

SAP NetWeaver Enqueue Replication 1 High Availability Cluster - SAP NetWeaver 7.40 and 7.50

Setup Guide

SUSE Linux Enterprise Server for SAP Applications 15
SAP NetWeaver 7.40 and 7.50

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

SAP NetWeaver Enqueue Replication 1 High Availability Cluster - SAP NetWeaver 7.40 and 7.50

Setup Guide

Date: 2025-03-13

SUSE® Linux Enterprise Server for SAP Applications is optimized in various ways for SAP* applications. This document explains how to deploy an SAP NetWeaver Enqueue Replication 1 High Availability Cluster solution. It is based on SUSE Linux Enterprise Server for SAP Applications 15 and related service packs.

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 18
- 5 Implementing the cluster 26
- 6 Administration 32
- 7 References 36
- 8 Appendix 40
- 9 Legal notice 46
- 10 GNU Free Documentation License 47

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 (HA). 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* S/4 HANA and for SAP* SAP NetWeaver.

SAP NetWeaver is a common stack of middleware functionality used to support the SAP business applications. The SAP Enqueue Replication Server constitutes application level redundancy for one of the most crucial components of the SAP NetWeaver 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 by SUSE Linux Enterprise Server for SAP Applications. The described concept has proven its maturity over several years of productive operations 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 obsoleted by the so-called simple-mount setup. The setup with cluster-controlled file systems is still supported for existing clusters. For deploying new HA clusters, please use the simple-mount setup, described in SUSE TID "Use of Filesystem resource for ASCS/ERS HA setup not possible" (<https://www.suse.com/support/kb/doc/?id=000019944>) and in the "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

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

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

This guide and other SAP-specific best practices documents can be downloaded from the documentation portal at <https://documentation.suse.com/sbp/sap> .

Here you can find guides for SAP HANA system replication automation and HA scenarios for SAP NetWeaver and SAP S/4 HANA.

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 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

This guide details how to:

- Plan a SUSE Linux Enterprise High Availability platform for SAP NetWeaver, including SAP Enqueue Replication Server.
- Set up a Linux high availability platform and perform a basic SAP NetWeaver installation including SAP Enqueue Replication Server on SUSE Linux Enterprise.
- Integrate the high availability cluster with the SAP control framework via `sap-suse-cluster-connector`, as certified by SAP.

This guide focuses on the high availability of the central services.

This guide does **not** address platform specific details. The solution described here is suitable for on-premises deployments and for public clouds. However, for details on deploying the solution in public clouds, refer also to the respective cloud provider documentation.

For SAP HANA system replication, follow the guides for the performance- or cost-optimized scenario.



Note

The HA setup described here is based on cluster-controlled file systems for the SAP central services working directories. This setup with cluster-controlled file system resources has been obsoleted by the so-called simple-mount setup. The setup with cluster-controlled file systems is still supported for existing clusters. For deploying new HA clusters, please use the simple-mount setup, described in SUSE TID "Use of Filesystem resource for ASCS/ERS HA setup not possible" (<https://www.suse.com/support/kb/doc/?id=000019944>) and in the "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).

3 Overview

This guide describes how to set up a pacemaker cluster using SUSE Linux Enterprise Server for SAP Applications 15 for the Enqueue Replication scenario. The goal is to match the SAP NW-HA-CLU 7.40 certification specifications and goals.

These goals 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

The updated certification SAP NW-HA-CLU 7.40 has redefined some of the test procedures and described new expectations how the cluster should behave in special conditions. These changes allowed us to improve the cluster architecture and to design it for easier usage and setup.

Shared SAP resources are on a central NFS server.

The SAP instances themselves are installed on a shared disk to allow switching over the file systems for proper functionality. The second need for a shared disk is that we are using the SBD for the cluster fencing mechanism STONITH.

3.1 Differences to previous cluster architectures

- The concept is different to the old stack with the multi-state architecture. With the new certification we switch to a more simple model with primitives. This means we have on one machine the ASCS with its own resources and on the other machine the ERS with its own resources.
- Use of native systemd integration for SAP hostagent and instance's *sapstartsrv*. Refer to 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 for ASCS, ERS, database and additional SAP Instances

This guide describes the installation of a distributed SAP system on three systems. In this setup, only two systems are in the cluster. The database and SAP dialog instances could also be added to the cluster by either adding the third node to the cluster or by installing the database on either of the nodes. However we recommend to install the database on a separate cluster.



Note

The cluster in this guide only manages the SAP instances ASCS and ERS, because of the focus of the SAP NW-HA-CLU 7.40 certification.

If your database is SAP HANA, we recommend to set up the performance optimized system replication scenario using our automation solution SAPHanaSR. The SAPHanaSR automation should be set up in an own two node cluster. The setup is described in a separate best practice document available at <http://documentation.suse.com/sbp/sap>.




FIGURE 1: THREE SYSTEMS FOR THE CERTIFICATION SETUP

CLUSTERED MACHINES

- One machine (hacert01) for ASCS
 - Host name: sapha1as
- One machine (hacert02) for ERS
 - Host name: sapha1er

NON-CLUSTERED MACHINE

- One machine (hacert03) for DB and DI

- Host name: sapha1db
- Host name: sapha1d1
- Host name: sapha1d2

3.3 High availability for the database

Depending on your needs you can also 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.

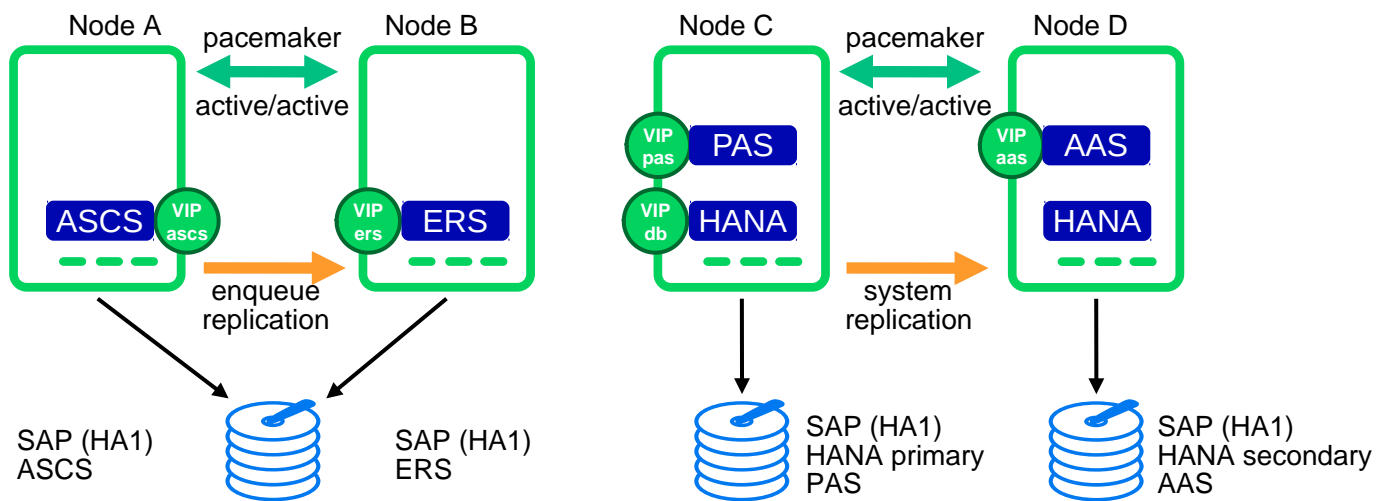


FIGURE 2: ONE CLUSTER FOR CENTRAL SERVICES, ONE FOR SAP HANA SR

TABLE 1: OS / DATABASE COMBINATIONS SUPPORTED WITH THE SYSTEM REPLICATION SCENARIO

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

SUSE Linux Enterprise Server for SAP Applications 15	
SAP HANA DATABASE 2.0	SAP HANA DATABASE 2.0

3.3.2 Simple stack

Another option is to implement a second cluster for a database without system replication, aka "ANYDB". The cluster resource agent SAPInstance uses the SAPHOSTAGENT to control and monitor the database.

