

SAP S/4HANA and SAP NetWeaver Multi-SID Cluster Guide

SUSE Linux Enterprise Server for SAP Applications 12 and 15

Bernd Schubert, Solution Architect SAP (SUSE)

Fabian Herschel, Distinguished Architect SAP (SUSE)

Date: 2024-12-12

SUSE® Linux Enterprise Server for SAP Applications is optimized in various ways for SAP* applications. This document explains how to implement multiple SAP NetWeaver and S/4HANA systems in a High Availability Cluster solution. It is based on SUSE Linux Enterprise Server for SAP Applications 12 and 15. The concept can also be used with newer service packs of SUSE Linux Enterprise Server for SAP Applications.

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

Contents

- 1 About This Guide 4
- 2 Scope of This Document 5
- 3 Multi-SID - Central Services 11
- 4 Multi-SID - Database Services 28
- 5 Legal Notice 29
- 6 GNU Free Documentation License 30

1 About This Guide

1.1 Introduction

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

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

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

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

1.2 Additional Documentation and Resources

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

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

More guides and best practices documents referring to SUSE Linux Enterprise Server and SAP can be found and downloaded at the SUSE Best Practices Web page:

<https://documentation.suse.com/sbp/sap> .

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

Additional resources, such as customer references, brochures or flyer, can be found at the SUSE Linux Enterprise Server for SAP Applications resource library:

<https://www.suse.com/products/sles-for-sap/#resources> .

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 integrate the central services with multi-SID setups in one cluster. Other architectures such as the SAP HANA performance optimized scenario will follow.

2.1 Architecture Central Services

2.1.1 Integration of SAP S/4HANA and 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 service **sapstartsrv** controls SAP instances since SAP Kernel versions 6.40. One of the classic problems running SAP instances in a highly available environment is the following: If an SAP administrator changes the status (start/stop) of an SAP instance without using the interfaces provided by the cluster software, the cluster framework will detect that as an error status. In consequence, it will bring the SAP instance into the old status by either starting or stopping the SAP instance. This can result in very dangerous situa-

tions, if the cluster changes the status of an SAP instance during some SAP maintenance tasks. The new updated solution enables the central component **sapstartsrv** to report state changes to the cluster software. This avoids dangerous situations as previously described. More details can be found in the blog article "Using sap_vendor_cluster_connector for interaction between cluster framework and sapstartsrv" at <https://blogs.sap.com/2014/05/08/using-sapvendorcluster-connector-for-interaction-between-cluster-framework-and-sapstartsrv/comment-page-1/>.



