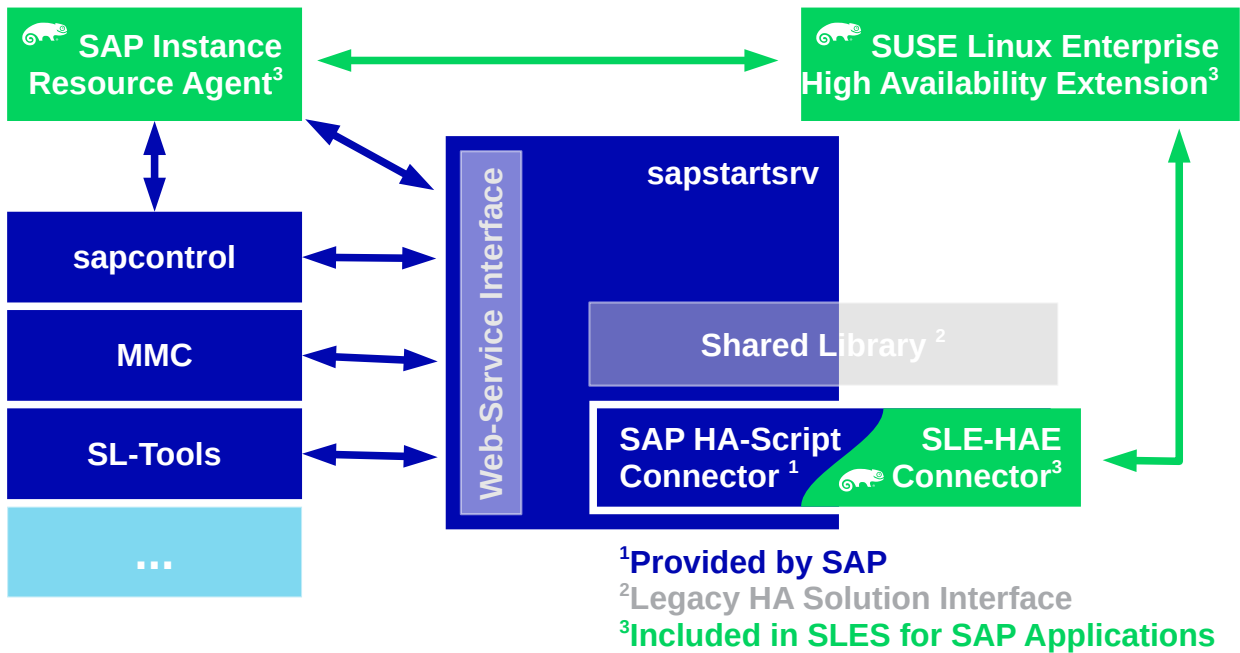


FIGURE 3: CLUSTER CONNECTOR TO INTEGRATE THE CLUSTER WITH THE SAP START FRAMEWORK



Note

For this scenario, an updated version of the **sap-suse-cluster-connector** is used. It implements the API version 3 for the communication between the cluster framework and the **sapstartsrv** service.

The new version of the **sap-suse-cluster-connector** allows starting, stopping and migrating an SAP instance. The integration between the cluster software and the **sapstartsrv** also implements the option to run checks of the HA setup using either the command line tool **sapcontrol** or even the SAP management consoles (SAP MMC or SAP MC). Since version 3.1.0 and later the maintenance mode of cluster resources triggered with SAP **sapcontrol** commands is supported. See also manual page `sap_suse_cluster_connector(8)`.

3.5 Disks and partitions

XFS is used for all SAP file systems besides the file systems on NFS.

3.5.1 Shared disk for cluster ASCS and ERS

The disk for the ASCS and ERS instances must be shared and assigned to the cluster nodes valuga11 and valuga12 in the two-node cluster example. In addition to the partitions for the file systems for the SAP instances, the disk also provides the partition to be used as SBD.

On valuga11 prepare the file systems for the shared disk. Create three partitions on the shared drive /dev/disk/by-id/SUSE-Example-A:

- Partition one (/dev/disk/by-id/SUSE-Example-A-part1) for SBD (7 MB)
- Partition two (/dev/disk/by-id/SUSE-Example-A-part2) for the first file system (10 GB) formatted with XFS
- Partition three (/dev/disk/by-id/SUSE-Example-A-part3) for the second file system (10 GB) formatted with XFS

To create the partitions, you can use either YaST or available command line tools. The following script can be used for non-interactive setups.

EXAMPLE 1: CREATE PARTITIONS AND FILE SYSTEMS FOR SBD AND ASCS/ERS ON VALUGA11

```
# parted -s /dev/disk/by-id/SUSE-Example-A print
# # we are on the 'correct' drive, right?
# parted -s /dev/disk/by-id/SUSE-Example-A mklabel gpt
# parted -s /dev/disk/by-id/SUSE-Example-A mkpart primary 1049k 8388k
# parted -s /dev/disk/by-id/SUSE-Example-A mkpart primary 8389k 10.7G
# parted -s /dev/disk/by-id/SUSE-Example-A mkpart primary 10.7G 21.5G
# mkfs.xfs /dev/disk/by-id/SUSE-Example-A-part2
# mkfs.xfs /dev/disk/by-id/SUSE-Example-A-part3
```

For these file systems, we recommend using normal partitions to keep the cluster configuration as simple as possible. However, you can also place these file systems in separate volume groups. In that case, you need to add further cluster resources to control the logical volume groups, and occasionally MD-Raid devices. The description of such a setup however is out of the scope of this guide.

After having partitioned the shared disk on valuga11, you need to request a partition table rescan on valuga12.

```
# partprobe; fdisk -l /dev/disk/by-id/SUSE-Example-A
```

During the SAP installation you need to mount /usr/sap/EN2/ASCS00 on valuga11 and /usr/sap/EN2/ERS10 on valuga12.

- valuga11: /dev/disk/by-id/SUSE-Example-A-part2 /usr/sap/EN2/ASCS00
- valuga12: /dev/disk/by-id/SUSE-Example-A-part3 /usr/sap/EN2/ERS10

3.5.2 Disk for database and dialog instances (HANA DB)

The disk for the database server is assigned to valuga01. The disk for the primary and secondary application server is assigned to valuga01.

VALUGA01

- Partition one (/dev/disk/by-id/SUSE-Example-B-part1) for the database (160 GB) formatted with XFS
- Partition two (/dev/disk/by-id/SUSE-Example-B-part2) for the PAS instance (10 GB) formatted with XFS
- Partition three (/dev/disk/by-id/SUSE-Example-B-part3) for the AAS instance (10 GB) formatted with XFS

You can either use YaST or available command line tools to create the partitions. The following script can be used for non-interactive setups.

EXAMPLE 2: CREATE PARTITIONS AND FILE SYSTEMS FOR DB AND APP SERVERS ON VALUGA01

```
# parted -s /dev/disk/by-id/SUSE-Example-B print
# # we are on the 'correct' drive, right?
# parted -s /dev/disk/by-id/SUSE-Example-B mklabel gpt
# parted -s /dev/disk/by-id/SUSE-Example-B mkpart primary 1049k 160G
# parted -s /dev/disk/by-id/SUSE-Example-B mkpart primary 160G 170G
# parted -s /dev/disk/by-id/SUSE-Example-B mkpart primary 170G 180G
# mkfs.xfs /dev/disk/by-id/SUSE-Example-B-part1
# mkfs.xfs /dev/disk/by-id/SUSE-Example-B-part2
# mkfs.xfs /dev/disk/by-id/SUSE-Example-B-part3
```

TO BE MOUNTED EITHER BY OS OR AN OPTIONAL CLUSTER

- valuga01: /dev/disk/by-id/SUSE-Example-B-part1 /hana
- valuga01: /dev/disk/by-id/SUSE-Example-B-part2 /usr/sap/EN2/D01
- valuga01: /dev/disk/by-id/SUSE-Example-B-part3 /usr/sap/EN2/D02



Note

D01: Since NetWeaver 7.5 the primary application server instance directory has been renamed to 'D<Instance_Number>'.