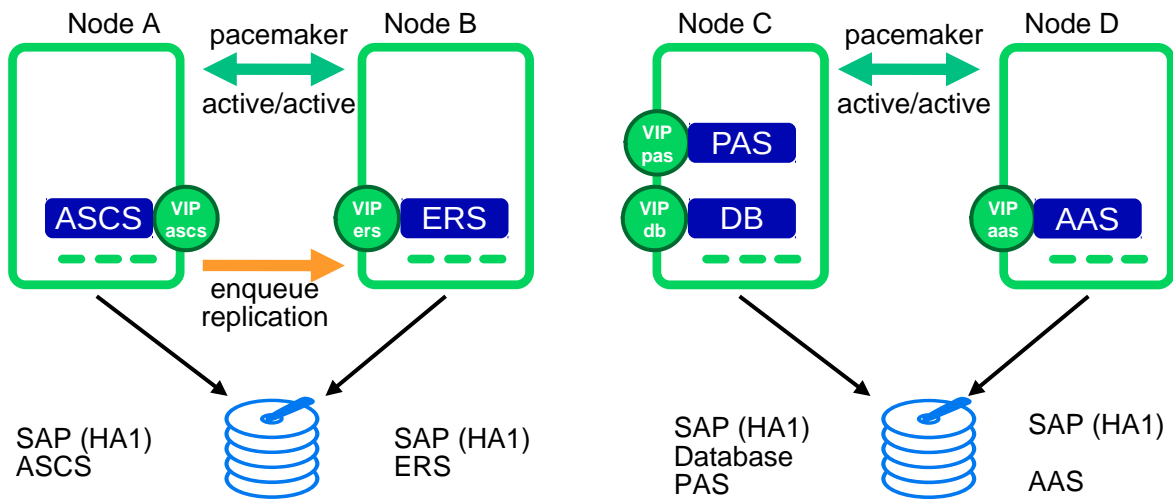


FIGURE 3: ONE CLUSTER FOR THE CENTRAL SERVICES AND ONE CLUSTER FOR THE ANY DATABASE

TABLE 2: OS / DATABASE COMBINATION EXAMPLES FOR THE SIMPLE STACK SCENARIO

SUSE Linux Enterprise Server for SAP Applications 15	
Intel X86_64	IBM PowerLE
SAP HANA DATABASE 2.0	SAP HANA DATABASE 2.0
DB2 FOR LUW 10.5	
MaxDB 7.9	
ORACLE 12.1	



Note

The first version for SAP NetWeaver on IBMPowerLE is 7.50. The first version for SAP HANA on IBM PowerLE is 2.0. More information about supported combinations of OS and databases for SAP NetWeaver can be found at the SAP Product Availability Matrix. (SAP PAM (<https://apps.support.sap.com/sap/support/pam>))

3.4 Integration of SAP NetWeaver 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 `sapstartsrv` controls SAP instances since SAP Kernel versions 6.40. One of the classical 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. It will then 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. This new updated solution enables the central component `sapstartsrv` to report state changes to the cluster software, and therefore avoids the previously described dangerous situations. (See also blog article "Using `sap_vendor_cluster_connector` for interaction between cluster framework and `sapstartsrv`") (<https://blogs.sap.com/2014/05/08/using-sapvendorclusterconnector-for-interaction-between-cluster-framework-and-sapstartsrv/comment-page-1/>).

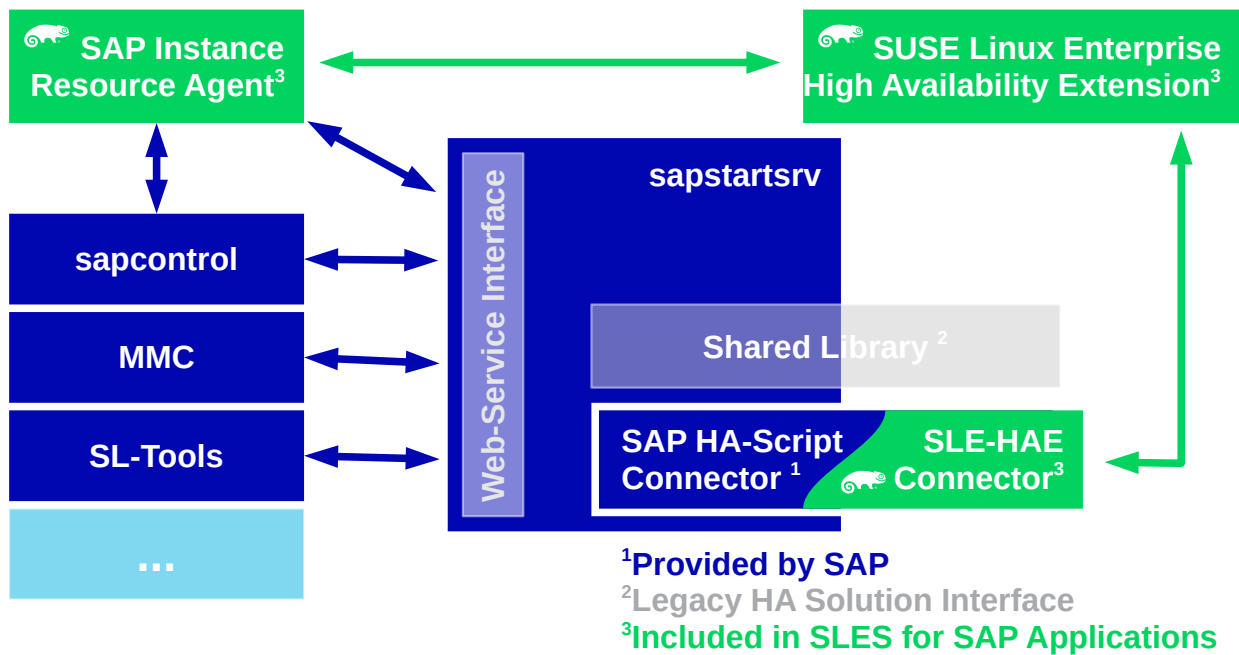


FIGURE 4: CLUSTER CONNECTOR TO INTEGRATE THE CLUSTER WITH THE SAP START FRAMEWORK



Note

For this scenario we are using an updated version of the `sap-suse-cluster-connector`. This version implements the API version 3 for the communication between the cluster framework and the `sapstartsrv`.