Note

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

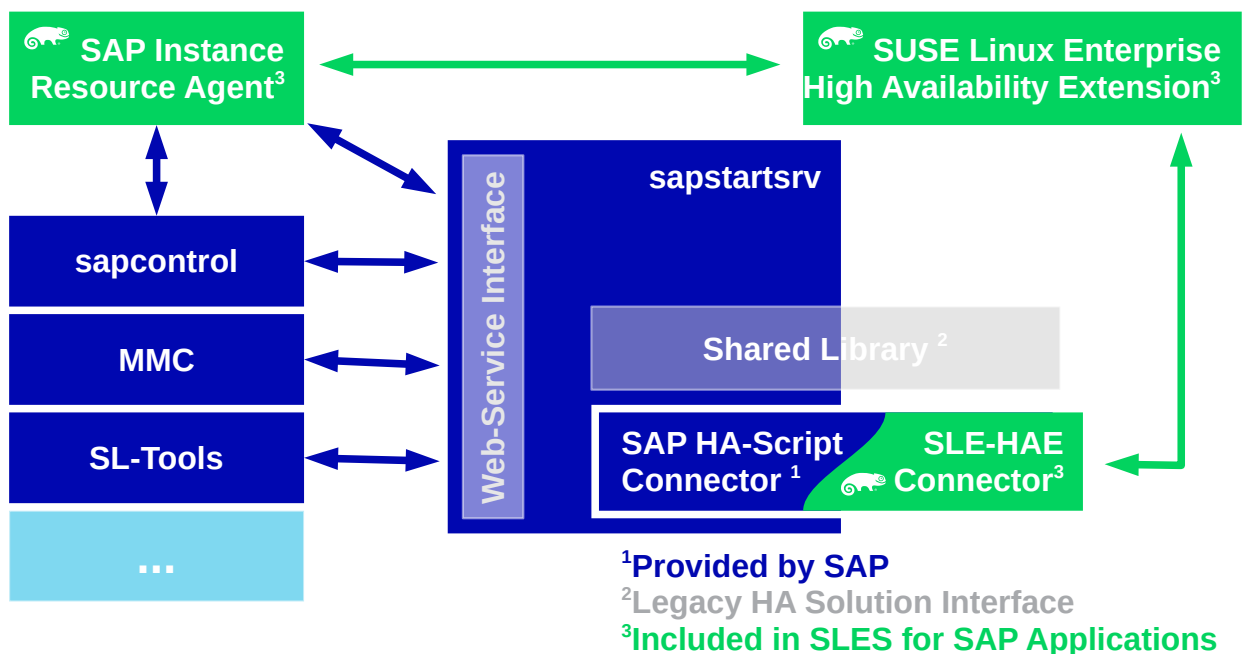


FIGURE 1: CLUSTER CONNECTOR TO INTEGRATE THE CLUSTER WITH THE SAP START FRAMEWORK



Note

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

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

2.2 Architecture DB Services

SUSE has developed multiple cluster architectures for SAP databases running with high availability. Not all of those architectures are planned for multi-SID configurations. The following sections describe common SAP HANA scenarios. You will also find a note if the architecture is available for multi-SID.

2.2.1 Scale-Up Scenarios and Resource Agents

SUSE has implemented the scale-up scenario with the SAPHana resource agent (RA), which performs the actual check of the SAP HANA database instances. This RA is configured as a master/slave resource. In the scale-up scenario, the master assumes responsibility for the SAP HANA databases running in primary mode. The slave is responsible for instances that are operated in synchronous (secondary) status.

To make configuring the cluster as simple as possible, SUSE developed the SAPHanaTopology resource agent. This RA runs on all nodes of a SUSE Linux Enterprise Server for SAP Applications cluster and gathers information about the statuses and configurations of SAP HANA system replications. It is designed as a normal (stateless) clone.

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

- **Performance optimized ($A \Rightarrow B$).** This scenario and setup **will be described in this document in a future update.**

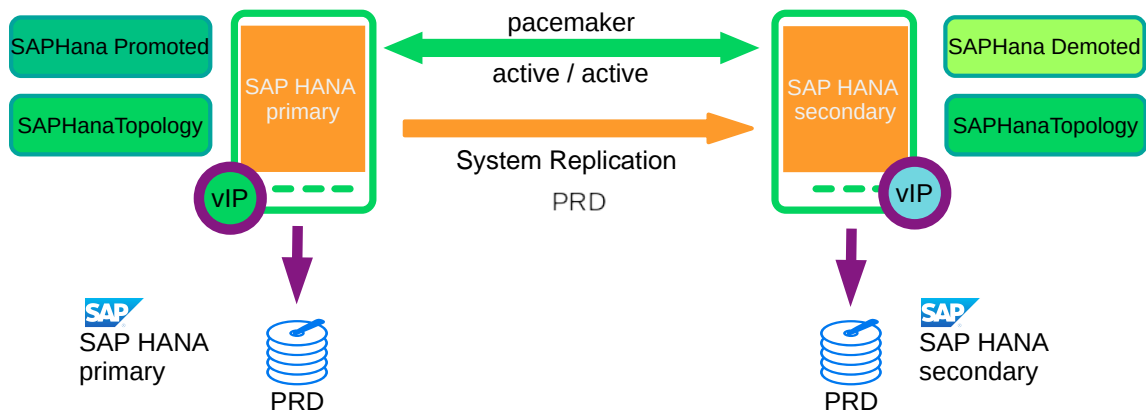


FIGURE 2: SAP HANA SYSTEM REPLICATION SCALE-UP IN THE CLUSTER - PERFORMANCE OPTIMIZED

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

One big advantage of the performance optimized scenario of SAP HANA is the possibility to allow read access on the secondary database site. To support this **read enabled** scenario, a second virtual IP address is added to the cluster and bound to the secondary role of the system replication.



Note

The performance optimized scenario could be implemented also for multi-SID setups, if it is guaranteed that all system replication pairs are handled by the cluster separately. It is not allowed to force an SAP HANA system to process a takeover only because another one is taking over. There is also no load balance for the primary sides.

- **Cost optimized** ($A \Rightarrow B, Q$). This scenario and setup is described in another document available from the documentation Web page (<https://documentation.suse.com/sbp/sap/>). The document for *cost optimized* is named "Setting up an SAP HANA SR Cost Optimized Infrastructure".

