

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

Setup Guide

SUSE Linux Enterprise Server for SAP applications 16
SAP S/4HANA

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

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

Setup Guide

Date: 2026-09-10

SUSE® Linux Enterprise Server for SAP applications is optimized in various ways for SAP* applications. This document explains how to deploy an S/4HANA Enqueue Replication 2 High Availability Cluster solution. It is based on SUSE Linux Enterprise Server for SAP applications 16.0. The concept, however, can also be used with other minor versions of SUSE Linux Enterprise Server for SAP applications 16.

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	5
3	Overview	6
4	Installing the SAP system	12
5	Implementing the cluster	20
6	Setting up multi-node cluster setups for SAP S/4HANA	26
7	Administration	34
8	References	41
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 S/4HANA ABAP Platform.

SAP S/4HANA ABAP Platform is a common stack of middleware functionalities 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.

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 S/4HANA ABAP Platform with SAP Enqueue Replication Server 2.

For additional information on the simple mount architecture, read:

- SUSE blog article "Simple Mount Structure for SAP Application Platform" (<https://www.suse.com/c/simple-mount-structure-for-sap-application-platform/> .
- SUSE knowledge base TID (Technical Information Document) 00019944 "Use of Filesystem resource for ASCS/ERS HA setup not possible" (<https://www.suse.com/support/kb/doc/?id=000019944> .
- Manual page `ocf_suse_SAPStartSrv(7)`.

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. Among them are:

- SUSE product documentation at <https://documentation.suse.com/> 
- SUSE best practices Web page at <https://documentation.suse.com/sbp/sap-16/> 
- Supported high availability solutions by SUSE Linux Enterprise Server for SAP applications overview at <https://documentation.suse.com/sles-sap/sap-ha-support/> 

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

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.
- plan and implement an NFS-based storage layout for 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 and the SAPStartSrv resource agent.

This guide focuses on high availability of the central services. For high availability of the SAP HANA database, consult the guide for the scale-up performance-optimized scenario (see [Section 1.2, “Additional documentation and resources”](#)).

3 Overview

This document describes setting up a pacemaker high-availability (HA) cluster using SUSE Linux Enterprise Server for SAP applications 16 for the SAP Enqueue Replication (ENSA2) scenario. For the setup described in this document, two or three nodes are used for the ASCS central services instance and ERS replicated enqueue instance. These nodes are controlled by the SUSE Linux Enterprise High Availability cluster. Additional nodes are used for running the database, and the PAS and AAS application server instances. Finally, you need a highly available NFS server.

The goals for the setup include:

- Implementation of a cluster with a shared SAP applications directory
- Integration of the **SAPStartSrv** resource agent
- Integration of the cluster with the native systemd 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

All shared SAP resources are located on a central NFS server.

Shared disks allow using SBD as the cluster fencing mechanism.

3.1 Installing typical systems for ASCS, ERS, database and additional SAP instances

The document at hand describes the installation of a distributed SAP system on three and more systems. In this setup, only two or three systems reside inside the cluster. The database and SAP dialog instances can be controlled by another cluster. We recommend installing 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

The cluster described in this guide manages only the SAP central services instances ASCS and ERS.

For the SAP HANA database, we recommend setting up the performance-optimized system replication scenario with SAPHanaSR-angi. The SAPHanaSR-angi solution should be set up in a separate two-node cluster. The setup is described in a separate best practices document available from the SUSE Best Practices on documentation portal at <https://documentation.suse.com/sbp/sap-16/>.




FIGURE 1: TYPICAL SYSTEMS

CLUSTERED MACHINES TWO-NODE SCENARIO

- One machine (valuga11) for ASCS; Virtual hostname: sapen2as
- One machine (valuga12) for ERS; Virtual hostname: sapen2er

ALSO NEEDED MACHINES - OPTIONALLY CLUSTERED

- One machine (valuga01) for DB; Virtual hostname: sapen2db
- One machine (valuga13) for the PAS; Virtual hostname: sapen2d1
- One machine (valuga14) for the AAS; Virtual hostname: sapen2d2

3.2 Implementing high availability for the SAP HANA database - optional

We recommend implementing SAP HANA system replication (SR) automation to enhance the scenario described in this document.




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

Version for SAP S/4HANA ABAP Platform on Linux on AMD64/Intel 64 and IBM PowerLE. Find more information about the supported combinations of OS and databases for SAP S/4HANA Server 2021 or newer 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).

3.3 Integrating 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 starting with 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 cluster software interfaces, the cluster framework detects that as an error status. The framework then restores the SAP instance to its previous status by either starting or stopping it. 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 status changes to the cluster software. This avoids dangerous situations as previously described. For more details, read 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/>.




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

For this scenario, the latest 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 latest 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 the SAP management consoles (SAP MMC or SAP MC). Starting with version 3.1.0, the maintenance mode of cluster resources triggered with SAP `sapcontrol` commands is supported. See also manual page `sap_suse_cluster_connector(8)`.

3.4 Sharing disks and NFS

XFS is used for all local file systems. For `/sapmnt/<SID>` and `/usr/sap/<SID>`, NFS is used.

3.4.1 Sharing disk for SBD and mounting NFS for cluster ASCS and ERS

The disk for the fencing mechanism SBD must be shared and assigned to the cluster nodes `valuga11` and `valuga12` in the two-node cluster example. The NFS file systems for the ASCS and ERS instances are mounted both on `valuga11` and `valuga12`. They can also be mounted on the SAP application servers (in this example `valuga13/PAS` and `valuga14/AAS`) to simplify the storage layout of the complete SAP system. Make sure that your shared SBD disk `/dev/disk/by-id/SUSE-Example-A` is visible on `valuga11` and `valuga12`:

```
# lsblk | grep /dev/disk/by-id/SUSE-Example-A
```

During the SAP software installation, you need to mount via NFS and entries in `/etc/fstab`.