The new version of the `sap-suse-cluster-connector` now allows to start, stop and 'move' 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 the SAP management consoles (SAP MMC or SAP MC).

3.5 Disks and partitions

For all SAP file systems beside the file systems on NFS we are using XFS.

3.5.1 Shared disk for cluster ASCS and ERS

The disk for the ASCS and ERS instances need to be shared and assigned to the cluster nodes hacert01 and hacert02. Beside the partitions for the file systems for the SAP instances, the disk also provides the partition to be used as SBD.

On hacert01, prepare the file systems for the shared disk. Create three partitions on the shared drive /dev/sdb:

- partition one (/dev/sdb1) for SBD (7M)
- partition two (/dev/sdb2) for the first file system (10GB) formatted with XFS
- partition three (/dev/sdb3) for the second file system (10GB) formatted with XFS

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

```
# parted -s /dev/sdb print
# # we are on the 'correct' drive, right?
# parted -s /dev/sdb mklabel gpt
# parted -s /dev/sdb mkpart primary 1049k 8388k
# parted -s /dev/sdb mkpart primary 8389k 10.7G
# parted -s /dev/sdb mkpart primary 10.7G 21.5G
# mkfs.xfs /dev/sdb2
# mkfs.xfs /dev/sdb3
```

For these file systems we recommend to use plain 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. This is out of the scope of this setup guide.

After you have partitioned the shared disk on hacert01, you need to request a partition table rescan on hacert02.

```
# partprobe; fdisk -l /dev/sdb
```

During the SAP installation, /usr/sap/HA1/ASCS00 needs to be mounted on hacert01 and /usr/sap/HA1/ERS10 needs to be mounted on hacert02.

- hacert01: /dev/sdb2 /usr/sap/HA1/ASCS00
- hacert02: /dev/sdb3 /usr/sap/HA1/ERS10

3.5.2 Disk for DB and dialog instances (MaxDB example)

The disk for the database and primary application server is assigned to hacert03. In an advanced setup, this disk should be shared between hacert03 and an optional additional node building an own cluster.

- partition one (/dev/sdb1) for SBD (7M) - not used here but a reservation for an optional second cluster
- partition two (/dev/sdb2) for the Database (60GB) formatted with XFS
- partition three (/dev/sdb3) for the second file system (10GB) formatted with XFS
- partition four (/dev/sdb4) for the third file system (10GB) formatted with XFS

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

```
# parted -s /dev/sdb print
# # we are on the 'correct' drive, right?
# parted -s /dev/sdb mklabel gpt
# parted -s /dev/sdb mkpart primary 1049k 8388k
# parted -s /dev/sdb mkpart primary 8389k 60G
# parted -s /dev/sdb mkpart primary 60G 70G
# parted -s /dev/sdb mkpart primary 70G 80G
# mkfs.xfs /dev/sdb2
# mkfs.xfs /dev/sdb3
# mkfs.xfs /dev/sdb4
```

TO BE MOUNTED EITHER BY OS OR AN OPTIONAL CLUSTER

- hacert03: /dev/sdb2 /sapdb
- hacert03: /dev/sdb3 /usr/sap/HA1/DVEBMGS01
- hacert03: /dev/sdb4 /usr/sap/HA1/D02



Note

D01: Since NetWeaver 7.5, the primary application server instance directory has been renamed. (D < Instance_Number >)

NFS SERVER

- `nfs1:/data/nfs/suseEnqReplNW7x/HA1/sapmnt /sapmnt`
- `nfs1:/data/nfs/suseEnqReplNW7x/HA1/usr/sapsys /usr/sap/HA1/SYS`

MEDIA

- `nfs1:/data/SCT/media/SAP-MEDIA/NW74 /sapcd`

or

- `nfs1:/data/SCT/media/SAP-MEDIA/NW75 /sapcd`

3.6 IP addresses and virtual names

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

```
192.168.201.111 hacert01
192.168.201.112 hacert02
192.168.201.113 hacert03
192.168.201.115 saphalas
192.168.201.116 saphaler
192.168.201.117 saphaldb
192.168.201.118 saphald1
192.168.201.119 saphald2
```

3.7 Mount points and NFS shares

In our 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` also contains the SAP control file `sapservices` and the `saphostagent`, the directory should not be placed on a shared file system between the cluster nodes.

Create the directory structure on all nodes which might be able to run the SAP resource. The `SYS` directory will be on an NFS share for all nodes.

- Creating mount points and mounting NFS share at all nodes

EXAMPLE 1: SAP NETWEAVER 7.4

```
# mkdir -p /sapcd
```

```
# mkdir -p /sapmnt
# mkdir -p /usr/sap/HA1/{ASCS00,D02,DVEBMGS01,ERS10,SYs}
# mount -t nfs nfs1:/data/nfs/suseEnqReplNW7x/HA1/sapmnt /sapmnt
# mount -t nfs nfs1:/data/nfs/suseEnqReplNW7x/HA1/ursapsys /usr/sap/HA1/SYS
# mount -t nfs nfs1:/data/SCT/media/SAP-MEDIA/NW74 /sapcd
```

EXAMPLE 2: SAP NETWEAVER 7.5

```
# mkdir -p /sapcd
# mkdir -p /sapmnt
# mkdir -p /usr/sap/HA1/{ASCS00,D01,D02,ERS10,SYs}
# mount -t nfs nfs1:/data/nfs/suseEnqReplNW7x/HA1/sapmnt /sapmnt
# mount -t nfs nfs1:/data/nfs/suseEnqReplNW7x/HA1/ursapsys /usr/sap/HA1/SYS
# mount -t nfs nfs1:/data/SCT/media/SAP-MEDIA/NW75 /sapcd
```

- Only MaxDB: creating mount points for the database at hacert03:

```
# mkdir -p /sapdb
```

- Only HANA: creating mount points for database at hacert03:

```
# mkdir -p /hana/{shared,data,log}
```

- Other databases: creating mount points based on there installation guide.

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




FIGURE 5: FILE SYSTEM LAYOUT INCLUDING NFS SHARES

Prepare the three servers for the distributed SAP installation. Server 1 (hacert01) will be used to install the ASCS SAP instance. Server 2 (hacert02) will be used to install the ERS SAP instance and server 3 (hacert03) will be used to install the dialog SAP instances and the database.