NFS SERVER

- 192.168.1.1:/data/export/S4_HA_CLU_10/EN2/sapmnt /sapmnt
- 192.168.1.1:/data/export/S4_HA_CLU_10/EN2/SYS /usr/sap/EN2/SYS

MEDIA

- 192.168.1.1:/Landscape /var/lib/Landscape

3.6 IP addresses and virtual names

Check if the file `/etc/hosts` contains at least the following address resolutions. Add those entries if they are missing.

```
192.168.1.100 valuga01
192.168.1.103 valuga11
192.168.1.104 valuga12

192.168.1.112 sapen2as
192.168.1.113 sapen2er
192.168.1.114 sapen2db
192.168.1.110 sapen2d1
192.168.1.111 sapen2d2
```

3.7 Mount points and NFS shares

In the present setup the directory `/usr/sap` is part of the root file system. You can also create a dedicated file system for that area and mount `/usr/sap` during the system boot. As `/usr/sap` contains the SAP control file `sapservices` and the **saphostagent**, the directory should not be placed on a shared file system between the cluster nodes.

You need to create the directory structure on all nodes that should run the SAP resource. The `SYS` directory will be located on an NFS share for all nodes.

- Create mount points and mounting NFS shares on all nodes.

EXAMPLE 3: MOUNT NFS SHARES ON ALL NODES

```
# mkdir -p /sapmnt /var/lib/Landscape
# mkdir -p /usr/sap/EN2/{ASCS00,D01,D02,ERS10,SYs}
# mount -t nfs 192.168.1.1:/data/export/S4_HA_CLU_10/EN2/sapmnt /sapmnt
# mount -t nfs 192.168.1.1:/data/export/S4_HA_CLU_10/EN2/SYS /usr/sap/EN2/SYS
# mount -t nfs 192.168.1.1:/Landscape /var/lib/Landscape
```

Mount options may depend on your particular environment. See also manual page `mount.nfs(8)`.

- For HANA: Create mount points for database at `valuga01`:

```
# mkdir /hana
# mount /dev/disk/by-id/SUSE-Example-B-part1 /hana/
# mkdir -p /hana/{shared,data,log}
# mount /dev/disk/by-id/SUSE-Example-B-part2 /usr/sap/EN2/D01
# mount /dev/disk/by-id/SUSE-Example-B-part3 /usr/sap/EN2/D02
```

You do not control the NFS shares via the cluster in the setup at hand. Thus you should add these file systems to `/etc/fstab` to get the file systems mounted during the next system boot.




FIGURE 4: FILE SYSTEM LAYOUT INCLUDING NFS SHARES

Prepare the three servers for the distributed SAP installation. Server 1 (`valuga01`) is used to install the SAP HANA database. Server 1 (`valuga01`) is used to install the PAS SAP instance. Server 1 (`valuga01`) is used to install the AAS SAP instances. Server 2 (`valuga11`) is used to install the ASCS SAP instances. Server 3 (`valuga12`) is used to install the ERS SAP instances.

- Mount the instance and database file systems on one specific node:

EXAMPLE 4: MOUNT SEVERAL FILE SYSTES FOR SAP S/4HANA ON DIFFERENT NODES

```
(ASCS valuga11) # mount /dev/disk/by-id/SUSE-Example-A-part2 /usr/sap/EN2/ASCS00
(ERS valuga12) # mount /dev/disk/by-id/SUSE-Example-A-part3 /usr/sap/EN2/ERS10
(DB valuga01) # mount /dev/disk/by-id/SUSE-Example-B-part1 /hana/
(Dialog valuga01) # mount /dev/disk/by-id/SUSE-Example-B-part2 /usr/sap/EN2/D01
(Dialog valuga01) # mount /dev/disk/by-id/SUSE-Example-B-part2 /usr/sap/EN2/D02
```

- As a result, the directory `/usr/sap/EN2/` should now look as follows:

```
# ls -la /usr/sap/EN2/
total 0
drwxr-xr-x 1 en2adm sapsys 70 28. Mar 17:26 ./
drwxr-xr-x 1 root sapsys 58 28. Mar 16:49 ../
drwxr-xr-x 7 en2adm sapsys 58 28. Mar 16:49 ASCS00/
drwxr-xr-x 1 en2adm sapsys 0 28. Mar 15:59 D02/
drwxr-xr-x 1 en2adm sapsys 0 28. Mar 15:59 D01/
drwxr-xr-x 1 en2adm sapsys 0 28. Mar 15:59 ERS10/
drwxr-xr-x 5 en2adm sapsys 87 28. Mar 17:21 SYS/
```



Note

The owner of the directory and files is changed during the SAP installation. By default, all of them are owned by *root*.

4 SAP installation

The overall procedure to install the distributed SAP system is as follows:

TASKS

1. Plan Linux user and group number scheme
2. Install the ASCS instance for the central services
3. Install the ERS to get a replicated enqueue scenario
4. Prepare the ASCS and ERS installations for the cluster take-over
5. Install the database

6. Install the primary application server instance (PAS)
7. Install additional application server instances (AAS)

The result will be a distributed SAP installation as illustrated here:




FIGURE 5: DISTRIBUTED INSTALLATION OF THE SAP SYSTEM

4.1 Linux user and group number scheme

Whenever asked by the SAP software provisioning manager (SWPM) which Linux User IDs or Group IDs to use, refer to the following table as an example:

Group	sapinst	1001
Group	sapsys	1002
Group	dbhshh	1003
User	en2adm	2001
User	sapadm	2002
User	dbhadm	2003

4.2 Installing ASCS on valuga11

Temporarily, as the local IP address, set the service IP address which you will later use in the cluster, because the installer needs to be able to resolve and use it. Make sure to use the correct virtual host name for each installation step. If applicable, make sure to also mount file systems like `/dev/disk/by-id/SUSE-Example-A-part2` and `/var/lib/Landscape/`.

```
# ip a a 192.168.1.112/24 dev eth0
# mount /dev/disk/by-id/SUSE-Example-A-part2 /usr/sap/EN2/ASCS00
# cd /var/lib/Landscape/media/SAP-media/SWPM20_P9/
```

```
# ./sapinst SAPINST_USE_HOSTNAME=sapen2as
```

- SWPM product installation path:
 - Installing SAP S/4HANA Server 1809 → SAP HANA DATABASE → Installation → Application Server ABAP → High-Availability System → ASCS Instance
- Use SID EN2
- Use instance number 00
- Deselect using FQDN
- All passwords: use <use-your-secure-pwd>
- Double-check during the parameter review, if virtual name **sapen2as** is used
- If you get an error during the installation about permissions, change the ownership of the ASCS directory

```
# chown -R en2adm:sapsys /usr/sap/EN2/ASCS00
```

4.3 Installing ERS on valuga12

Temporarily, as the local IP address, set the service IP address which you will later use in the cluster, because the installer needs to be able to resolve and use it. Make sure to use the correct virtual host name for each installation step.

```
# ip a a 192.168.1.113/24 dev eth0
# mount /dev/disk/by-id/SUSE-Example-A-part3 /usr/sap/EN2/ERS10
# cd /var/lib/Landscape/media/SAP-media/SWPM20_P9/
# ./sapinst SAPINST_USE_HOSTNAME=sapen2er
```

- SWPM product installation path:
 - Installing SAP S/4HANA Server 1809 → SAP HANA DATABASE → Installation → Application Server ABAP → High-Availability System → ERS Instance
- Use instance number 10
- Deselect using FQDN

- Double-check during the parameter review that virtual name **sapen2er** is used
- If you get an error during the installation about permissions, change the ownership of the ERS directory

```
# chown -R en2adm:sapsys /usr/sap/EN2/ERS10
```

- If you get a prompt to manually stop/start the ASCS instance, log in to *valuga11* as user *en2adm* and call 'sapcontrol'.

```
# sapcontrol -nr 00 -function Stop    # to stop the ASCS
# sapcontrol -nr 00 -function Start  # to start the ASCS
```

4.4 Subsequent steps for ASCS and ERS

After installation, you can perform several subsequent steps on the ASCS and ERS instances.