3.4.2 Preparing the disk for database and dialog instances (HANA DB)

The disk `/dev/disk/by-id/SUSE-Example-B` for the database (260 GB) is assigned to `valuga01` and formatted with XFS.

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

EXAMPLE 1: CREATE PARTITIONS AND FILE SYSTEMS FOR DB AND APP SERVERS ON VALUGA01

```
# lsblk
# parted -s /dev/disk/by-id/SUSE-Example-B print
# # we are on the 'correct' drive, right?
# mkfs.xfs /dev/disk/by-id/SUSE-Example-B
# mkdir /hana
```

```
# echo "/dev/disk/by-id/SUSE-Example-B /hana xfs defaults 0 2" >> /etc/fstab
# mount /dev/disk/by-id/SUSE-Example-B
```



Note

D01: Starting with NetWeaver 7.5, the primary application server instance directory is named 'D<Instance_Number>'.

NFS SERVER

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

MEDIA

- 192.168.1.1:/sapmedia /sapmedia

3.5 Adding 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
192.168.1.105 valuga13

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
```

3.6 Creating 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.

- Create mount points and mounting NFS shares on all non-HANA nodes, for example on valuga11, valuga12. If the application servers should also profit from the simplified storage setup, they can use the share as well.

EXAMPLE 2: MOUNT NFS SHARES ON ALL NON-HANA NODES

```
# mkdir -p /sapmnt/EN2 /usr/sap/EN2 /sapmedia
# echo "192.168.1.1:/data/export/S4_HA_CLU_10/EN2/sapmnt/EN2 /sapmnt/EN2 nfs
defaults 0 0" >> /etc/fstab
# echo "192.168.1.1:/data/export/S4_HA_CLU_10/EN2/usr/sap/EN2 /usr/sap/EN2 nfs
defaults 0 0" >> /etc/fstab
# mount /sapmnt/EN2
# mount /usr/sap/EN2
```

Mount options may depend on your particular environment. See also manual pages SAPStartsrv_basic_cluster(8) and mount.nfs(8).

4 Installing the SAP system

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 is a distributed SAP installation as illustrated here:




FIGURE 4: 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	halshm	1003
User	en2adm	2001
User	sapadm	2002
User	haladm	2003

4.2 Installing ASCS on valuga11

Temporarily set the service IP address which you will later use in the cluster as the local IP address, because the installer needs to be able to resolve and use it. Use the correct virtual hostname for each installation step. If applicable, mount file systems such as */sapmedia/*.

```
# ip a a 192.168.1.112/24 dev ens1
# # if not mounted yet, mount these now
# mount /sapmnt/EN2
# mount /usr/sap/EN2
# cd /sapmedia/SWPM20_P9/
# ./sapinst SAPINST_USE_HOSTNAME=sapen2as
```

- SWPM product installation path:
 - Installing SAP S/4HANA Server 2021 → 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>.
- Verify during the parameter review that 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 set the service IP address which you will later use in the cluster as the local IP address, because the installer needs to be able to resolve and use it. Use the correct virtual hostname for each installation step.

```
# ip a a 192.168.1.113/24 dev ens1
# # if not mounted yet, mount these now
# mount /sapmnt/EN2
# mount /usr/sap/EN2
# cd /sapmedia/SWPM20_P9/
# ./sapinst SAPINST_USE_HOSTNAME=sapen2er
```

- SWPM product installation path:
 - Installing SAP S/4HANA Server 2021 → SAP HANA DATABASE → Installation → Application Server ABAP → High-Availability System → ERS Instance
- Use instance number 10.
- Deselect using FQDN.
- Verify 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 Performing 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*, run:

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

On *valuga12*, run:

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

4.4.2 Register ASCS and ERS

ASCS is installed on the first cluster node, ERS is installed on the second. To integrate ASCS and ERS with systemd on all cluster nodes, the SAP instances need to be registered on each node.

On the first cluster node ERS needs to be registered, on the second cluster node ASCS needs to be registered.

- Register the ERS SAP instance on the first cluster node (*valuga11*):

```
# 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
```

- Register the ASCS SAP instance on the second cluster node (*valuga12*):

```
# 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>_ASCS<instanceNumberAscs>_<virtHostNameAscs> -reg
```

After registration, the list of systemd services must be the same and should look like:

```
# systemctl list-unit-files | grep SAP  
SAPEN2_00.service disabled disabled  
SAPEN2_10.service disabled disabled  
SAP.slice static -
```

After registration, the service might still be enabled. The next section describes how to disable the SAP<sid>_<ino> services.

4.4.3 Disabling systemd services of the ASCS and the ERS SAP instance

This step is required to transfer control of 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 the instance services will stop the SAP instance as well. Starting the instance services will not start the SAP instance.

- Check the SAP systemd integration:

```
# systemctl list-unit-files | grep SAP  
SAPEN2_00.service disabled disabled  
SAPEN2_10.service disabled 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 remains for compatibility. It lists native systemd commands, with one line per registered instance.

4.4.4 Integrating the cluster framework using `sap-suse-cluster-connector`

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/`.

Configure the `sapstartsrv` service to load the HA script connector library and use the connector `sap-suse-cluster-connector`. Ensure 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)`.