- Mounting the instance and database file systems at one specific node:

EXAMPLE 3: **SAP NETWEAVER 7.40 ON X86_64 ARCHITECTURE WITH MAXDB**

```
(ASCS hacert01) # mount /dev/sdb2 /usr/sap/HA1/ASCS00
(ERS hacert02) # mount /dev/sdb3 /usr/sap/HA1/ERS10
(DB hacert03) # mount /dev/sdb2 /sapdb
(Dialog hacert03) # mount /dev/sdb3 /usr/sap/HA1/DVEBMGS01
(Dialog hacert03) # mount /dev/sdb4 /usr/sap/HA1/D02
```

EXAMPLE 4: **SAP NETWEAVER 7.50 ON POWERLE ARCHITECTURE WITH HANA**

```
(ASCS hacert01) # mount /dev/sdb2 /usr/sap/HA1/ASCS00
(ERS hacert02) # mount /dev/sdb3 /usr/sap/HA1/ERS10
(DB hacert03) # mount /dev/sdc1 /hana/shared
(DB hacert03) # mount /dev/sdc2 /hana/log
(DB hacert03) # mount /dev/sdc3 /hana/data
(Dialog hacert03) # mount /dev/sdb3 /usr/sap/HA1/D01
(Dialog hacert03) # mount /dev/sdb4 /usr/sap/HA1/D02
```

- As a result, the directory `/usr/sap/HA1/` should now look like:

```
# ls -la /usr/sap/HA1/
total 0
drwxr-xr-x 1 haladm sapsys 70 28. Mar 17:26 ./
drwxr-xr-x 1 root sapsys 58 28. Mar 16:49 ../
drwxr-xr-x 7 haladm sapsys 58 28. Mar 16:49 ASCS00/
drwxr-xr-x 1 haladm sapsys 0 28. Mar 15:59 D02/
drwxr-xr-x 1 haladm sapsys 0 28. Mar 15:59 D01/
drwxr-xr-x 1 haladm sapsys 0 28. Mar 15:59 ERS10/
drwxr-xr-x 5 haladm 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 is:

- Installing the ASCS instance for the central services
- Installing the ERS to get a replicated enqueue scenario
- preparing the ASCS and ERS installations for the cluster take-over
- Installing the Database
- Installing the primary application server instance (PAS)
- Installing additional application server instances (AAS)

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




FIGURE 6: 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 which is only an example.

Group	sapinst	1000
Group	sapsys	1001
Group	sapadm	3000
Group	sdba	3002
User	haladm	3000
User	sdb	3002
User	sqdha1	3003
User	sapadm	3004

4.2 Installing ASCS on hacert01

Temporarily, you need to set the service IP address used later in the cluster as local IP, because the installer wants to resolve or use it. Make sure to use the right virtual host name for each installation step. Take care for file systems like `/dev/sdb2` and `/sapcd/` which might also need to be mounted.

```
# ip a a 192.168.201.115/24 dev eth0
# mount /dev/sdb2 /usr/sap/HA1/ASCS00
# cd /sapcd/SWPM/
# ./sapinst SAPINST_USE_HOSTNAME=sapha1as
```

- SWPM option depends on SAP NetWeaver version and architecture
 - Installing SAP NetWeaver 7.40 SR2 → MaxDB → SAP-Systems → Application Server ABAP → High-Availability System → ASCS Instance
 - Installing SAP NetWeaver 7.5 → SAP HANA Database → Installation → Application Server ABAP → High-Availability System → ASCS Instance
- SID id HA1
- Use instance number 00
- Deselect using FQDN
- All passwords: use `<yourSecurePwd>`
- Double-check during the parameter review if virtual name **sapha1as** is used

4.3 Installing ERS on hacert02

Temporarily, you need to set the service IP address used later in the cluster as local IP because the installer wants to resolve or use it. Make sure to use the right virtual host name for each installation step.

```
# ip a a 192.168.201.116/24 dev eth0
# mount /dev/sdb3 /usr/sap/HA1/ERS10
# cd /sapcd/SWPM/
```

```
# ./sapinst SAPINST_USE_HOSTNAME=saphaler
```

- SWPM option depends on SAP NetWeaver version and architecture
 - Installing SAP NetWeaver 7.40 SR2 → MaxDB → SAP-Systems → Application Server ABAP → High-Availability System → Enqueue Replication Server Instance
 - Installing SAP NetWeaver 7.5 → SAP HANA Database → Installation → Application Server ABAP → High-Availability System → Enqueue Replication Server Instance
- Use instance number 10
- Deselect using FQDN
- Double-check during the parameter review if virtual name **saphaler** is used
- If you get an error during the installation about permissions, change the ownership of the ERS directory

```
# chown -R haladm:sapsys /usr/sap/HA1/ERS10
```

- If you get a prompt to manually stop/start the ASCS instance, log in at hacert01 as user haladm and call sapcontrol.

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

4.4 Poststeps for ASCS and ERS

4.4.1 Stopping ASCS and ERS

On hacert01

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

On hacert02

```
# su - haladm  
# 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 SAPHA1_00.service
# systemctl stop SAPHA1_00.service
# systemctl disable SAPHA1_10.service
# systemctl stop SAPHA1_10.service
```



Note

Stopping these 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
SAPHA1_10.service disabled
SAPHA1_10.service disabled
```

The instance services are now 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 SAPHA1_00 # sapstartsrv pf=/usr/sap/HA1/SYS/profile/
HA1_ASCS00_saphalas
systemctl --no-ask-password start SAPHA1_10 # sapstartsrv pf=/usr/sap/HA1/SYS/profile/
HA1_ERS10_saphaler
```

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 our repositories:

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



Note

Be careful as there are two packages available. The package `sap_suse_cluster_connector` continues to contain the old version 1.1.0 (SAP API 1). The package `sap-suse-cluster-connector` contains the new version 3.1.x (SAP API 3). The package `sap-suse-cluster-connector` with version 3.1.x implements the SUSE SAP API version 3. New features like SAP Rolling Kernel Switch (RKS) and the move of ASCS are only supported with this new version.

For the ERS and ASCS instance edit the instance profiles `HA1_ASCS00_sapha1as` and `HA1_ERS10_sapha1er` in the profile directory `/usr/sap/HA1/SYS/profile/`.

You need to tell `sapstartsrv` to load the HA script connector library and to use the `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 `haladm` to the unix user group `haclient`.

```
# usermod -aG haclient haladm
```

4.4.4 Adapting SAP profiles to match the SAP NW-HA-CLU 7.40 certification

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

EXAMPLE 5: FILE `/USR/SAP/HA1/SYS/PROFILE/HA1_ASCS00_SAPHA1AS`

```
Start_Program_01 = local $_EN pf=$_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, change the start command from *Restart_Program_xx* to *Start_Program_xx* for the enqueue replication server (enrepsvr).

EXAMPLE 6: FILE /USR/SAP/HA1/SYS/PROFILE/HA1_ERS10_SAPHA1ER

```
Start_Program_00 = local $_ER pf=$_PFL NR=$(SCSID)
```

4.4.5 Starting ASCS and ERS

On hacert01

```
# su - haladm  
# sapcontrol -nr 00 -function StartService HA1  
# sapcontrol -nr 00 -function Start
```

On hacert02

```
# su - haladm  
# sapcontrol -nr 10 -function StartService HA1  
# sapcontrol -nr 10 -function Start
```

4.5 Installing DB on hacert03 (example MaxDB)

The MaxDB needs minimum 40 GB. Use /dev/sdb2 and mount the partition to /sapdb.

```
# ip a a 192.168.201.117/24 dev eth0  
# mount /dev/sdb2 /sapdb  
# cd /sapcd/SWPM/  
# ./sapinst SAPINST_USE_HOSTNAME=sapha1db
```

- Install SAP NetWeaver 7.40 SR2 → MaxDB → SAP-Systems → Application Server ABAP → High Availability System → DB
- Profile directory is /sapmnt/HA1/profile
- DB ID is HA1
- Volume Media Type **keep** File (not raw)
- Deselect using FQDN
- Double-check during the parameter review if virtual name **sapha1db** is used