4.4.1 Stopping ASCS and ERS

To stop the ASCS and ERS instances, use the commands below. On *valuga11*, do the following:

```
# su - en2adm
# sapcontrol -nr 00 -function Stop
# sapcontrol -nr 00 -function StopService
```

On *valuga12*, do the following:

```
# su - en2adm
# sapcontrol -nr 10 -function Stop
# sapcontrol -nr 10 -function StopService
```

4.4.2 Disabling **systemd** services of the ASCS and the ERS SAP instance

This is mandatory for giving control over the instance to the HA cluster. See also manual pages *ocf_suse_SAPStartSrv(7)* and *SAPStartSrv_basic_Cluster(7)*.

```
# systemctl disable SAPEN2_00.service
# systemctl stop SAPEN2_00.service
# systemctl disable SAPEN2_10.service
# systemctl stop SAPEN2_10.service
```



Note

Stopping this instance services will stop the SAP instance as well. Starting the instance services will not start the SAP instances.

- Check the SAP systemd integration:

```
# systemctl list-unit-files | grep SAP
SAPEN2_10.service disabled
SAPEN2_10.service disabled
```

The instance services are indeed disabled, as required.

```
# systemctl list-unit-files | grep sap
saphostagent.service enabled
sapinit.service generated
saprouter.service disabled
saptune.service enabled
```

The mandatory `saphostagent` service is enabled. This is the installation default. Some more SAP related services might be enabled, for example the recommended `saptune`.

```
# cat /usr/sap/sapservices
systemctl --no-ask-password start SAPEN2_00 # sapstartsrv pf=/usr/sap/EN2/SYS/profile/
EN2_ASCS00_sapen2as
systemctl --no-ask-password start SAPEN2_10 # sapstartsrv pf=/usr/sap/EN2/SYS/profile/
EN2_ERS10_sapen2er
```

The `sapservices` file is still there for compatibility. It shows native systemd commands, one per line for each registered instance. You will find a SystemV style example in the appendix.

4.4.3 Integrating the cluster framework using *sap-suse-cluster-connector*

Install the package **sap-suse-cluster-connector** version 3.1.0 from the SUSE repositories:

```
# zypper in sap-suse-cluster-connector
```



Note

The package **sap-suse-cluster-connector** contains the version 3.x.x (SAP API 3). The package **sap-suse-cluster-connector** with version 3.0.x implements the SUSE SAP API version 3. New features like SAP Rolling Kernel Switch (RKS) and migration of ASCS

are only supported with this new version. The package **sap-suse-cluster-connector** with version 3.1.x supports in addition the maintenance mode of cluster resources triggered from SAP tools.

For the ERS and ASCS instances, edit the instance profiles `EN2_ASCS00_sapen2as` and `EN2_ERS10_sapen2er` in the profile directory `/usr/sap/EN2/SYS/profile/`.

Tell the **sapstartsrv** service to load the HA script connector library and to use the connector **sap-suse-cluster-connector**. Additionally, make sure the feature *Autostart* is **not** used.

```
service/halib = $(DIR_EXECUTABLE)/saphascriptco.so
service/halib_cluster_connector = /usr/bin/sap_suse_cluster_connector
```

Add the user `en2adm` to the unix user group `haclient`.

```
# usermod -a -G haclient en2adm
```

See also manual pages `sap_suse_cluster_connector(8)`, `usermod(8)` and `groupmod(8)`.

4.4.4 Adapting SAP profiles to match the SAP S/4-HA-CLU 1.0 certification

For the ASCS instance, change the start command from `Restart_Program_xx` to `Start_Program_xx` for the enqueue server (Enqueue Server 2). This change tells the SAP start framework **not** to self-restart the enqueue process. Such a restart would result in a loss of the locks.

File `/usr/sap/EN2/SYS/profile/EN2_ASCS00_sapen2as`

```
Start_Program_01 = local $_ENQ pf=$_PF
```

Optionally, you can limit the number of restarts of services (in the case of ASCS, this limits the restart of the message server).

For the ERS instance, change the start command from `Restart_Program_xx` to `Start_Program_xx` for the enqueue replication server (Enqueue Replicator 2).

File `/usr/sap/EN2/SYS/profile/EN2_ERS10_sapen2er`

```
Start_Program_00 = local $_ENQR pf=$_PF NR=$(SCSID)
```

4.4.5 Starting ASCS and ERS

To start the ASCS and ERS instances, use the commands below.

On *valuga11*, do the following:

```
# su - en2adm
# sapcontrol -nr 00 -function StartService EN2
# sapcontrol -nr 00 -function Start
```

On *valuga12*, do the following

```
# su - en2adm
# sapcontrol -nr 10 -function StartService EN2
# sapcontrol -nr 10 -function Start
```

4.5 Installing the database on *valuga01*

The HANA DB has very strict HW requirements. The storage sizing depends on many indicators. Check the supported configurations at [SAP HANA Hardware Directory \(https://support.sap.com/en/release-upgrade-maintenance.html#section_1969201630\)](https://support.sap.com/en/release-upgrade-maintenance.html#section_1969201630) and [SAP HANA TDI \(https://www.sap.com/documents/2016/05/e8705aae-717c-0010-82c7-eda71af511fa.html\)](https://www.sap.com/documents/2016/05/e8705aae-717c-0010-82c7-eda71af511fa.html).

```
# ip a a 192.168.1.114/24 dev eth0
# mount /dev/disk/by-id/SUSE-Example-B-part1 /hana
# mount /dev/disk/by-id/SUSE-Example-B-part2 /usr/sap/EN2/D01
# mount /dev/disk/by-id/SUSE-Example-B-part3 /usr/sap/EN2/D02
# cd /var/lib/Landscape/media/SAP-media/SWPM20_P9/
# ./sapinst SAPINST_USE_HOSTNAME=sapen2db
```

- SWPM product installation path:
 - Installing SAP S/4HANA Server 1809 → SAP HANA DATABASE → Installation → Application Server ABAP → High-Availability System → Database Instance
- Profile directory `/sapmnt/EN2/profile`
- Deselect using FQDN
- Database parameters : Database ID (DBSID) is DBH; Database Host is sapen2db; Instance Number is 00
- Database System ID enter Instance Number is 00; SAP Mount Directory is `/sapmnt/EN2/profile`
- Account parameters: change them in case of custom values needed

- Clean-up: select **Yes**, remove operating system users from the group `sapinst`.
- Double-check during the parameter review, if virtual name **sapen2db** is used

4.6 Installing the primary application server (PAS) on valuga01

```
# ip a a 192.168.1.110/24 dev eth0
# mount /dev/disk/by-id/SUSE-Example-B-part2 /usr/sap/EN2/D01
# cd /var/lib/Landscape/media/SAP-media/SWPM20_P9/
# ./sapinst SAPINST_USE_HOSTNAME=sapen2d1
```

- SWPM product installation path:
 - Installing SAP S/4HANA Server 1809 → SAP HANA DATABASE → Installation → Application Server ABAP → High-Availability System → Primary Application Server Instance
- Use instance number 01.
- Deselect using FQDN.
- For this example setup, we have used a default secure store key.
- Do not install **Diagnostic Agent**.
- Do not use SLD.
- Double-check during the parameter review if virtual name **sapen2d1** is used.

4.7 Installing an Additional Application Server (AAS) on valuga01

```
# ip a a 192.168.1.111/24 dev eth0
# mount /dev/disk/by-id/SUSE-Example-B-part3 /usr/sap/EN2/D02
# cd /var/lib/Landscape/media/SAP-media/SWPM20_P9/
# ./sapinst SAPINST_USE_HOSTNAME=sapen2d2
```

- SWPM product installation path:
 - Installing SAP S/4HANA Server 1809 → SAP HANA DATABASE → Installation → Application Server ABAP → High-Availability System → Additional Application Server Instance
- Use instance number 02.

- Deselect using FQDN.
- Do not install **Diagnostic Agent**.
- Double-check during the parameter review if virtual name **sapen2d2** is used.

5 Implementing the cluster

The main procedure to implement the cluster is as follows:

TASKS

1. OS preparation and install the cluster software
2. Configure the cluster base including corosync and resource manager
3. Configure the cluster resources
4. Tune the cluster timing in special for the SBD.



Note

Before you continue to set up the cluster, perform the following actions: First stop all SAP instances. Then remove the (manually added) IP addresses on the cluster nodes. Finally unmount the file systems which will be controlled by the cluster later.



Note

The SBD device/partition needs to be created beforehand. Double-check which device/partition to use! In this setup guide a partition `/dev/disk/by-id/SUSE-Example-A-part1` is already reserved for SBD usage.