4.4.5 Adapting SAP profiles

For the ASCS instance, change the start command from *Restart_Program_xx* to *Start_Program_xx* for the enqueue server (ENSA2). This setting prevents the SAP start framework from automatically restarting the enqueue process. A restart results in lost 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 (ENSA2).

File **/usr/sap/EN2/SYS/profile/EN2_ERS10_sapen2er:**

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

4.4.6 Starting ASCS and ERS

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

On *valuga12*, type the following

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

On *valuga11*, type the following:

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

4.5 Installing 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 \(https://help.sap.com/docs/SAP_HANA_PLATFORM/eb3777d5495d46c5b2fa773206bbf-b46/0b15a92d554c4941a452c9ca127f8c70.html\)](https://help.sap.com/docs/SAP_HANA_PLATFORM/eb3777d5495d46c5b2fa773206bbf-b46/0b15a92d554c4941a452c9ca127f8c70.html).

```
# ip a a 192.168.1.114/24 dev ens1
# mount /dev/disk/by-id/SUSE-Example-B /hana
# cd /sapmedia/SWPM20_P9/
# ./sapinst SAPINST_USE_HOSTNAME=sapen2db
```

- SWPM product installation path:
 - Installing SAP S/4HANA Server 2021 → SAP HANA DATABASE → Installation → Application Server ABAP → High-Availability System → Database Instance
- Profile directory is */sapmnt/EN2/profile*.

- Deselect using FQDN.
- Database parameters : Database ID (DBSID) is HA1; Database Host is sapen2db; Instance Number is 53.
- Database System ID enter Instance Number is 53; 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.
- Verify during the parameter review that virtual name **sapen2db** is used.

4.6 Installing the primary application server (PAS) on valuga13

```
# ip a a 192.168.1.110/24 dev ens1
# mount /dev/disk/by-id/SUSE-Example-B-part2 /usr/sap/EN2/D01
# cd /sapmedia/SWPM20_P9/
# ./sapinst SAPINST_USE_HOSTNAME=sapen2d1
```

- SWPM product installation path:
 - Installing SAP S/4HANA Server 2021 → 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.
- No SLD is used.
- Verify during the parameter review that virtual name **sapen2d1** is used.

4.7 Installing an additional application server (AAS) on valuga14

```
# ip a a 192.168.1.111/24 dev ens1
# mount /dev/disk/by-id/SUSE-Example-B-part3 /usr/sap/EN2/D02
# cd /sapmedia/SWPM20_P9/
# ./sapinst SAPINST_USE_HOSTNAME=sapen2d2
```

- SWPM product installation path:
 - Installing SAP S/4HANA Server 2021 → 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.
- Verify during the parameter review that virtual name **sapen2d2** is used.

4.8 Optional: Preparing additional cluster nodes

To control the ASCS and ERS in a cluster with three or more nodes, prepare each node to run the SAP instances. Ensure SAP Linux users and system configurations are in place. SWPM 2.0 already includes an installation option to prepare additional cluster nodes. In this case, `sapinst` is called without `SAPINST_USE_HOSTNAME`.

Ensure that the SAP NFS shares are also mounted on the additional cluster node.

```
# cd /sapmedia/SWPM20_P9/
# ./sapinst
```

- SWPM product installation path:
 - Installing SAP S/4HANA Server 2021 → SAP HANA DATABASE → Installation → Application Server ABAP → High-Availability System → Prepare Additional Cluster Node
- Provide SAP profile path.
- Verify the parameter review.

5 Implementing the cluster

The main procedure to implement the cluster is as follows:

TASKS

1. Prepare the operating system and install the cluster software.
2. Configure the cluster base including corosync and resource manager.

3. Configure the cluster resources.
4. Tuning 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

Create the SBD device/partition before starting. Verify which device/partition to use. In this setup guide, the disk `/dev/disk/by-id/SUSE-Example-A` is reserved for SBD usage.

5.1 Preparing the operating system and installing the cluster software

- Set up and enable `chrony`.
- Install the RPM pattern `sles_sap_HAAPP` on both cluster nodes.

```
# zypper in -t pattern sles_sap_HAAPP
```

5.2 Configuring the cluster base

TASKS

- To configure the cluster base, use the interactive command line tool `crm` (crm cluster init).
- Install and configure the watchdog device on the first machine.

Instead of deploying the software-based solution, rather 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 | grep "(wd|dog|i6|iT|ibm)"
```

- Install and configure the cluster stack on the first machine

```
# crm cluster init -s /dev/disk/by-id/SUSE-Example-A
```

- 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, use the interactive command line tool `crm`.

```
# crm cluster join -c valugall
```

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

```
Stack: corosync
Current DC: valugall (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 valugall

2 nodes configured
1 resource configured

Online: [ valugall valuga12 ]

fencing-sbd (stonith:fence_sbd): Started valugall
```

5.3 Configuring cluster resources

The `SAPInstance` resource configuration is needed to start and stop the ASCS and the ERS instances themselves. See manual page `ocf_heartbeat_SAPInstance(7)` for details.

The `SAPStartSrv` resource starts and stops the `sapstartsrv` service and guarantees that only one instance is running per cluster at the same time. See manual page `ocf_suse_SAPStartSrv(7)` for details.

In the latest version of ENSA2, the ASCS instance receives the enqueue lock table over the network from the ERS instance. Usually, both instances run on different nodes. However, if only node is available, the ASCS instance starts on the same host as the the ERS instance.




FIGURE 5: RESOURCES AND CONSTRAINTS

5.3.1 Preparing the cluster for adding the resources

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

As user *root*, type the following command:

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

5.3.2 Preparing the new SAPStartSrv resource agent implementation

As a prerequisite for having a single NFS mount for ASCS and ERS, the `sapstartsrv` instance agents of ASCS and ERS must not be started by the `sapinit` service during system start-up, as these services are started and stopped by dedicated cluster resources.

With the `sapstartsrv-resource-agents` RPM package, two systemd services called `sapping` and `sappong` are included. `sapping` runs before `sapinit` and moves `/usr/sap/sapservices` out of the way. Consequently, the `sapstartsrv` instance agents are not started automatically by `sapinit`. `sappong` runs after `sapinit` and moves `/usr/sap/sapservices` back to its original location.

On `valuga11` and `valuga12`, enable the `sapping` and `sappong` services.