4.6 Installing DB on hacert03 (example SAP HANA)

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 192.168.201.117/24 dev eth0
# mount /dev/sdc1 /hana/shared
# mount /dev/sdc2 /hana/log
# mount /dev/sdc3 /hana/data
# cd /sapcd/SWPM/
# ./sapinst SAPINST_USE_HOSTNAME=sapha1db
```

- Install SAP NetWeaver 7.5 → SAP HANA Database → Installation → Application Server ABAP → High-Availability System → Database Instance
- Profile directory is */sapmnt/HA1/profile*
- Deselect using FQDN
- Database parameters: enter DBSID is H04; Database Host is sapha1db; Instance Number is 00
- Database System ID: enter Instance Number is 00; SAP Mount Directory is /hana/shared
- Account parameters: change them in case custom values are needed
- Clean up: select **Yes**, remove operating system users from group'sapinst'....
- Double-check during the parameter review if virtual name **sapha1db** is used

4.7 Installing the primary application server (PAS) on hacert03

```
# ip a a 192.168.201.118/24 dev eth0
# mount /dev/sdb3 /usr/sap/HA1/DVEBMGS01
# cd /sapcd/SWPM/
# ./sapinst SAPINST_USE_HOSTNAME=sapha1d1
```

or alternatively:

```
# ip a a 192.168.201.118/24 dev eth0
# mount /dev/sdb3 /usr/sap/HA1/D01
# cd /sapcd/SWPM/
# ./sapinst SAPINST_USE_HOSTNAME=sapha1d1
```


- SWPM option depends on SAP NetWeaver version and architecture
 - Installing SAP NetWeaver 7.40 SR2 → MaxDB → SAP-Systems → Application Server ABAP → High-Availability System → Primary Application Server Instance (PAS)
 - Installing SAP NetWeaver 7.5 → SAP HANA Database → Installation → Application Server ABAP → High-Availability System → Primary Application Server Instance (PAS)
- Use instance number 01
- Deselect using FQDN
- For our hands-on setup use a default secure store key
- Do not install Diagnostic Agent
- No SLD
- Double-check during the parameter review if virtual name **sapha1d1** is used

4.8 Installing an additional application server (AAS) on hacert03

```
# ip a a 192.168.201.119/24 dev eth0
# mount /dev/sdb4 /usr/sap/HA1/D02
# cd /sapcd/SWPM/
# ./sapinst SAPINST_USE_HOSTNAME=sapha1d2
```

- SWPM option depends on SAP NetWeaver version and architecture
 - Installing SAP NetWeaver 7.40 SR2 → MaxDB → SAP-Systems → Application Server ABAP → High-Availability System → Additional Application Server Instance (AAS)
 - Installing SAP NetWeaver 7.5 → SAP HANA Database → Installation → Application Server ABAP → High-Availability System → Additional Application Server Instance (AAS)
- Use instance number 02
- Deselect using FQDN
- Do not install Diagnostic Agent
- Double-check during the parameter review if virtual name **sapha1d2** is used

5 Implementing the cluster

The main procedure to implement the cluster is as follows:

- Install the cluster software if not already done during the installation of the operating system
- Configure the cluster communication framework corosync.
- Configure the cluster resource manager.
- Configure the cluster resources.
- Tune the cluster timing in special for the SBD.



Note

Before you continue to set up the cluster, first stop all SAP instances, remove the (manually added) IP addresses on the cluster nodes and unmount the file systems which will be controlled by the cluster later.



Note

The SBD device/partition needs to be created beforehand. In this setup guide, partition `/dev/sdb1` is already reserved for SBD usage.

TASKS

1. Setup Chrony (best with YaST) and enable it
2. Install pattern `ha_sles` on both cluster nodes

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

5.1 Configuring the cluster base

TASKS

- Install and configure the cluster stack at the first machine

You can use either YaST or the interactive command line tool `ha-cluster-init` to configure the cluster base. The following script can be used for automated setups.

```
# modprobe softdog
# echo "softdog" > /etc/modules-load.d/softdog.conf
# systemctl enable sbd
# ha-cluster-init -y -i eth0 -u -s /dev/sdb1
```

Keep in mind that a hardware watchdog is preferred instead of the softdog method.

- Join the second node

Find below some preparation steps on the second node.

```
# modprobe softdog
# echo "softdog" > /etc/modules-load.d/softdog.conf
# systemctl enable sbd
# rsync 192.168.201.111:/etc/sysconfig/sbd /etc/sysconfig
```

You can use either YaST to configure the cluster base or the interactive command line tool [ha-cluster-join](#). The following script can be used for automated setups.

```
# ha-cluster-join -y -c 192.168.201.111 -i eth0
```

- The `crm_mon -lr` output should look like this:

```
Stack: corosync
Current DC: hacert01 (version 1.1.18+20180430.b12c320f5-1.14-b12c320f5) - partition with
quorum
Last updated: Wed Apr  3 13:53:40 2019
Last change: Wed Apr  3 13:44:40 2019 by root via cibadmin on hacert01

2 nodes configured
1 resource configured

Online: [ hacert01 hacert02 ]

Full list of resources:

stonith-sbd    (stonith:external/sbd): Started hacert01
```

- After both nodes are listed in the overview, verify the property setting of the basic cluster configuration. Very important here is the setting: **record-pending = true**.

```
# crm configure show
...
property cib-bootstrap-options: \
    have-watchdog=true \
    dc-version=1.1.18+20180430.b12c320f5-1.14-b12c320f5 \
```