5.1 Preparing the operating system and installing the cluster software

- Set up and enable **chrony** with **YaST**.
- Install the RPM pattern **ha_sles** on both cluster nodes.


```
# zypper in -t pattern ha_sles
```

5.2 Configuring the cluster base

TASKS

- To configure the cluster base, you can use either YaST or the interactive command line tool **ha-cluster-init**. The example below uses the command line wizard.
- Install and configure the watchdog device on the first machine.

Instead of deploying the software based solution, a hardware-based watchdog device should preferably be used. The following example uses the software device but can be easily adapted to the hardware device.

```
# modprobe softdog
# echo softdog > /etc/modules-load.d/watchdog.conf
# systemctl restart systemd-modules-load
# lsmod | egrep "(wd|dog|i6|iT|ibm)"
```

- Install and configure the cluster stack on the first machine

```
# ha-cluster-init -u -s /dev/disk/by-id/SUSE-Example-A-part1
```

- Join the second node

On the second node, perform some preparation steps.

```
# modprobe softdog
# echo softdog > /etc/modules-load.d/watchdog.conf
# systemctl restart systemd-modules-load
# lsmod | egrep "(wd|dog|i6|iT|ibm)"
```

To configure the cluster base, you can use either YaST or the interactive command line tool **ha-cluster-join**. The example below uses the command line wizard.

```
# ha-cluster-join -c valuga11
```

- The *crm_mon -1r* output should look as follows:

```
Stack: corosync
Current DC: valuga11 (version 1.1.18+20180430.b12c320f5-1.14-b12c320f5) - partition with
quorum
```

```

Last updated: Mon Jan 28 13:10:37 2019
Last change: Wed Jan 23 09:52:57 2019 by root via cibadmin on valuga11

2 nodes configured
1 resource configured

Online: [ valuga11 valuga12 ]

stonith-sbd (stonith:external/sbd): Started valuga11

```

5.3 Configuring cluster resources

A changed SAPIInstance resource configuration is needed for SAP S/4HANA to **not** use the multi-state construct anymore. Move to a more cluster-like construct to start and stop the ASCS and the ERS instances themselves, but **not** the complete multi-state resource.

With the new version of ENSA2, the ASCS instance can be started on the same host. There is no longer a need to follow the ERS instance. The ASCS instance receives the enqueue lock table over the network from the ERS instance. If no other node is available, the ASCS instance will be started on the same host where the ERS instance is running.

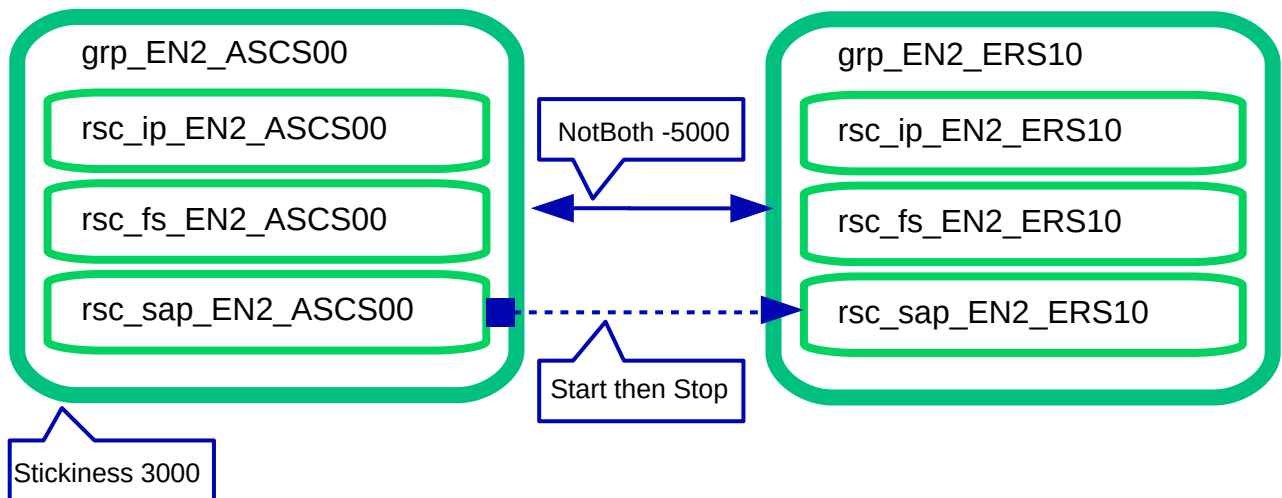


FIGURE 6: RESOURCES AND CONSTRAINTS

Another benefit of this concept is that you can work with native (mountable) file systems instead of a shared (NFS) file system for the SAP instance directories.

5.3.1 Preparing the cluster for adding the resources

To prevent the cluster from starting partially defined resources, set the cluster to the maintenance mode. This deactivates all monitor actions.

As user *root*, type the following command:

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

5.3.2 Configuring the resources for the ASCS instance

First, configure the resources for the file system, IP address and the SAP instance. You need to adapt the parameters for your specific environment. The shown file system and SAPInstance monitor timeouts are a trade-off between fast recovery vs. resilience against sporadic temporary NFS issues. You may slightly increase it to fit your infrastructure. The SAPInstance timeout needs to be higher than the file system timeout. Consult your storage or NFS server documentation for appropriate timeout values. See also manual pages `ocf_heartbeat_Fileystem(7)`, `ocf_heartbeat_SAPInstance(7)` and `nfs(5)`.

EXAMPLE 5: ASCS PRIMITIVE

```
primitive rsc_fs_EN2_ASCS00 Filesystem \  
  params device="/dev/disk/by-id/SUSE-Example-A-part2" directory="/usr/sap/EN2/  
ASC00" \  
  fstype=xfstype=xfstype \  
  op start timeout=60s interval=0 \  
  op stop timeout=60s interval=0 \  
  op monitor interval=20s timeout=40s  
primitive rsc_ip_EN2_ASCS00 IPAddr2 \  
  params ip=192.168.1.112 \  
  op monitor interval=10s timeout=20s  
primitive rsc_sap_EN2_ASCS00 SAPInstance \  
  op monitor interval=11 timeout=60 on-fail=restart \  
  params InstanceName=EN2_ASCS00_sapen2as \  
    START_PROFILE="/usr/sap/EN2/SYS/profile/EN2_ASCS00_sapen2as" \  
    AUTOMATIC_RECOVER=false \  
  meta resource-stickiness=5000
```

See also manual pages `ocf_heartbeat_SAPInstance(7)`, `ocf_heartbeat_IPAddr2(7)` and `ocf_heartbeat_Fileystem(7)`.

EXAMPLE 6: ASCS GROUP

```
group grp_EN2_ASCS00 \  
  primitive rsc_fs_EN2_ASCS00 Filesystem \  
  primitive rsc_ip_EN2_ASCS00 IPAddr2 \  
  primitive rsc_sap_EN2_ASCS00 SAPInstance
```

```
rsc_ip_EN2_ASCS00 rsc_fs_EN2_ASCS00 rsc_sap_EN2_ASCS00 \  
meta resource-stickiness=3000
```

Create a *txt* file (like *crm_ascs.txt*) with your preferred text editor. Add both examples (primitives and group) to that file and load the configuration to the cluster manager configuration.

As user *root*, type the following command:

```
# crm configure load update crm_ascs.txt
```

5.3.3 Configuring the resources for the ERS instance

Second, configure the resources for the file system, IP address and the SAP instance. You need to adapt the parameters for your specific environment. The shown file system and SAPInstance monitor timeouts are a trade-off between fast recovery versus resilience against sporadic temporary NFS issues. You may slightly increase it to fit your infrastructure. The SAPInstance timeout needs to be higher than the file system timeout. Consult your storage or NFS server documentation for appropriate timeout values. See also manual pages *ocf_heartbeat_Fileystem(7)*, *ocf_heartbeat_SAPInstance(7)* and *nfs(5)*.

The specific parameter *IS_ERS = true* must only be set for the ERS instance.