```
# systemctl enable sapping
# systemctl enable sappong
```

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

5.3.3 Configuring resources for the ASCS Instance

First, configure the resources for the IP address, the SAP instance agent and the SAP instance. You need to adapt the parameters for your specific environment.

Ensure the parameter **MINIMAL_PROBE** is set to **true** in the SAP instance definition .

EXAMPLE 3: ASCS PRIMITIVE

```
primitive rsc_ip_EN2_ASCS00 IPAddr2 \  
  params ip=192.168.1.112 \  
  op monitor interval=21 timeout=20  
primitive rsc_SAPStartSrv_EN2_ASCS00 ocf:suse:SAPStartSrv \  
  params InstanceName=EN2_ASCS00_sapen2as  
primitive rsc_sap_EN2_ASCS00 SAPInstance \  
  op monitor interval=22 timeout=60 on-fail=restart \  
  params InstanceName=EN2_ASCS00_sapen2as \  
    START_PROFILE="/usr/sap/EN2/SYS/profile/EN2_ASCS00_sapen2as" \  
    AUTOMATIC_RECOVER=false MINIMAL_PROBE=true \  
  meta resource-stickiness=5000
```

The shown SAPInstance monitor timeout is a trade-off between fast recovery of the ASCS vs. resilience against sporadic temporary NFS issues. You may slightly increase it to fit your infrastructure. Consult the storage or NFS server documentation for appropriate timeout values. Make sure the **SAPStartSrv** resource has **NO** monitor operation configured. See also manual pages `ocf_heartbeat_SAPInstance(7)`, `ocf_heartbeat_IPAddr2(7)`, `ocf_suse_SAPStartSrv(7)` and `nfs(5)`.

EXAMPLE 4: ASCS GROUP

```
group grp_EN2_ASCS00 \  
  rsc_ip_EN2_ASCS00 rsc_SAPStartSrv_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
```

5.3.4 Configuring resources for the ERS Instance

Next, configure the resources for the IP address, the SAP instance agent and the SAP instance. You need to adapt the parameters for your specific environment.

Ensure the parameter **MINIMAL_PROBE** is set to **true** in the SAP instance definition .

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

EXAMPLE 5: ERS PRIMITIVE

```
primitive rsc_ip_EN2_ERS10 IPAddr2 \  
  params ip=192.168.1.113 \  
  op monitor interval=10 timeout=20  
primitive rsc_SAPStartSrv_EN2_ERS10 ocf:suse:SAPStartSrv \  
  params InstanceName=EN2_ERS10_sapen2er  
primitive rsc_sap_EN2_ERS10 SAPInstance \  
  op monitor interval=22 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 MINIMAL_PROBE=true
```

The shown SAPInstance monitor timeout is a trade-off between fast recovery of the ERS vs. resilience against sporadic temporary NFS issues. You may slightly increase it to fit your infrastructure. Consult the storage or NFS server documentation for appropriate timeout values. Ensure the **SAPStartSrv** resource has **NO** monitor operation configured. See also manual pages `ocf_heartbeat_SAPInstance(7)`, `ocf_heartbeat_IPAddr2(7)` `ocf_suse_SAPStartSrv(7)` and `nfs(5)`.

EXAMPLE 6: ERS GROUP

```
group grp_EN2_ERS10 \  
  rsc_ip_EN2_ERS10 rsc_SAPStartSrv_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*, type the following command:

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

5.3.5 Configuring the colocation constraints between ASCS and ERS

An ASCS instance should avoid starting up on the cluster node running the ERS instance if any other node is available. 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 takeover.

EXAMPLE 7: LOCATION CONSTRAINT

```
colocation col_sap_EN2_separate -5000: grp_EN2_ERS10 grp_EN2_ASCS00
order ord_sap_EN2_ascs_first 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
```

5.3.6 Activating the cluster

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

As user *root*, type the following command:

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

6 Setting up multi-node cluster setups for SAP S/4HANA

Multi-node cluster setups are cluster configurations with more than two nodes. Depending on the starting point, you can or extend a two-node cluster setup or directly start with more than two nodes for an ASCS / ERS high availability setup. The following examples show how to set up a multi-node cluster and how to extend an existing two-node cluster. They outline major configuration changes and the basic preparation of a new member node. The task list for setting up a three-node cluster is similar to the task list for a two-node cluster. However, some details differ to allow for a diskless SBD setup. A diskless SBD setup is an optional improvement for three nodes, but does not work on two nodes. Conversely, priority fencing is an optional improvement for two nodes, but does not work on three nodes. An example priority fencing configuration for a two-node cluster is available in the appendix. See the SUSE Linux Enterprise High Availability product documentation for details.