```

cluster-infrastructure=corosync \
cluster-name=hacluster \
stonith-enabled=true \
last-lrm-refresh=1494346532
rsc_defaults rsc-options: \
resource-stickiness=1 \
migration-threshold=3
op_defaults op-options: \
timeout=600 \
record-pending=true

```

5.2 Configuring cluster resources

You need a changed SAPInstance resource agent for SAP NetWeaver to **not** use the multi-state construct but to move to a more cluster-like construct, to start and stop the ASCS and the ERS itself and **not** the complete multi-state construct.

For this there is a new functionality for the ASCS needed to follow the ERS. The ASCS needs to mount the shared memory table of the ERS to avoid the loss of locks.

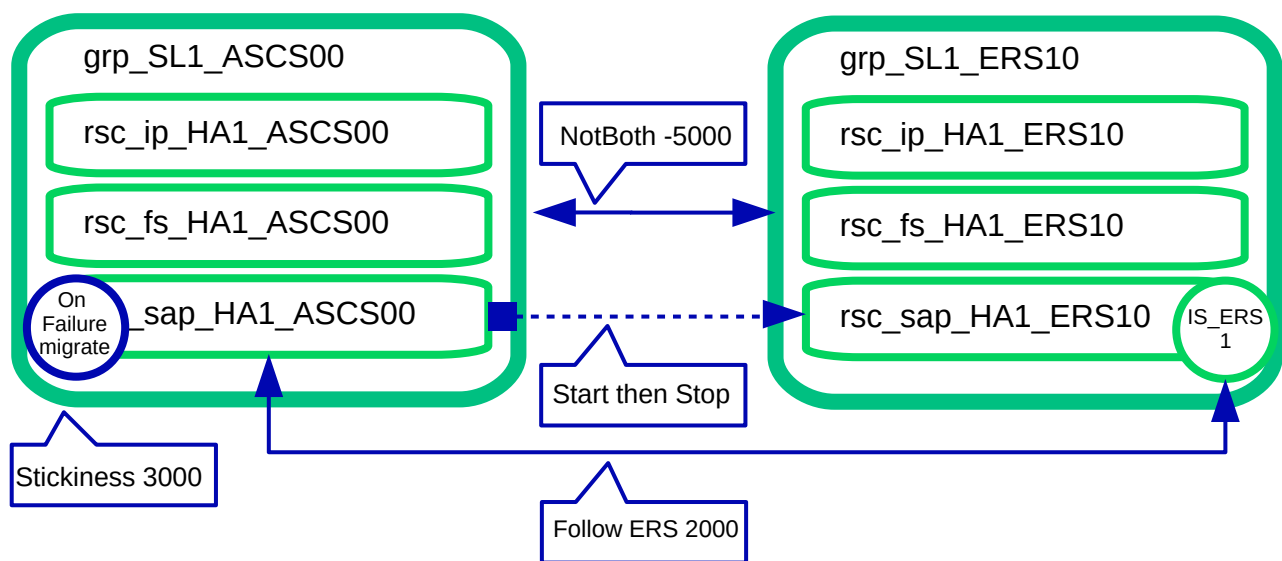


FIGURE 7: RESOURCES AND CONSTRAINTS

The implementation is done using the new flag "runs_ers_\$SID" within the RA, enabled using the resource parameter "IS_ERS = TRUE".

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

5.2.1 Preparing the cluster for adding the resources

To avoid that the cluster starts partially defined resources, set the cluster to the maintenance mode. This deactivates all monitor actions.

As user root

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

5.2.2 Configuring the resources for the ASCS

First, configure the resources for the file system, IP address and the SAP instance. Make sure you adapt the parameters to your environment. The shown file system and SAPIInstance 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 SAPIInstance 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_SAPIInstance(7)` and `nfs(5)`.

EXAMPLE 7: ASCS PRIMITIVE

```
primitive rsc_fs_HA1_ASCS00 Filesystem \  
  params device="/dev/sdb2" directory="/usr/sap/HA1/ASC00" \  
  fstype=xfse \  
  op start timeout=60s interval=0 \  
  op stop timeout=60s interval=0 \  
  op monitor interval=20s timeout=40s \  
primitive rsc_ip_HA1_ASCS00 IPAddr2 \  
  params ip=192.168.201.115 \  
  op monitor interval=10s timeout=20s \  
primitive rsc_sap_HA1_ASCS00 SAPIInstance \  
  op monitor interval=11 timeout=60 on-fail=restart \  
  params InstanceName=HA1_ASCS00_saphalas \  
  START_PROFILE="/sapmnt/HA1/profile/HA1_ASCS00_saphalas" \  
  AUTOMATIC_RECOVER=false \  
  meta resource-stickiness=5000 failure-timeout=60 \  
  migration-threshold=1 priority=10
```

EXAMPLE 8: ASCS GROUP

```
group grp_HA1_ASCS00 \  
  rsc_ip_HA1_ASCS00 rsc_fs_HA1_ASCS00 rsc_sap_HA1_ASCS00 \  
  meta resource-stickiness=3000
```

Create a txt file (like *crm_ascs.txt*) with your preferred text editor, enter 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_ascs.txt
```

5.2.3 Configuring the resources for the ERS

Second, configure the resources for the file system, IP address and the SAP instance. Make sure you adapt the parameters to your 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_FileSystem(7)`, `ocf_heartbeat_SAPInstance(7)` and `nfs(5)`.

The specific parameter `IS_ERS=true` should only be set for the ERS instance.

EXAMPLE 9: ERS PRIMITIVE

```
primitive rsc_fs_HA1_ERS10 Filesystem \  
  params device="/dev/sdb3" directory="/usr/sap/HA1/ERS10" fstype=xfst \  
  op start timeout=60s interval=0 \  
  op stop timeout=60s interval=0 \  
  op monitor interval=20s timeout=40s \  
primitive rsc_ip_HA1_ERS10 IPAddr2 \  
  params ip=192.168.201.116 \  
  op monitor interval=10s timeout=20s \  
primitive rsc_sap_HA1_ERS10 SAPInstance \  
  op monitor interval=11 timeout=60 on-fail=restart \  
  params InstanceName=HA1_ERS10_saphaler \  
    START_PROFILE="/sapmnt/HA1/profile/HA1_ERS10_saphaler" \  
    AUTOMATIC_RECOVER=false IS_ERS=true \  
  meta priority=1000
```

EXAMPLE 10: ERS GROUP

```
group grp_HA1_ERS10 \  
  rsc_ip_HA1_ERS10 rsc_fs_HA1_ERS10 rsc_sap_HA1_ERS10
```

Create a txt file (like *crm_ers.txt*) with your preferred text editor, enter 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
```

5.2.4 Configuring the colocation constraints between ASCS and ERS

The constraints between the ASCS and ERS instance are needed to define that the ASCS instance starts exactly on the cluster node running the ERS instance after a failure (`loc_sap_HA1_failover_to_ers`). This constraint is needed to ensure that the locks are not lost after an ASCS instance (or node) failure.

If the ASCS instance has been started by the cluster, the ERS instance should be moved to an "other" cluster node (`col_sap_HA1_no_both`). This constraint is needed to ensure that the ERS will synchronize the locks again and the cluster is ready for an additional take-over.

EXAMPLE 11: LOCATION CONSTRAINT

```
colocation col_sap_HA1_no_both -5000: grp_HA1_ERS10 grp_HA1_ASCS00
location loc_sap_HA1_failover_to_ers rsc_sap_HA1_ASCS00 \
  rule 2000: runs_ers_HA1 eq 1
order ord_sap_HA1_first_start_ascs Optional: rsc_sap_HA1_ASCS00:start \
  rsc_sap_HA1_ERS10:stop symmetrical=false
```

Create a txt file (like `crm_col.txt`) with your preferred text editor, enter all three constraints to that file and load the configuration to the cluster manager configuration.

As user root

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

5.2.5 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

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

6 Administration

6.1 Dos and don'ts

6.1.1 Never stop the ASCS instance

For normal operation, **do not stop** the ASCS SAP instance with any tool such as cluster tools or SAP tools. The stop of the ASCS instance might lead to a loss of enqueue locks. Because following the new SAP NW-HA-CLU 7.40 certification the cluster must allow local restarts of the ASCS. This feature is needed to allow rolling kernel switch (RKS) updates without reconfiguring the cluster.



Warning

Stopping the ASCS instance might lead to the loss of SAP enqueue locks during the start of the ASCS on the same node.

6.1.2 How to move ASCS

To **move** the ASCS SAP instance, you should use the SAP tools such as the SAP management console. This will trigger `sapstartsrv` to use `sap-suse-cluster-connector` to move the ASCS instance. As user `ha1adm` you might call the following command to move away the ASCS. The move-away will always move the ASCS to the ERS side which will keep the SAP enqueue locks.

As ha1adm

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

6.1.3 Never block resources

With SAP NW-HA-CLU 7.40 it is **not longer allowed to block resources** from being controlled manually. This using the variable `BLOCK_RESOURCES` in `/etc/sysconfig/sap_suse_cluster_connector` is not allowed anymore.

6.1.4 Always use unique instance numbers

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

6.1.5 How to set cluster into maintenance mode

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

As user root

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

As user ha1adm (the full path is needed)

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

6.1.6 Procedure to end the cluster maintenance

As user root

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

6.1.7 Cleaning up resources

Next is how to **clean up resource failures**. Failures of the ASCS will automatically be 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