EXAMPLE 7: ERS PRIMITIVE

```
primitive rsc_fs_EN2_ERS10 Filesystem \  
  params device="/dev/disk/by-id/SUSE-Example-A-part3" directory="/usr/sap/EN2/  
ERS10" fstype=trfs \  
  op start timeout=60s interval=0 \  
  op stop timeout=60s interval=0 \  
  op monitor interval=20s timeout=40s  
primitive rsc_ip_EN2_ERS10 IPAddr2 \  
  params ip=192.168.1.113 \  
  op monitor interval=10s timeout=20s  
primitive rsc_sap_EN2_ERS10 SAPInstance \  
  op monitor interval=11 timeout=60 on-fail=restart \  
  params InstanceName=EN2_ERS10_sapen2er \  
    START_PROFILE="/usr/sap/EN2/SYS/profile/EN2_ERS10_sapen2er" \  
    AUTOMATIC_RECOVER=false IS_ERS=true
```

See also manual pages *ocf_heartbeat_SAPInstance(7)*, *ocf_heartbeat_IPAddr2(7)* and *ocf_heartbeat_Fileystem(7)*.

EXAMPLE 8: ERS GROUP

```
group grp_EN2_ERS10 \  
  primitive rsc_ip_EN2_ERS10 IPAddr2 \  
  primitive rsc_sap_EN2_ERS10 SAPInstance \  
  primitive rsc_fs_EN2_ERS10 Filesystem
```

```
rsc_ip_EN2_ERS10 rsc_fs_EN2_ERS10 rsc_sap_EN2_ERS10
```

Create a *txt* file (like *crm_ers.txt*) with your preferred text editor. Add both examples (primitives and group) to that file and load the configuration to the cluster manager configuration.

As user *root*

```
# crm configure load update crm_ers.txt
```

5.3.4 Configuring the colocation constraints between ASCS and ERS

Compared to the ENSA1 configuration, the constraints between the ASCS and ERS instances are changed. An ASCS instance should avoid starting up on the cluster node running the ERS instance if any other node is available. Today the ENSA2 setup can resynchronize the lock table over the network.

If the ASCS instance has been started by the cluster on the ERS node, the ERS instance should be moved to another cluster node (*col_sap_EN2_no_both*). This constraint is needed to ensure that the ERS instance will synchronize the locks again and the cluster is ready for an additional take-over.

EXAMPLE 9: LOCATION CONSTRAINT

```
colocation col_sap_EN2_no_both -5000: grp_EN2_ERS10 grp_EN2_ASCS00
order ord_sap_EN2_first_start_ascsc Optional: rsc_sap_EN2_ASCS00:start \
rsc_sap_EN2_ERS10:stop symmetrical=false
```

Create a *txt* file (like *crm_col.txt*) with a text editor. Add both constraints to that file and load the configuration to the cluster manager configuration.

As user *root*, type the following command:

```
# crm configure load update crm_col.txt
```

5.3.5 Activating the cluster

The last step is to end the cluster maintenance mode and to allow the cluster to detect already running resources.

As user *root*, type the following command:

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

6 Administration

6.1 Dos and don'ts



Note

Before each test, verify that the cluster is in idle state, no migration constraints are active, and no resource failure messages are visible. Start each procedure with a clean set-up. See also manual pages `sap_suse_cluster_connector(8)`, `cs_clusterstate(8)`, `crm(8)` and `crm_mon(8)`.

6.1.1 Migrating the ASCS instance

To **migrate** the ASCS SAP instance, you should use SAP tools such as the SAP management console. This will trigger `sapstartsrv` to use the `sap-suse-cluster-connector` to migrate the ASCS instance. As user `en2adm` you can run the command below to migrate the ASCS. This will always migrate the ASCS to the ERS side which will keep the SAP enqueue locks.

As user `en2adm`, type the command:

```
# sapcontrol -nr 00 -function HAFailoverToNode ""
```

6.1.2 Never block resources

With SAP S/4-HA-CLU 1.0 it is **no longer allowed to block resources** from being controlled manually. This means using the variable `BLOCK_RESOURCES` in `/etc/sap_suse_cluster_connector` is not allowed anymore.

6.1.3 Always use unique instance numbers

Currently all SAP **instance numbers controlled by the cluster must be unique**. If you need multiple dialog instances, such as D00, running on different systems, they should not be controlled by the cluster.

6.1.4 Setting the cluster to maintenance mode

The procedure to set the cluster into maintenance mode can be executed as user *root* or *sidadm*.

As user *root*, type the following command:

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

See also manual page `crm(8)`.

As user *en2adm*, type the following command (the full path is needed):

```
# /usr/sbin/crm configure property maintenance-mode="true"
```

6.1.5 Stopping the cluster maintenance

The procedure to end the maintenance mode for the cluster can be executed as user *root*. Type the following command:

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

6.1.6 Starting the resource maintenance mode

The procedure to start the resource maintenance mode can be executed as user *en2adm*. This sets the ASCS and ERS cluster resource to **unmanaged**.

As user *en2adm*, type the command:

```
# sapcontrol -nr 00 -function HASETMaintenanceMode 1
```

6.1.7 Stopping the Resource Maintenance Mode

The procedure to start the resource maintenance mode can be executed as user *en2adm*. This sets the ASCS and ERS cluster resource to **managed**.

As user *en2adm*, type the command:

```
# sapcontrol -nr 00 -function HASETMaintenanceMode 0
```

6.1.8 Cleaning up resources

You can also **clean-up resource failures**. Failures of the ASCS will be automatically deleted to allow a failback after the configured period of time. For all other resources, you can clean-up the status including the failures.

As user *root*, type the following command:

```
# crm resource cleanup RESOURCE-NAME
```

6.2 Testing the cluster

It is strongly recommended to perform at least the following tests before you go into production with your cluster:

6.2.1 Checking Product Names with HAGetFailoverConfig

Check if the name of the SUSE cluster solution is shown in the output of **sapcontrol** or the SAP management console. This test checks the status of the SAP NetWeaver cluster integration.

As user *en2adm*, type the following command:

```
# sapcontrol -nr 00 -function HAGetFailoverConfig
```

6.2.2 Running SAP checks using HACheckConfig and HACheckFailoverConfig

Check if the HA configuration tests are passed successfully and do not produce error messages.

As user *en2adm*, type the following commands:

```
# sapcontrol -nr 00 -function HACheckConfig  
# sapcontrol -nr 00 -function HACheckFailoverConfig
```

6.2.3 Manually migrating ASCS

Check if manually migrating the ASCS instance using HA tools works properly.

As user *root*, run the following commands:

```
# crm resource migrate rsc_sap_EN2_ASCS00 force  
## wait until the ASCS is been migrated to the ERS host
```



```
# crm resource unmigrate rsc_sap_EN2_ASCS00
```

6.2.4 Migrating ASCS using HAFailoverToNode

Check if moving the ASCS instance using SAP tools like sapcontrol works properly.

As user *en2adm*, type the following command:

```
# sapcontrol -nr 00 -function HAFailoverToNode ""
```

6.2.5 Testing ASCS migration after operating system failure

Check if the ASCS instance moves correctly after a node failure. This test will immediately trigger a hard reboot of the node.

As user *root*, type the following command:

```
## on the ASCS host  
# echo b >/proc/sysrq-trigger
```

6.2.6 Restarting ASCS in-place using *Stop* and *Start*

Check if the in-place restart of the SAP resources have been processed correctly. The SAP instance should not failover to an other node, it must start on the same node where it has been stopped.

As user *en2adm*, do the following:

```
## example for ASCS  
# sapcontrol -nr 00 -function Stop  
## wait till the ASCS is completely down  
# sapcontrol -nr 00 -function Start
```

6.2.7 Automated restart of the ASCS instance (simulating rolling kernel switch)

The next test should proof that the cluster solution did not interact neither try to restart the ASCS instance during a maintenance procedure. In addition, it should verify that no locks are lost during the restart of an ASCS instance during an RKS procedure. The cluster solution should recognize that the restart of the ASCS instance was expected. No failure or error should be reported or counted.

Optionally, you can set locks and verify that they still exist after the maintenance procedure. There are multiple ways to do that. One example test can be performed as follows:

1. Log in to your SAP system and open the transaction SU01.
2. Create a new user. Do not finish the transaction to see the locks.
3. With the SAP MC / MMC, check if there are locks available.
4. Open the ASCS instance entry and go to *Enqueue Locks*.
5. With the transaction SM12, you can also see the locks.

Do this test multiple times in a short time frame. The restart of the ASCS instance in the example below happens five times.