Note

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

6.1 Extending an existing two-node cluster configuration

TASKS

1. Backing up the current cluster
2. Installing the operating system of the new node
3. Patching the existing nodes
4. Preparing the new node's operating system
5. Installing the cluster software on the new node
6. Preparing **SAPStartSrv** resource agent on the new node
7. Preparing the SAP installation on the new node
8. Adding the new node to the cluster
9. Testing the new cluster configuration

6.1.1 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 With Simple Mount - Setup Guide*.

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

```
# crm configure save 2nodes_crm.txt
# supportconfig
```

Back up the existing `/etc/corosync/corosync.conf` and all other files required to restore the system. The following example is one method for creating a backup. The important step is storing the backup on an external destination.

```
# tar cvf /<path to external storage>/backup_2nodes_configuration.tar \  
  /etc/corosync/corosync.conf \  
  /etc/corosync/authkey \  
  /etc/sysconfig/sbd \  
  2nodes_crm.txt \  
  /var/log/scc_*.txz  
<add your additional files here>
```

6.1.2 Installing the operating system of the new node

We recommend using automated installation to ensure that the system setup across nodes is identical. Document any additional steps taken beyond the automated setup. In this example, we deploy machines with Agama and use Ansible for configuration.

6.1.3 Patching the 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 Multi Linux Manager, add the new system to the same repositories to ensure they have the same patch level as your existing nodes.

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

- We recommend 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 to do so:

```
# zypper patch  
## or  
# zypper update
```

6.1.4 Preparing the new node's operating system

EXAMPLE 8: CHECK FOR VALID DNS, TIME SYNCHRONIZATION, NETWORK SETTINGS

- Verify that DNS is working:

```
# ping <hostname>
```

- Verify `chrony` is active and working:

```
# systemctl status chronyd
```

```
# chronyc sources
```

- 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.

6.1.5 Installing cluster software on the new node

EXAMPLE 9: INSTALLING PACKAGES

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

```
# zypper in -t pattern sles_sap_HAAPP
```

6.1.6 Preparing **SAPStartSrv** resource agent on the new node

- Enable the sapping and sappong services.

```
# systemctl enable sapping  
# systemctl enable sappong
```

After installing all necessary packages, compare the installed package versions.

EXAMPLE 10: CHECK IF ALL NODES HAVE THE SAME SOFTWARE PACKAGES AND VERSIONS:

- On the existing cluster nodes, type:

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

- On the new node, type:

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

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

```
# vimdiff rpm_valuga13.txt rpm_valuga11.txt
```



Note

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

EXAMPLE 11: INSTALL AND CONFIGURE THE WATCHDOG DEVICE ON THE NEW MACHINE.

Instead of deploying the software-based solution, 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 identic SBD configurations as the two existing nodes

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

6.1.7 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 helps to prepare a new host which can later run either ASCS or ERS in the cluster environment. Create the directory structure that should run the SAP resource. The instance directory is located on an NFS share for all nodes.

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

EXAMPLE 12: MOUNT NFS SHARES ON VALUGA13

```
# mkdir -p /sapmnt/EN2 /usr/sap/EN2 /sapmedia
# mount -t nfs 192.168.1.1:/data/export/S4_HA_CLU_10/EN2/sapmnt/EN2 /sapmnt/EN2
# mount -t nfs 192.168.1.1:/data/export/S4_HA_CLU_10/EN2/usr/sap/EN2 /usr/sap/EN2
# mount -t nfs 192.168.1.1:/sapmedia /sapmedia
```

As the directories */sapmnt/EN2*, */usr/sap/EN2* 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*.

The next step requires the following information:

- profile directory
- password for SAP System Administrator
- UID for sapadm

```
# cd /sapmedia/SWPM20_P9/  
# ./sapinst
```

- SWPM product installation path:
 - Installing SAP S/4HANA Server 2021 → SAP HANA DATABASE → Installation → Application Server ABAP → High-Availability System → Prepare Additional Cluster Node
- Use /sapmnt/EN2/profile for the profile directory
- All passwords: <use-your-secure-pwd>
- UID for sapadm: 2002

EXAMPLE 13: 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>_<virtHostNameAscs> -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
```

6.1.8 Adding the new node to the cluster

EXAMPLE 14: ADDING THE NEW NODE TO THE CLUSTER

Check if the SBD device is available in case the SBD fencing method is in place for the two nodes. If the existing cluster uses a different, supported fencing mechanism, check and verify it also for the new cluster node.

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

- Joining the cluster can be done with **crm cluster join**:

```
# crm cluster join -c valugall
```

After the new node has joined the cluster, the configuration must be adapted to the new situation. Double-check if joining the cluster was successful and verify the file `/etc/corosync/corosync.conf`.

```
# awk '/quorum/,/}/' /etc/corosync/corosync.conf
# corosync-quorumtool -s
```

The values **expected_votes** and **two_node** should now look like the following 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
```

6.1.9 Testing the new cluster configuration

We recommend running test to verify that the new configuration is working as expected. Find a list of tests in the basic setup for the two node cluster above.

6.2 Pros and Cons for odd and even numbers of cluster nodes

There are certain use cases and infrastructure requirements which end up in different setups. Below, find a list of advantages and disadvantages of special configurations:

- 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 itself)
 - Advantage: cluster can 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 planning effort and complexity for infrastructure planning

7 Administration

7.1 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 setup.

A minimal example sequence for checking the cluster status might look like the following:

```
# crm_mon -lr
# crm configure show | grep cli-
# cs_clusterstate -i
```

See also manual pages `cs_clusterstate(8)`, `crm(8)` and `crm_mon(8)`.