```
# crm resource refresh RESOURCE-NAME
```



Warning

You should not clean up the complete group of the ASCS resource as this might lead into an unwanted cluster action to take-over the complete group to the node where ERS instance is running.

6.2 Testing the cluster

We strongly recommend that you at least process the following tests before you plan going 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 SAP management console. This test checks the status of the SAP NetWeaver cluster integration.

As user ha1adm

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

6.2.2 Starting SAP checks using *HACheckConfig* and *HACheckFailoverConfig*

Check if the HA configuration tests are showing no errors.

As user ha1adm

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

6.2.3 Manually moving ASCS

Check if manually moving the ASCS using HA tools works properly.

As user root

```
# crm resource move rsc_sap_HA1_ASCS00 force  
## wait until the ASCS is been moved to the ERS host  
# crm resource clear rsc_sap_HA1_ASCS00
```

6.2.4 Migrating ASCS using *HAFailoverToNode*

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

As user ha1adm

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

6.2.5 Testing ASCS move after failure

Check if the ASCS instance moves correctly after a node failure.

As user root

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

6.2.6 Inplacing restart of ASCS using stop and start

Check if the inplace 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.



Warning

This test will force the SAP system to **lose** the enqueue locks. **This test should not be processed during production.**

As user ha1adm

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

6.2.7 Additionally recommended tests

- Automated restart of the ASCS (simulating RKS)
- 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
- Check the outage and restart of sapstartsrv
- Check the rolling kernel switch procedure (RKS), if possible
- Check the simulation of an upgrade
- Check the simulation of cluster resource failures

7 References

For more information, see the documents listed below.

7.1 Pacemaker

- Pacemaker 1.1 Configuration Explained: https://clusterlabs.org/pacemaker/doc/en-US/Pacemaker/1.1/html/Pacemaker_Explained/ ↗
- Pacemaker 2.0 Configuration Explained: https://clusterlabs.org/pacemaker/doc/en-US/Pacemaker/2.0/html-single/Pacemaker_Explained/index.html ↗

7.2 Related Manual Pages

- corosync.conf(8)
- corosync_overview(8)
- crm(8)
- crm_mon(8)
- crm_simulate(8)
- cs_clusterstate(8)
- cs_man2pdf(8)
- mount.nfs(8)
- ocf_heartbeat_IPaddr2(7)
- ocf_heartbeat_SAPInstance(7)
- ocf_suse_SAPStartSrv(7)
- sapping(7)
- sapservices-move(8)
- SAPStartSrv_basic_cluster(7)
- sap_suse_cluster_connector(8)

- `sbd(8)`
- `stonith_sbd(7)`
- `systemctl(8)`
- `usermod(8)`
- `votequorum(5)`
- `zypper (8)`

7.3 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>)

7.4 Related SAP Notes

- 405827 - Linux: Recommended file systems (<https://launchpad.support.sap.com/#/notes/405827>)
- 953653 - Rolling Kernel Switch (<https://launchpad.support.sap.com/#/notes/953653/E>)
- 1056161 - SUSE Priority Support for SAP applications (<https://launchpad.support.sap.com/#/notes/1056161>)

- 1092448 - IBM XL C/C++ runtime environment for Linux on system p (<https://launchpad.support.sap.com/#/notes/1092448/E>)
- 1153713 - Problems with SAP Management Console (Java) (<https://launchpad.support.sap.com/#/notes/1153713/E>)
- 1275776 - Linux: Preparing SLES for SAP environments (<https://launchpad.support.sap.com/#/notes/1275776>)
- 1763512 - Support details for SUSE Linux Enterprise for SAP Applications (<https://launchpad.support.sap.com/#/notes/1763512/E>)
- 1944799 - SAP HANA Guidelines for SLES Operating System Installation (<https://launchpad.support.sap.com/#/notes/1944799>)
- 1984787 - SUSE LINUX Enterprise Server 12: Installation notes (<https://launchpad.support.sap.com/#/notes/1984787/E>)
- 2077934 - Rolling kernel switch in HA environments (<https://launchpad.support.sap.com/#/notes/2077934/E>)
- 2205917 - SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12 (<https://launchpad.support.sap.com/#/notes/2205917>)
- 2235581 - SAP HANA: Supported Operating Systems (<https://launchpad.support.sap.com/#/notes/2235581/E>)
- 2254173 - Linux: Rolling Kernel Switch in Pacemaker based NetWeaver HA environments (<https://launchpad.support.sap.com/#/notes/2254173/E>)
- 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/E>)
- 2578899 - SUSE Linux Enterprise Server 15: Installation Note (<https://launchpad.support.sap.com/#/notes/2578899>)
- 2684254 - SAP HANA DB: Recommended OS settings for SLES 15 / SLES for SAP Applications 15 (<https://launchpad.support.sap.com/#/notes/2684254>)
- 2855499 - FAIL: RKS Warning(s): Unsupported SCS instance with additional gateway found (<https://launchpad.support.sap.com/#/notes/2855499>)

- 2992028 - SYB: Fault Manager: install Fault Manager in separated instance (<https://launchpad.support.sap.com/#/notes/2992028>)
- 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>)

8 Appendix

8.1 CRM configuration of the two-node cluster

Find below the complete crm configuration for SAP system HA1. This example is for the two node cluster, but without the simple mount setup.

```
#
# nodes
#
node 1084753931: hacert01
node 1084753932: hacert02
#
# primitives for ASCS and ERS
#
primitive rsc_fs_HA1_ASCS00 Filesystem \
  params device="/dev/sdb2" directory="/usr/sap/HA1/ASC00" fstype=xfst \
  op start timeout=60s interval=0 \
  op stop timeout=60s interval=0 \
  op monitor interval=20s timeout=40s
primitive rsc_fs_HA1_ERS10 Filesystem \
  params device="/dev/sdb3" directory="/usr/sap/HA1/ERS10" fstype=xfst \
  op start timeout=60s interval=0 \
  op stop timeout=60s interval=0 \
  op monitor interval=20s timeout=40s
primitive rsc_ip_HA1_ASCS00 IPAddr2 \
  params ip=192.168.201.115 \
  op monitor interval=10s timeout=20s
primitive rsc_ip_HA1_ERS10 IPAddr2 \
  params ip=192.168.201.116 \
  op monitor interval=10s timeout=20s
primitive rsc_sap_HA1_ASCS00 SAPIInstance \
  op monitor interval=11 timeout=60 on-fail=restart \
  params InstanceName=HA1_ASCS00_saphalas \
  START_PROFILE="/sapmnt/HA1/profile/HA1_ASCS00_saphalas" \
  AUTOMATIC_RECOVER=false \
  meta resource-stickiness=5000 failure-timeout=60 migration-threshold=1 \
  priority=10
primitive rsc_sap_HA1_ERS10 SAPIInstance \
  op monitor interval=11 timeout=60 on-fail=restart \
  params InstanceName=HA1_ERS10_saphaler \
  START_PROFILE="/sapmnt/HA1/profile/HA1_ERS10_saphaler" \
  AUTOMATIC_RECOVER=false IS_ERS=true \
  meta priority=1000
#
```