As user *en2adm*, create and execute the following script:

```
$ cat ascs_restart.sh
#!/bin/bash
for lo in 1 2 3 4 5; do
  echo LOOP "$lo - Restart ASCS00"
  sapcontrol -host sapen2as -nr 00 -function StopWait 120 1
  sleep 1
  sapcontrol -host sapen2as -nr 00 -function StartWait 120 1
  sleep 1
done
```

```
$ bash ascs_restart.sh
```

6.2.8 Rolling kernel switch procedure

The rolling kernel switch (RKS) is an automated procedure that enables the kernel in an ABAP system to be exchanged without any system downtime. During an RKS, all instances of the system, and generally all SAP start services (sapstartsrv), are restarted.

1. Check in SAP note 953653 whether the new kernel patch is RKS compatible to your currently running kernel.
2. Check SAP note 2077934 - Rolling kernel switch in HA environments.
3. Download the new kernel from the SAP service market place.

4. Make a backup of your current central kernel directory.
5. Extract the new kernel archive to the central kernel directory.
6. Start the RKS via SAP MMC, system overview (transaction SM51) or via command line.
7. Monitor and check the version of your SAP instances with the SAP MC / MMC or with **sapcontrol**.

As user *en2adm*, type the following commands:

```
## sapcontrol [-user <sidadm psw>] -host <host> -nr <INSTANCE_NR> -function UpdateSystem
120 300 1
# sapcontrol -user en2adm <use-your-secure-pwd> -host sapen2as -nr 00 -function
UpdateSystem 120 300 1
# sapcontrol -nr 00 -function GetSystemUpdateList -host sapen2as \
-user en2adm <use-your-secure-pwd>
# sapcontrol -nr 00 -function GetVersionInfo -host sapen2as \
-user en2adm <use-your-secure-pwd>
# sapcontrol -nr 10 -function GetVersionInfo -host sapen2er \
-user en2adm <use-your-secure-pwd>
# sapcontrol -nr 01 -function GetVersionInfo -host sapen2d1 \
-user en2adm <use-your-secure-pwd>
# sapcontrol -nr 02 -function GetVersionInfo -host sapen2d2 \
-user en2adm <use-your-secure-pwd>
```

6.2.9 Additional tests

In addition to the already performed tests, you should do the following:

- Check the recoverable and non-recoverable outage of the message server process.
- Check the non-recoverable outage of the SAP enqueue server process.
- Check the outage of the SAP Enqueue Replication Server 2.
- Check the outage and restart of sapstartsrv.
- Check the simulation of an upgrade.
- Check the simulation of cluster resource failures.

7 Multi-node cluster setups for an SAP S/4HANA

Multinode cluster setups mean cluster configurations with more than two nodes. Depending from the starting point it is possible to extend a two node cluster setup or directly start with more than two nodes for an ASCS / ERS high availability setup. The examples below will show the setting up of Multi-node cluster and the extension of an existing two node cluster pair. The major configuration changes will be shown and the basic preparation of the new cluster member node. The task list to set up the three node cluster is similar to the task list of the two node cluster. But some details are described different here to get a diskless SBD setup. Such a diskless SBD setup is an optional improvement for three nodes, but does not work for two nodes. On the other hand, priority fencing is an optional improvement for two nodes, but does not work for three nodes. See the SUSE Linux Enterprise High Availability product documentation for details (<https://documentation.suse.com/sle-ha/15-SP3/single-html/SLE-HA-administration/#pro-ha-storage-protect-fencing>).



Note

When extending a cluster from two to three nodes, make sure to not use priority fencing.

7.1 Extending an existing two-node cluster configuration

TASKS

1. Back up the current cluster.
2. Install the OS on the new node.
3. Patch the existing nodes.
4. Prepare the new node's operating system.
5. Prepare SAP.
6. Install the cluster software on the new node.
7. Joining the cluster.
8. Test the new cluster configuration.

7.1.1 Backing up the current cluster

To **back up the current cluster**, perform a

- Backup of your system including
 - cluster configuration
 - corosync.conf
 - all data and configuration which are important and customized and not default

The system is configured as described in the SUSE Best Practices document *SAP S/4 HANA - Enqueue Replication 2 High Availability Cluster - Setup Guide*.

To **back up the cluster configuration**, go to one of the cluster nodes and save the cluster configuration with *crm* and *crm report* commands:

```
# crm configure save 2node.txt
```

Back up the existing /etc/corosync/corosync.conf and all other files which may be important for a restore. This example is one method creating a backup. The important point is using an **external** destination.

```
# tar cvf /<path to an external storage>/bck_2nodes_configuration.tar \  
  /etc/corosync/corosync.conf \  
  2node.txt \  
  <add your additional files here>
```

7.1.2 Installing the operating system of the new cluster node member

We recommend using automating the installation to ensure that the system setup across nodes is identical. Make sure to document any additional steps you take beyond the automated setup. In our example, we deploy our machines with an AutoYaST configuration file and run a post step script which does the basic configuration.

7.1.3 Patching existing nodes

If applicable, install the latest updates and patches on your existing nodes. Alternatively, if you are using frozen repositories such as those provided by SUSE Manager, add the new system to the same repositories, so they have the same patch level as your existing nodes.

Use **zypper patch** or **zypper update** depending on your company's rules.

- It is recommended installing the latest available patches to guarantee system stability and hardening. Bug fixing and security patches help avoid unplanned outages and make the system less vulnerable.

There are multiple ways:

```
# zypper patch --category security
or
# zypper patch --severity important
or
# zypper patch
## or
# zypper update
```

7.1.4 Preparing the new node's operating system

EXAMPLE 10: CHECK FOR VALID DNS, TIME SYNCHRONIZATION, NETWORK SETTINGS

- Verify that DNS is working:

```
# ping <hostname>
```

- Set up **chrony** (this is best done with **YaST**) and enable it:

```
# yast ntp-client
```

- Check the network settings:

```
# ip r
```

You may run into trouble if there is no valid or no default route configured.

- Patch the new system like the existing nodes, verify that your systems have the same patch level and that all required reboots have been performed.

7.1.5 Installing cluster software on the new node

EXAMPLE 11: INSTALLING PACKAGES

- Install the pattern **ha_sles** on the new cluster node:

```
# zypper in -t pattern ha_sles
```

- Install the package **sap-suse-cluster-connector** version 3.1.0 from the SUSE repositories:

```
# zypper in sap-suse-cluster-connector
```

After installing all necessary packages, compare installed package versions.

EXAMPLE 12: CHECK THAT ALL NODES HAVE THE SAME SOFTWARE PACKAGES AND VERSIONS

- On the existing cluster nodes:

```
# rpm -qa | sort >rpm_valuga11.log
```

- On the new node:

```
# rpm -qa | sort >rpm_valuga13.log
```

- Copy the file from one node to the other and compare the two versions:

```
# vimdiff rpm_valuga13.log rpm_valuga11.log
```



Note

If there are any differences, fix them first before you proceed.

EXAMPLE 13: INSTALL AND CONFIGURE THE WATCHDOG DEVICE ON THE NEW MACHINE.

Instead of deploying the software-based solution, preferably use a hardware-based watchdog device. The following example uses the software device but can be easily adapted to the hardware device.

```
# modprobe softdog
# echo softdog > /etc/modules-load.d/watchdog.conf
# systemctl restart systemd-modules-load
# lsmod | egrep "(wd|dog|i6|iT|ibm)"
```



Note

Ensure that the new node is connected to the same SBD disk connected to the existing two nodes. Ensure that the new node has the same exact SBD configurations already exist on the two existing nodes

```
# sbd -d /dev/disk/by-id/SUSE-Example-A-part1 dump
```

7.1.6 Preparing the SAP installation on the new node

With SWPM 2.0 (SP4 or later), which is part of the SL Toolset, SAP provides a new option which can perform all necessary steps to prepare a fresh install server to be able to fit into an existing SAP system. This new option will help us to prepare a new host which can later run either the ASCS or ERS in the cluster environment.

You need to create the directory structure that should run the SAP resource. The SYS directory is located on an NFS share for all nodes.

- Create mount points and mount NFS shares on the new added node (valuga13):

EXAMPLE 14: MOUNT NFS SHARES ON VALUGA13

```
# mkdir -p /sapmnt /var/lib/Landscape
# mkdir -p /usr/sap/EN2/{ASCS00,ERS10,SYS}
# mount -t nfs 192.168.1.1:/data/export/S4_HA_CLU_10/EN2/sapmnt /sapmnt
# mount -t nfs 192.168.1.1:/data/export/S4_HA_CLU_10/EN2/SYS /usr/sap/EN2/SYS
# mount -t nfs 192.168.1.1:/Landscape /var/lib/Landscape
```

As this directory needs to be available at all time, make sure it is mounted during boot. This can be achieved by put the information into the `/etc/fstab`.