7.1.1 Maintenance procedure for a Linux cluster or operating system with ASCS and ERS instances remain running

Check the state of the Linux cluster and ASCS/ERS:

```
# cs_clusterstate -i
```

The command returns the following output:

```
Cluster state: S_IDLE
```

Display the cluster summary, node list and resource list with the following command:

```
# crm_mon -lr
```

Check for any undesired location constraints:

```
# crm configure show | grep cli-
```

If the command returns no output, no undesired constraints exist.

Display a list of system instances:

```
# su - en2adm -c "sapcontrol -nr 00 -function GetSystemInstanceList"
```

Display a list of processes running on the ASCS and ERS instances:

```
# su - en2adm -c "sapcontrol -nr 00 -function GetProcessList"
```

Check whether maintenance mode is set in the cluster configuration:

```
# su - en2adm -c "sapcontrol -nr 00 -function HACheckMaintenanceMode"
```

Display details for the cluster, available HA nodes and the active instance node:

```
# su - en2adm -c "sapcontrol -nr 00 -function HAGetFailoverConfig"
```

Before you set the Linux cluster into maintenance mode, check its state by running the `cs_clusterstate -i` command. Then run the command below to set the cluster into maintenance mode:

```
# crm maintenance on
```

Stop the Linux cluster on all nodes:

```
# crm cluster stop --all
```

You can now perform maintenance on the Linux cluster or system. Before performing maintenance, stop any SAP instances as necessary. If you stopped the SAP instances, start them before activating the cluster.

When maintenance is complete, start the Linux cluster on all nodes:

```
# crm cluster start
```

Let the Linux cluster detect the status of the ASCS and ERS resources:

```
# crm resource refresh rsc_sap_EN2_ASCS00  
# crm resource refresh rsc_sap_EN2_ERS10
```

Set the cluster to ready for operations:

```
# cs_clusterstate -i  
# crm maintenance off
```

Check the status of the Linux cluster and ASCS/ERS:

```
# cs_clusterstate -i  
# crm_mon -lr  
# crm configure show | grep cli-  
# su - en2adm -c "sapcontrol -nr 00 -function GetSystemInstanceList"
```

```
# su - en2adm -c "sapcontrol -nr 00 -function GetProcessList"
# su - en2adm -c "sapcontrol -nr 10 -function GetProcessList"
# su - en2adm -c "sapcontrol -nr 00 -function HACheckMaintenanceMode"
# su - en2adm -c "sapcontrol -nr 00 -function HAGetFailoverConfig"
# cs_clusterstate -i
```

7.1.2 Migrating the ASCS instance

To **migrate** the ASCS SAP instance, 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*, 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 ""
```

7.1.3 Using unique instance numbers

All SAP instance numbers controlled by the cluster must be unique. If you need multiple dialog instances with the same instance number running on different systems, they must not be controlled by the cluster.

7.1.4 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"
```

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

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

7.1.5 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"
```

See also manual page `crm(8)`.

7.1.6 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
```

7.1.7 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
```

7.1.8 Cleaning up resources

You can also **clean up resource failures**. Failures are automatically deleted to allow a failback after a specified period of time. You can also clean up the status, including the failures.

Run the following command as user *root*:

```
# crm resource cleanup _RESOURCE-NAME_
```

7.2 Testing the cluster

We recommend to perform at least the following tests before you go into production with your cluster.

7.2.1 Checking product names with `HAGetFailoverConfig`

Verify that 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 S/4HANA cluster integration.

As user *en2adm*, type the following command:

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

7.2.2 Running SAP checks using HACheckConfig and HACheckFailoverConfig

Verify that the HA configuration tests pass without errors.

As user *en2adm*, type the following commands:

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

7.2.3 Manually migrating ASCS

Verify that 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
```

7.2.4 Migrating ASCS using HAFailoverToNode

Verify that moving the ASCS instance using SAP tools like *sapcontrol* works properly.

As user *en2adm*, type the following command:

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

7.2.5 Test ASCS migration after operating system failure

Verify that 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
```

7.2.6 Restarting ASCS in-place using Stop and Start

Verify that the in-place restart of the SAP resources succeeded. The SAP instance should not failover to another node. It starts on the same node where it stopped.

As user *en2adm*, do the following:

```
## example for ASCS
# sapcontrol -nr 00 -function Stop
# sapcontrol -nr 00 -function WaitforStopped 60 20
## cs_clusterstate -i
# sapcontrol -nr 00 -function Start
```

7.2.7 Performing an automated restart of the ASCS instance (simulating rolling kernel switch)

The following test verifies that the cluster solution neither interacts with nor restarts the ASCS instance during a maintenance procedure. In addition, it verifies that no locks are lost when restarting an ASCS instance during a rolling kernel switch (RKS) procedure. The cluster solution should recognize that the restart of the ASCS instance was expected. No failures or errors 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 so. 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.

Perform this test multiple times within a short time frame. In the following example, the ASCS instance restarts 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
```