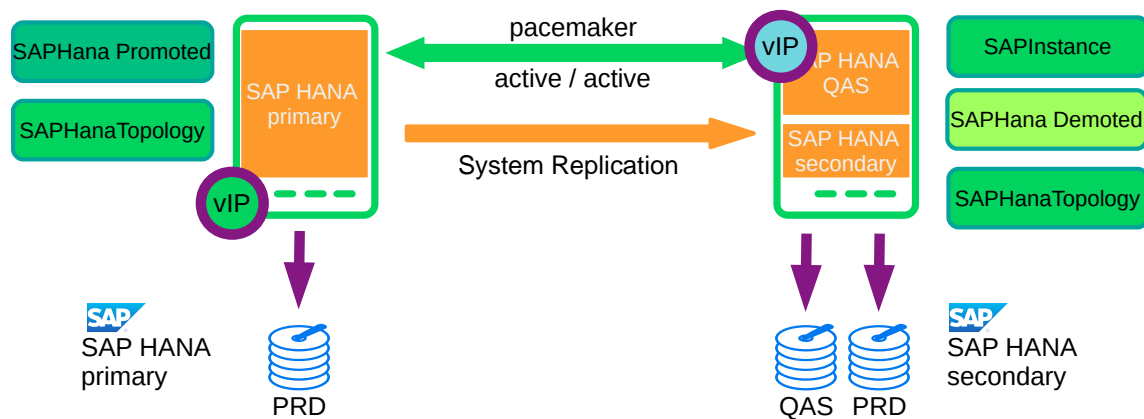


FIGURE 3: SAP HANA SYSTEM REPLICATION SCALE-UP IN THE CLUSTER - COST OPTIMIZED

In the cost optimized scenario the second node is also used for a non-productive SAP HANA RDBMS system (like QAS or TST). Whenever a takeover is needed, the non-productive system must be stopped first. As the productive secondary system on this node must be limited in using system resources, the table preload must be switched off. A possible takeover needs longer than in the performance optimized use case.

In the cost optimized scenario the secondary needs to be running in a reduced memory consumption configuration. This is why *read enabled* must not be used in this scenario.



Note

The cost optimized scenario is not intended to be used for multi-SID configuration. The reason is that, whenever one of the primary systems would be moved to the secondary node, a second memory setup for the primary systems would be needed. Otherwise the removed memory limitation would tell the SAP HANA to consume all available memory. Currently there is no globally defined space where to set such parameters to be used after a takeover.

- **Multi-Tier** ($A \Rightarrow B \rightarrow C$) and **Multi-Target** ($B \Leftarrow A \Rightarrow C$).




FIGURE 4: SAP HANA SYSTEM REPLICATION SCALE-UP IN THE CLUSTER - PERFORMANCE OPTIMIZED CHAIN

A *Multi-tier* system replication has an additional target. In the past this third side must have been connected to the secondary (chain topology). With current SAP HANA versions also *multiple target topology* is allowed by SAP.

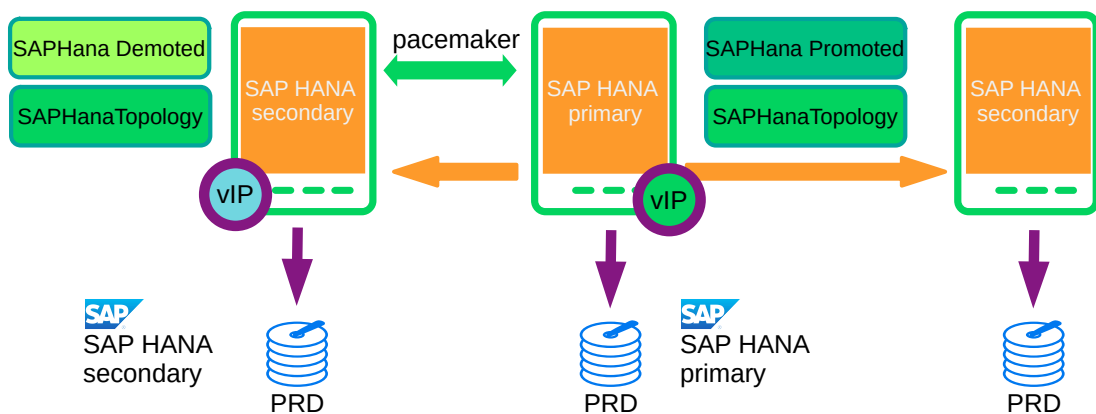


FIGURE 5: SAP HANA SYSTEM REPLICATION SCALE-UP IN THE CLUSTER - PERFORMANCE OPTIMIZED MULTI-TARGET

Multi-tier and multi-target systems are implemented as described in this document. Only the first replication pair (A and B) is handled by the cluster itself. The main difference to the plain performance optimized scenario is that the auto registration must be switched off.



Note

From a cluster perspective, this type of scenario is very similar to the performance optimized scenario. As long as only A and B are driven by the cluster this scenario is also expected to work with multi-SID setups. However, this has not been tested yet in the labs. Support for this scenario is expected to be added in the future.

- **Multi-tenancy or MDC.**

Multi-tenancy is supported for all of the above scenarios and use cases. This scenario is supported since SAP HANA SPS09. The setup and configuration from a cluster point of view is the same for multi-tenancy and single container. Thus you can use the above documents for both kinds of scenarios.