In the next step the following information are required: - profile directory - password for SAP System Administrator - UID for sapadm

```
# cd /var/lib/Landscape/media/SAP-media/SWPM20_P9/
# ./sapinst
```


- SWPM product installation path:
 - Installing SAP S/4HANA Server 1809 → SAP HANA DATABASE → Installation → Application Server ABAP → High-Availability System → Prepare Additional Host
- Use /sapmnt/EN2/profile for the profile directory
- All passwords: <use-your-secure-pwd>
- UID for sapadm: 2002

EXAMPLE 15: POST STEP PROCEDURE AFTER SAP PREPARATION STEP

- Add the user *en2adm* to the unix user group *haclient*.

```
# usermod -a -G haclient en2adm
```

- Register the ASCS and the ERS SAP instance:

```
# LD_LIBRARY_PATH=/usr/sap/hostctrl/exe:$LD_LIBRARY_PATH; export
LD_LIBRARY_PATH
# /usr/sap/hostctrl/exe/sapstartsrv pf=/usr/sap/<SID>/SYS/profile/
<SID>_ERS<instanceNumberErs>_<virtHostNameErs>
> -reg
# /usr/sap/hostctrl/exe/sapstartsrv pf=/usr/sap/<SID>/SYS/profile/
<SID>_ASCS<instanceNumberAscs>_<virtHostNameA
scs> -reg
```



Note

This must be done for each instance. Call `sapstartsrv` with parameters **pf=** <profile-of-the-sap-instance> and **-reg**.

- Disable systemd services of the ASCS and the ERS SAP instance:

```
# systemctl disable SAPEN2_00.service
# systemctl disable SAPEN2_10.service
```



Note

This is mandatory for giving control over the instance to the HA cluster.

- Check the SAP systemd integration:

```
# systemctl list-unit-files | grep sap
saphostagent.service enabled
sapinit.service generated
saprouter.service disabled
saptune.service enabled
# systemctl list-unit-files | grep SAP
SAPEN2_10.service disabled
SAPEN2_10.service disabled
# cat /usr/sap/sapservices
systemctl --no-ask-password start SAPEN2_00 # sapstartsrv pf=/usr/sap/EN2/SYS/
profile/EN2_ASCS00_sapen2as
systemctl --no-ask-password start SAPEN2_10 # sapstartsrv pf=/usr/sap/EN2/SYS/
profile/EN2_ERS10_sapen2er
```

7.1.7 Adding the new node to the cluster

EXAMPLE 16: ADD THE NEW NODE TO THE CLUSTER

Check if the SBD device is available in case the SBD stonith method is in place for the two nodes. If the existing cluster using a different, supported stonith mechanism check and verify them too for the new cluster node.

```
# sbd -d /dev/disk/by-id/SUSE-Example-A-part1 dump
```

- Joining the cluster with *sleha-join*

```
# sleha-join -c valuga11
```

After the new node has joined into the cluster the configuration must be adapted to the new situation. Double check if the join was successful and verify the */etc/corosync/corosync.conf*

```
# grep votes -n2 /etc/corosync/corosync.conf
```

The values **expected_votes** and **two_node** should now look like this on all nodes:

- expected_votes: 3
- two_node: 0

Modify the cluster configuration and set a new colocation rule with *crm*:

```
# crm configure delete col_sap_EN2_no_both
# crm configure colocation ASCS00_ERS10_separated_EN2 -5000: grp_EN2_ERS10
grp_EN2_ASCS00
```

7.1.8 Testing the new cluster configuration

It is highly recommended to run certain test to verify that the new configuration is working as expected. A list of test can be found in the basic setup for the two node cluster above.

7.2 Pro's and con's for odd and even numbers of cluster nodes

There are certain use cases and infrastructure requirements which end up in different installation setup's. We will cover some advantages and disadvantages of special configuration below:

- The two node cluster and two locations
 - Advantage: symmetric spread of all nodes over all locations
 - Disadvantage: no diskless SBD feature allowed for all two node clusters
- The two node cluster and more than two locations
 - Advantage: SBD device can be provided from there (must be HA himself)
 - Advantage: cluster could operate with three SBD devices from different locations
 - Disadvantage: no diskless SBD feature allowed for all two node clusters
- The three node cluster and two locations

- Advantage: less complex infrastructure
- Advantage: diskless SBD feature is allowed
- Disadvantage: "pre selected" location (two node + one node)
- The three node cluster and three locations
 - Advantage: symmetric spread of all nodes over all locations
 - Advantage: diskless SBD feature is allowed
 - Disadvantage: higher planing effort and complexity for infrastructure planning





8 References

For more information, see the documents listed below.

8.1 Pacemaker

- Pacemaker 1.1 Configuration Explained: https://clusterlabs.org/pacemaker/doc/en-US/Pacemaker/1.1/html/Pacemaker_Explained/ 

8.2 Related SUSE TID

- SUSE SAP Best Practice Guide Errata (<https://www.suse.com/support/kb/doc?id=7023713> )
- SAP S/4 HANA - Enqueue Replication 2 High Availability Cluster - Setup Guide (<https://www.suse.com/support/kb/doc?id=7023714> )
- Integration of sap-suse-cluster-connector does not work as expected (<https://www.suse.com/support/kb/doc?id=000019244> )
- Long Client hang to Cluster after failover of ERS Instance (<https://www.suse.com/support/kb/doc?id=000019293> )

- `sap_suse_cluster_connector` stuck as HAActive: FALSE (<https://www.suse.com/support/kb/doc/?id=000019924>)
- Use of Filesystem resource for ASCS/ERS HA setup not possible (<https://www.suse.com/support/kb/doc/?id=000019944>)

8.3 Related SAP Notes

- 953653 - Rolling Kernel Switch (<https://launchpad.support.sap.com/#/notes/953653>)
- 1092448 - IBM XL C/C++ runtime environment for Linux on system p (<https://launchpad.support.sap.com/#/notes/1092448>)
- 1153713 - Problems with SAP Management Console (Java) (<https://launchpad.support.sap.com/#/notes/1153713>)
- 1552925 - Linux: High Availability Cluster Solutions (<https://launchpad.support.sap.com/#/notes/1552925>)
- 1619879 - licensing of HA/DR/cluster → invalid license (<https://launchpad.support.sap.com/#/notes/1619879>)
- 1763512 - Support details for SUSE Linux Enterprise for SAP Applications (<https://launchpad.support.sap.com/#/notes/1763512>)
- 1984787 - SUSE LINUX Enterprise Server 12: Installation notes (<https://launchpad.support.sap.com/#/notes/1984787>)
- 2077934 - Rolling kernel switch in HA environments (<https://launchpad.support.sap.com/#/notes/2077934>)
- 2235581 - SAP HANA: Supported Operating Systems (<https://launchpad.support.sap.com/#/notes/2235581>)
- 2254173 - Linux: Rolling Kernel Switch in Pacemaker based NetWeaver HA environments (<https://launchpad.support.sap.com/#/notes/2254173>)
- 2308598 - Error "Ifconfig is obsolete" happens when you run `startsap` to start SAP system (<https://launchpad.support.sap.com/#/notes/2308598>)
- 2369910 - SAP Software on Linux: General information (<https://launchpad.support.sap.com/#/notes/2369910>)