7.2.8 Performing an RKS

The 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 with 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>
```

7.2.9 Additional tests

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

- Check the recoverable and non-recoverable outage of the SAP 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.
- Check the outage of network between cluster nodes.

8 References

For more information, see the documents listed below.

8.1 SUSE product documentation

- SUSE product documentation: <https://documentation.suse.com/> 
- SUSE release notes: <https://documentation.suse.com/?tab=release-notes> 
- SUSE best practices: <https://documentation.suse.com/sles-sap/sap-ha-support/>  <https://documentation.suse.com/sbp/sap-16/> 

8.2 Pacemaker

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

8.3 Related Manual Pages

- `chronyc(8)`
- `corosync.conf(8)`

- [corosync_overview\(8\)](#)
- [crm\(8\)](#)
- [crm_mon\(8\)](#)
- [crm_simulate\(8\)](#)
- [cs_clusterstate\(8\)](#)
- [cs_man2pdf\(8\)](#)
- [cs_show_sbd_devices\(8\)](#)
- [cs_wait_for_idle\(8\)](#)
- [ha_related_sap_notes\(7\)](#)
- [ha_related_suse_tids\(7\)](#)
- [mount.nfs\(8\)](#)
- [ocf_heartbeat_IPaddr2\(7\)](#)
- [ocf_heartbeat_SAPInstance\(7\)](#)
- [ocf_suse_SAPStartSrv\(7\)](#)
- [sapping\(7\)](#)
- [sapservices-move\(8\)](#)
- [SAPStartSrv_basic_cluster\(7\)](#)
- [SAPStartSrv_maintenance_procedures\(7\)](#)
- [sap_suse_cluster_connector\(8\)](#)
- [saptune\(8\)](#)
- [sbd\(8\)](#)
- [stonith_sbd\(7\)](#)
- [supportconfig\(8\)](#)
- [systemctl\(8\)](#)

- `systemd-cgls(8)`
- `usermod(8)`
- `votequorum(5)`
- `zypper(8)`

8.4 Related SUSE TIDs

- Use of Filesystem resource for ASCS/ERS HA setup not possible (<https://www.suse.com/support/kb/doc/?id=000019944> )
- 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 `HAAActive: FALSE` (<https://www.suse.com/support/kb/doc/?id=000019924> )
- Identify unplanned start/stop of the ASCS/ERS services (<https://www.suse.com/support/kb/doc/?id=000020661> )
- Combining SAP HANA database and SAP S/4HANA application platform in one cluster. (<https://www.suse.com/support/kb/doc/?id=000022133> )

8.5 Related SUSE blogs

- Simple Mount Structure for SAP Application Platform (<https://www.suse.com/c/simple-mount-structure-for-sap-application-platform/> )
- First certified simple-mount cluster ever! (<https://www.suse.com/c/first-certified-simple-mount-cluster-ever/> )
- Handover for the Next Round – SAP on SUSE Cluster and systemd Native Integration (<https://www.suse.com/c/handover-for-the-next-round-sap-on-suse-cluster-and-systemd-native-integration/> )

8.6 Related SAP Documentation

- SAP Product Availability Matrix (https://support.sap.com/en/release-upgrade-maintenance.html#section_1969201630)
- Important Resources about the HA-Interface Certification for Partners (<https://wiki.scn.sap.com/wiki/display/SI/Important+Resources+about+the+HA-Interface+Certification+for+Partners>)

8.7 Related SAP Notes

- 768727 - Automatic restart functions in sapstart for processes (<https://launchpad.support.sap.com/#/notes/768727>)
- 927637 - Web service authentication in sapstartsrv as of Release 7.00 (<https://launchpad.support.sap.com/#/notes/927637>)
- 953653 - Rolling Kernel Switch (<https://launchpad.support.sap.com/#/notes/953653>)
- 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>)
- 1693245 - SAP HA Script Connector Library (<https://launchpad.support.sap.com/#/notes/1693245>)
- 1763512 - Support details for SUSE Linux Enterprise for SAP applications (<https://launchpad.support.sap.com/#/notes/1763512>)
- 1864705 - Compatibility tests for certified HA setups (<https://launchpad.support.sap.com/#/notes/1864705>)
- 1872602 Rolling kernel switch using kernel release limits (<https://launchpad.support.sap.com/#/notes/1872602>)
- 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> )
- 2625407 - SAP S/4HANA 1809: Release Information Note (<https://launchpad.support.sap.com/#/notes/2625407> )
- 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> )
- 2711036 - Usage of the Standalone Enqueue Server 2 in an HA Environment (<https://launchpad.support.sap.com/#/notes/2711036> )
- 2714839 - New security settings for S/4HANA 1909 (and later) (<https://launchpad.support.sap.com/#/notes/2714839> )
- 2717369 - Download files for installing of SAP S/4HANA 1809, SAP S/4HANA 1909, BW S/4HANA 2.0 (<https://launchpad.support.sap.com/#/notes/2717369> )
- 2769531 - SAP S/4HANA 1909: Release Information Note (<https://launchpad.support.sap.com/#/notes/2769531> )
- 2884313 - SAP S/4HANA 2020: Release Information Note (<https://launchpad.support.sap.com/#/notes/2884313> )
- 2855499 - FAIL: RKS Warning(s): Unsupported SCS instance with additional gateway found (<https://launchpad.support.sap.com/#/notes/2855499> )
- 3015539 - SAP S/4HANA 2021: Release Information Note (<https://launchpad.support.sap.com/#/notes/3015539> )

- 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> )
- 3145277 - SAP S/4HANA 2022: Release Information Note (<https://launchpad.support.sap.com/#/notes/3145277> )
- 3307222 - SAP S/4HANA 2023: Release Information Note (<https://launchpad.support.sap.com/#/notes/3145277> )
- 3493301 - SAP S/4HANA 2025: Release Information Note (<https://launchpad.support.sap.com/#/notes/3493301> )
- 3565382 - SUSE Linux Enterprise Server 16: Installation Note (<https://launchpad.support.sap.com/#/notes/3565382> )

9 Appendix

9.1 CRM configuration of a two-node cluster

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