```

# SBD with adapted timing
#
primitive stonith-sbd stonith:external/sbd \
  params pcmk_delay_max=30
#
# group definitions for ASCS and ERS
#
group grp_HA1_ASCS00 rsc_ip_HA1_ASCS00 rsc_fs_HA1_ASCS00 rsc_sap_HA1_ASCS00 \
  meta resource-stickiness=3000
group grp_HA1_ERS10 rsc_ip_HA1_ERS10 rsc_fs_HA1_ERS10 rsc_sap_HA1_ERS10
#
# constraints between ASCS and ERS
#
colocation col_sap_HA1_not_both -5000: grp_HA1_ERS10 grp_HA1_ASCS00
location loc_sap_HA1_failover_to_ers rsc_sap_HA1_ASCS00 \
  rule 2000: runs_ers_HA1 eq 1
order ord_sap_HA1_first_ascs Optional: rsc_sap_HA1_ASCS00:start rsc_sap_HA1_ERS10:stop
symmetrical=false
#
# crm properties and more
#
property cib-bootstrap-options: \
  have-watchdog=true \
  dc-version=1.1.18+20180430.b12c320f5-1.14-b12c320f5 \
  cluster-infrastructure=corosync \
  cluster-name=hacluster \
  stonith-enabled=true \
  last-lrm-refresh=1494346532
rsc_defaults rsc-options: \
  resource-stickiness=1 \
  migration-threshold=3
op_defaults op-options: \
  timeout=600 \
  record-pending=true

```

8.2 Example for the two-node cluster with simple-mount setup

In contrast to the traditional setups, this setup uses an additional NFS mount for the SAP application layer without the need to have dedicated block devices and cluster-controlled file systems. That greatly simplifies the overall architecture, implementation and maintenance of a SUSE Linux Enterprise High Availability cluster for SAP NetWeaver with SAP Enqueue Replication Server. See also SUSE TID <https://www.suse.com/support/kb/doc/?id=000019944>.

8.2.1 systemd services for simple-mount setup

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

```
# systemctl disable SAPHA1_00.service
# systemctl disable SAPHA1_10.service
```

With the `sapstartsrv-resource-agents` RPM package there come two **systemd** services called `sapping` and `sappong`. `sapping` runs before `sapinit` and moves `/usr/sap/sapservices` out of the way. `sappong` runs after `sapinit` and moves `/usr/sap/sapservices` back to its original location.

```
# zypper info sapstartsrv-resource-agents
# systemctl enable sapping
# systemctl enable sappong
```

See manual pages `ocf_suse_SAPStartSrv(7)`, `sapping(8)` and `SAPStartSrv_basic_cluster(7)` for details.

8.2.2 fstab entries for simple-mount setup

As the directories `/sapmnt/HA1`, `/usr/sap/HA1` need to be available at all time, make sure they are mounted during boot. This can be achieved by putting the information into the `/etc/fstab`. Mount options may depend on your particular environment.

```
nfs1:/x/sapmnt/HA1 /sapmnt/HA1 nfs 0 0
nfs1:/x/ursap/HA1 /usr/sap/HA1 nfs 0 0
```

See also manual pages `SAPStartsrv_basic_cluster(8)` and `mount.nfs(8)`.

8.2.3 CRM configuration fragments with simple-mount setup

Find below crm configuration fragments for SAP system HA1. This example shows the specific items for the two-node cluster with priority fencing. This configuration is basically the same as above, except the file system resources are replaced by `SAPStartSrv` resources.

```
...
#
# primitives for ASCS and ERS, SAPStartSrv resources replacing Filesystem
#
primitive rsc_SAPStartSrv_HA1_ASCS00 ocf:suse:SAPStartSrv \
    params InstanceName=HA1_ASCS00_sapha1as
primitive rsc_SAPStartSrv_HA1_ERS10 ocf:suse:SAPStartSrv \
    params InstanceName=HA1_ERS10_sapha1er
```

```

#
# primitives for ASCS and ERS, SAPInstance option MINIMAL_PROBE=true
#
primitive rsc_sap_HA1_ASCS00 SAPInstance \
    op monitor interval=11 timeout=60 on-fail=restart \
    params InstanceName=HA1_ASCS00_saphalas \
        START_PROFILE="/usr/sap/HA1/SYS/profile/HA1_ASCS00_saphalas" \
        AUTOMATIC_RECOVER=false MINIMAL_PROBE=true \
meta resource-stickiness=5000 failure-timeout=60 migration-threshold=1 \
    priority=10
primitive rsc_sap_HA1_ERS10 SAPInstance \
    op monitor interval=11 timeout=60 on-fail=restart \
    params InstanceName=HA1_ERS10_saphaler \
        START_PROFILE="/usr/sap/HA1/SYS/profile/HA1_ERS10_saphaler" \
        AUTOMATIC_RECOVER=false IS_ERS=true MINIMAL_PROBE=true \
meta priority=1000
...
#
# group definitions for ASCS and ERS
#
group grp_HA1_ASCS00 rsc_ip_HA1_ASCS00 \
    rsc_SAPStartSrv_HA1_ASCS00 rsc_sap_HA1_ASCS00 \
    meta resource-stickiness=3000
group grp_HA1_ERS10 rsc_ip_HA1_ERS10 \
    rsc_SAPStartSrv_HA1_ERS10 rsc_sap_HA1_ERS10
...

```

See also manual page `ocf_suse_SAPStartSrv(7)`.

8.3 Corosync configuration of the two-node cluster

Find below a corosync configuration example for one corosync ring. Ideally, two rings would be used.

```

# 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: 192.168.201.111
        nodeid: 1
    }
    node {
        ring0_addr: 192.168.201.112
        nodeid: 2
    }
}

quorum {
    provider: corosync_votequorum
    expected_votes: 2
    two_node: 1
}

```

8.4 /usr/sap/sapservices without native systemd integration


```
#!/bin/sh
```

```
LD_LIBRARY_PATH=/usr/sap/HA1/ASCS00/exe:$LD_LIBRARY_PATH; export LD_LIBRARY_PATH; /usr/
sap/HA1/ASCS00/exe/sapstartsrv pf=/usr/sap/HA1/SYS/profile/HA1_ASCS00_saph1as -D -u
ha1adm
LD_LIBRARY_PATH=/usr/sap/HA1/ERS10/exe:$LD_LIBRARY_PATH; export LD_LIBRARY_PATH; /usr/
sap/HA1/ERS10/exe/sapstartsrv pf=/usr/sap/HA1/ERS10/profile/HA1_ERS10_saph1er -D -u
ha1adm
```

9 Legal notice

Copyright © 2006–2025 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.

10 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.