- 2464065 - Check of automatic maintenance mode for HA solutions (<https://launchpad.support.sap.com/#/notes/2464065>)
- 2578899 - SUSE Linux Enterprise Server 15: Installation Note (<https://launchpad.support.sap.com/#/notes/2578899>)
- 2630416 - Support for Standalone Enqueue Server 2 (<https://launchpad.support.sap.com/#/notes/2630416>)
- 2641019 - Installation of ENSA2 and update from ENSA1 to ENSA2 in SUSE HA environment (<https://launchpad.support.sap.com/#/notes/2641019>)
- 2684254 - SAP HANA DB: Recommended OS settings for SLES 15 / SLES for SAP Applications 15 (<https://launchpad.support.sap.com/#/notes/2684254>)
- 2711036 - Usage of the Standalone Enqueue Server 2 in an HA Environment (<https://launchpad.support.sap.com/#/notes/2711036>)
- 2855499 - FAIL: RKS Warning(s): Unsupported SCS instance with additional gateway found (<https://launchpad.support.sap.com/#/notes/2855499>)
- 3075829 - New sapstarstrv Web service method "ABAPSetServerInactive" (<https://launchpad.support.sap.com/#/notes/3075829>)
- 3091152 - sapstartsrv - improved deregistration for UNIX/Linux (<https://launchpad.support.sap.com/#/notes/3091152>)
- 3115889 - SAP Web Dispatcher embedded deployment in an ASCS/SCS instance (<https://launchpad.support.sap.com/#/notes/3115889>)
- 3139184 - Linux: systemd integration for sapstartsrv and SAP Hostagent (<https://launchpad.support.sap.com/#/notes/3139184>)
- 3145200 - SAP Host Agent 7.22 PL57 (<https://launchpad.support.sap.com/#/notes/3145200>)

9 Appendix

9.1 CRM configuration

Find below the complete crm configuration for SAP system EN2. This example is for the two node cluster, but without priority fencing. In multi node cluster you will find additional node entries like **node 3: valuga13**.

```
node 1: valuga11
node 2: valuga12
primitive rsc_fs_EN2_ASCS00 Filesystem \
  params device="/dev/disk/by-id/SUSE-Example-A-part2" \
  directory="/usr/sap/EN2/ASC00" \
  fstype=xfsp \
  op start timeout=60s interval=0 \
  op stop timeout=60s interval=0 \
  op monitor interval=20s timeout=40s
primitive rsc_fs_EN2_ERS10 Filesystem \
  params device="/dev/disk/by-id/SUSE-Example-A-part3" \
  directory="/usr/sap/EN2/ERS10" \
  fstype=xfsp \
  op start timeout=60s interval=0 \
  op stop timeout=60s interval=0 \
  op monitor interval=20s timeout=40s
primitive rsc_ip_EN2_ASCS00 IPAddr2 \
  params ip=192.168.1.112 \
  op monitor interval=10s timeout=20s
primitive rsc_ip_EN2_ERS10 IPAddr2 \
  params ip=192.168.1.113 \
  op monitor interval=10s timeout=20s
primitive rsc_sap_EN2_ASCS00 SAPIInstance \
  op monitor interval=11 timeout=60 on-fail=restart \
  params InstanceName=EN2_ASCS00_sapen2as \
  START_PROFILE="/sapmnt/EN2/profile/EN2_ASCS00_sapen2as" \
  AUTOMATIC_RECOVER=false \
  meta resource-stickiness=5000
primitive rsc_sap_EN2_ERS10 SAPIInstance \
  op monitor interval=11 timeout=60 on-fail=restart \
  params InstanceName=EN2_ERS10_sapen2er \
  START_PROFILE="/sapmnt/EN2/profile/EN2_ERS10_sapen2er" \
  AUTOMATIC_RECOVER=false IS_ERS=true
primitive stonith-sbd stonith:external/sbd \
  params pcmk_delay_max=30s
group grp_EN2_ASCS00 rsc_ip_EN2_ASCS00 rsc_fs_EN2_ASCS00 \
```

```

rsc_sap_EN2_ASCS00 \
meta resource-stickiness=3000
group grp_EN2_ERS10 rsc_ip_EN2_ERS10 rsc_fs_EN2_ERS10 \
rsc_sap_EN2_ERS10
colocation col_sap_EN2_not_both -5000: grp_EN2_ERS10 grp_EN2_ASCS00
order ord_sap_EN2_first_ascs Optional: rsc_sap_EN2_ASCS00:start \
rsc_sap_EN2_ERS10:stop symmetrical=false
property cib-bootstrap-options: \
have-watchdog=true \
cluster-infrastructure=corosync \
cluster-name=hacluster \
stonith-enabled=true \
placement-strategy=balanced
rsc_defaults rsc-options: \
resource-stickiness=1 \
migration-threshold=3
op_defaults op-options: \
timeout=600 \
record-pending=true

```

9.2 Corosync configuration of the two-node cluster

Find below the corosync configuration for one corosync ring. Ideally two rings would be used.

```

valugall:~ # cat /etc/corosync/corosync.conf
# Read the corosync.conf.5 manual page
totem {
    version: 2
    secauth: on
    crypto_hash: sha1
    crypto_cipher: aes256
    cluster_name: hacluster
    clear_node_high_bit: yes
    token: 5000
    token_retransmits_before_loss_const: 10
    join: 60
    consensus: 6000
    max_messages: 20
    interface {
        ringnumber: 0
        mcastport: 5405
        ttl: 1
    }
}

transport: udpu

```



```

}

logging {
    fileline: off
    to_stderr: no
    to_logfile: no
    logfile: /var/log/cluster/corosync.log
    to_syslog: yes
    debug: off
    timestamp: on
    logger_subsys {
        subsys: QUORUM
        debug: off
    }
}

nodelist {
    node {
        ring0_addr: 192.168.1.103
        nodeid: 1
    }

    node {
        ring0_addr: 192.168.1.104
        nodeid: 2
    }
}

quorum {

    # Enable and configure quorum subsystem (default: off)
    # see also corosync.conf.5 and votequorum.5
    provider: corosync_votequorum
    expected_votes: 2
    two_node: 1
}

```

9.3 Corosync configuration of the multi-node cluster

Find below the corosync configuration for one corosync ring. Ideally two rings would be used.

```

# Please read the corosync.conf.5 manual page
totem {

```

```

version: 2
secauth: on
crypto_hash: sha1
crypto_cipher: aes256
cluster_name: hacluster
clear_node_high_bit: yes
token: 5000
token_retransmits_before_loss_const: 10
join: 60
consensus: 6000
max_messages: 20
interface {
  ringnumber: 0
  mcastport: 5405
  ttl: 1
}

transport: udpu
}

logging {
  fileline: off
  to_stderr: no
  to_logfile: no
  logfile: /var/log/cluster/corosync.log
  to_syslog: yes
  debug: off
  timestamp: on
  logger_subsys {
    subsys: QUORUM
    debug: off
  }
}

nodelist {
  node {
    ring0_addr: 192.168.1.103
    nodeid: 1
  }

  node {
    ring0_addr: 192.168.1.104
    nodeid: 2
  }

  node {

```

```
ring0_addr: 192.168.1.105
nodeid: 3
}

}

quorum {

# Enable and configure quorum subsystem (default: off)
# see also corosync.conf.5 and votequorum.5
provider: corosync_votequorum
expected_votes: 3
two_node: 0
}
```

9.4 /usr/sap/sapservices without native systemd integration

```
#!/bin/sh
LD_LIBRARY_PATH=/usr/sap/EN2/ASCS00/exe:$LD_LIBRARY_PATH; export LD_LIBRARY_PATH; /usr/
sap/EN2/ASCS00/exe/sapstartsrv pf=/usr/sap/EN2/SYS/profile/EN2_ASCS00_sapen2as -D -u
en2adm
LD_LIBRARY_PATH=/usr/sap/EN2/ERS10/exe:$LD_LIBRARY_PATH; export LD_LIBRARY_PATH; /usr/
sap/EN2/ERS10/exe/sapstartsrv pf=/usr/sap/EN2/ERS10/profile/EN2_ERS10_sapen2er -D -u
en2adm
```

10 Legal notice

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

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

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

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

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

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

11 GNU Free Documentation License

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

0. PREAMBLE

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

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

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

1. APPLICABILITY AND DEFINITIONS

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

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

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

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

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

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

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

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

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

2. VERBATIM COPYING

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

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

3. COPYING IN QUANTITY

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

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

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

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

4. MODIFICATIONS

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

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

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

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

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

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

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

5. COMBINING DOCUMENTS

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

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

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

6. COLLECTIONS OF DOCUMENTS

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

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

7. AGGREGATION WITH INDEPENDENT WORKS

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

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

8. TRANSLATION

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

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

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

9. TERMINATION

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

10. FUTURE REVISIONS OF THIS LICENSE

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

ADDENDUM: How to use this License for your documents

Copyright (c) YEAR YOUR NAME.

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

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

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

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

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

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