```
node 1: valuga11
node 2: valuga12
primitive rsc_ip_EN2_ASCS00 IPAddr2 \
    params ip=192.168.1.112 \
    op monitor interval=10 timeout=20
primitive rsc_ip_EN2_ERS10 IPAddr2 \
    params ip=192.168.1.113 \
    op monitor interval=10 timeout=20
primitive rsc_SAPStartSrv_EN2_ASCS00 ocf:suse:SAPStartSrv \
    params InstanceName=EN2_ASCS00_sapen2as
primitive rsc_SAPStartSrv_EN2_ERS10 ocf:suse:SAPStartSrv \
    params InstanceName=EN2_ERS10_sapen2er
primitive rsc_sap_EN2_ASCS00 SAPIInstance \
    op monitor interval=21 timeout=60 on-fail=restart \
    params InstanceName=EN2_ASCS00_sapen2as \
        START_PROFILE="/usr/sap/EN2/SYS/profile/EN2_ASCS00_sapen2as" \
        AUTOMATIC_RECOVER=false MINIMAL_PROBE=true \
        meta resource-stickiness=5000
primitive rsc_sap_EN2_ERS10 SAPIInstance \
    op monitor interval=22 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 MINIMAL_PROBE=true
primitive fencing-sbd stonith:fence_sbd \
    params pcmk_delay_max=30
group grp_EN2_ASCS00 rsc_ip_EN2_ASCS00 \
    rsc_SAPStartSrv_EN2_ASCS00 rsc_sap_EN2_ASCS00 \
    meta resource-stickiness=3000
group grp_EN2_ERS10 rsc_ip_EN2_ERS10 \
    rsc_SAPStartSrv_EN2_ERS10 rsc_sap_EN2_ERS10
colocation col_sap_EN2_separate -5000: \
    grp_EN2_ERS10 grp_EN2_ASCS00
order ord_sap_EN2_ascs_first 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=2nodes \
fencing-enabled=true \
fencing-timeout=150 \
placement-strategy=balanced
rsc_defaults rsc-options: \
resource-stickiness=1 \
migration-threshold=3 \
failure-timeout=86400
op_defaults op-options: \
timeout=600 \
record-pending=true

```

9.2 CRM configuration fragments of the two-node cluster with priority fencing

Find below `crm` configuration fragments for SAP system EN2. This example shows the specific items for the two-node cluster with priority fencing. This configuration matches the previous one, except for the following items:

```

...
primitive rsc_sap_EN2_ASCS00 SAPInstance \
  op monitor interval=21 timeout=60 on-fail=restart \
  params InstanceName=EN2_ASCS00_sapen2as \
    START_PROFILE="/usr/sap/EN2/SYS/profile/EN2_ASCS00_sapen2as" \
    AUTOMATIC_RECOVER=false MINIMAL_PROBE=true \
    meta resource-stickiness=5000 priority=100
...
primitive fencing-sbd stonith:fence_sbd \
  params pcmk_delay_max=15
...
property cib-bootstrap-options: \
  have-watchdog=true \
  cluster-infrastructure=corosync \
  cluster-name=2nodes \
  fencing-enabled=true \
  fencing-timeout=150 \
  placement-strategy=balanced \
  priority-fencing-delay=30
...

```

9.3 Corosync configuration of the two-node cluster

Find below the Corosync configuration for one corosync ring. Ideally, two rings are used.

```
valuga11:~ # cat /etc/corosync/corosync.conf
totem {
    version: 2
    cluster_name: 2nodes
    transport: knet
    token: 5000
    join: 60
    max_messages: 20
    token_retransmits_before_loss_const: 10
    crypto_hash: sha256
    crypto_cipher: aes256
}

quorum {
    provider: corosync_votequorum
    two_node: 1
}

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

nodelist {
    node {
        ring0_addr: 192.168.1.103
        name: {my2nd1
        nodeid: 1
    }

    node {
        ring0_addr: 192.168.1.104
        name: valuga12
        nodeid: 2
    }
}
```

9.4 Corosync configuration of the multi-node cluster

Find below the Corosync configuration for one corosync ring. Ideally, two rings are used.

```
valuga11:~ # cat /etc/corosync/corosync.conf
totem {
    version: 2
    cluster_name: 3nodes
    transport: knet
    token: 5000
    join: 60
    max_messages: 20
    token_retransmits_before_loss_const: 10
    crypto_hash: sha256
    crypto_cipher: aes256
}

quorum {
    provider: corosync_votequorum
    two_node: 0
    expected_votes: 3
}

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

nodelist {
    node {
        ring0_addr: 192.168.1.103
        name: valuga11
        nodeid: 1
    }

    node {
        ring0_addr: 192.168.1.104
        name: valuga12
        nodeid: 2
    }

    node {
        ring0_addr: 192.168.1.105
        name: valuga13
        nodeid: 3
    }
}
```

```
}  
}
```

10 Legal Notice

Copyright © 2006-2026 SUSE LLC and contributors. All rights reserved.

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

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

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

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

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

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 retile any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

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

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

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

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

5. COMBINING DOCUMENTS

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

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

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

6. COLLECTIONS OF DOCUMENTS

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

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

7. AGGREGATION WITH INDEPENDENT WORKS

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

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

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

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

9. TERMINATION

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

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <https://www.gnu.org/copyleft/> .

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been

published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

```
Copyright (c) YEAR YOUR NAME.
```

```
Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.2
or any later version published by the Free Software Foundation;
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.
A copy of the license is included in the section entitled "GNU
Free Documentation License".
```

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

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

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

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