Note

Multi-tenancy has no direct side effect to multi-SID setups. This means that multi-tenancy is supported in combination with multi-SID, if the basic architecture is supported with multi-SID.

3 Multi-SID - Central Services

This guide will cover the following scenarios:

- Example ENSA1 in a Two-Node Cluster
- Example ENSA2 in a Two-Node Cluster
- Example ENSA1 in a Multi-Node Cluster
- Example ENSA2 in a Multi-Node Cluster
- Example ENSA1 and ENSA2 in a Two-Node Cluster
- Example ENSA1 and ENSA2 in a Multi-Node Cluster

Configuration examples in this guide are given with cluster-controlled filesystem resources. Nevertheless, also the simple-mount concept based on SAPStartSrv resources (<https://documentation.suse.com/sbp/sap/html/SAP-S4HA10-setupguide-simplemount-sle15/>) fits well for Multi-SID scenarios.

3.1 Two-Node Setup

- The hardware configuration of each node must be able to run all resources in case of a fail over
- Each SAP system (SID) must be independent from each other

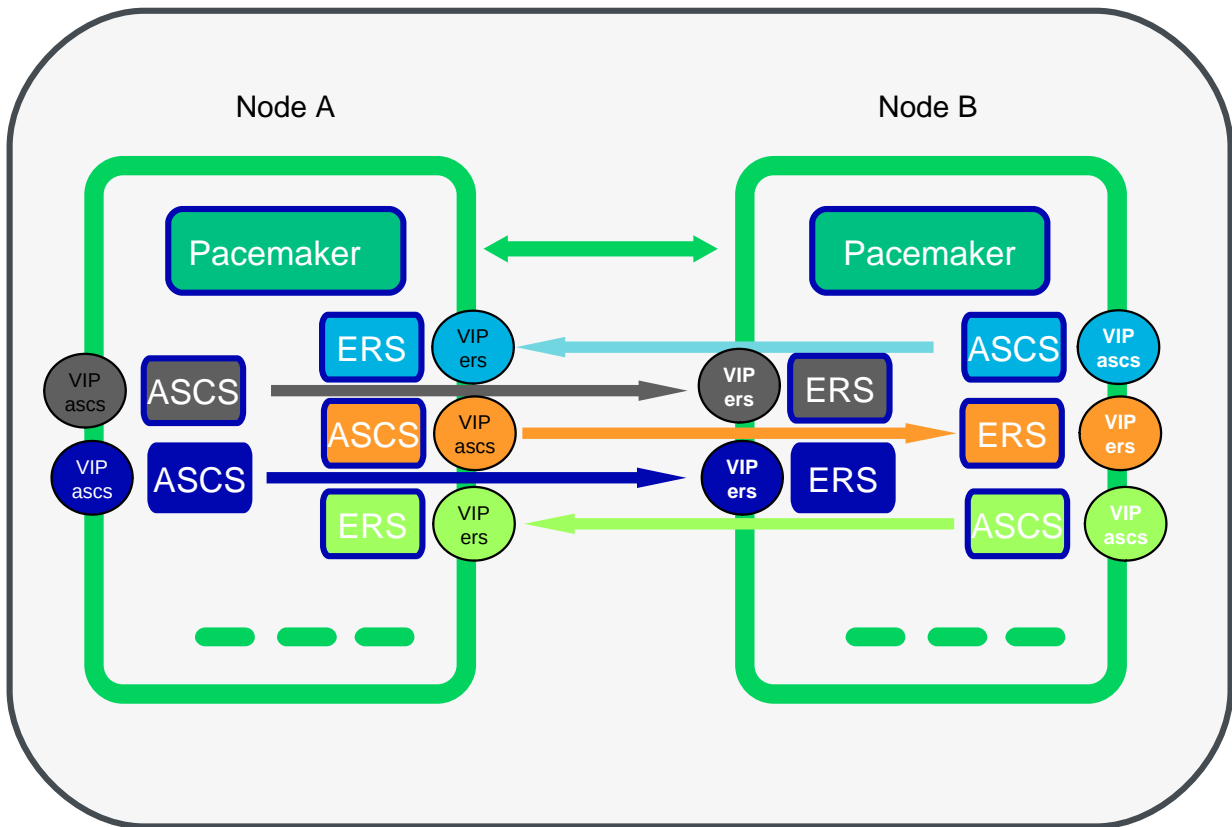


FIGURE 6: MULTIPLE SAP CENTRAL SERVICES INSTANCES RUNNING IN A TWO-NODE CLUSTER

3.2 Multi-Node Setup

- The hardware configuration of each node must be able to run additional resources in case of a fail over
- Each SAP system (SID) must be independent from each other
- ENSA2 load balancing (under observation)

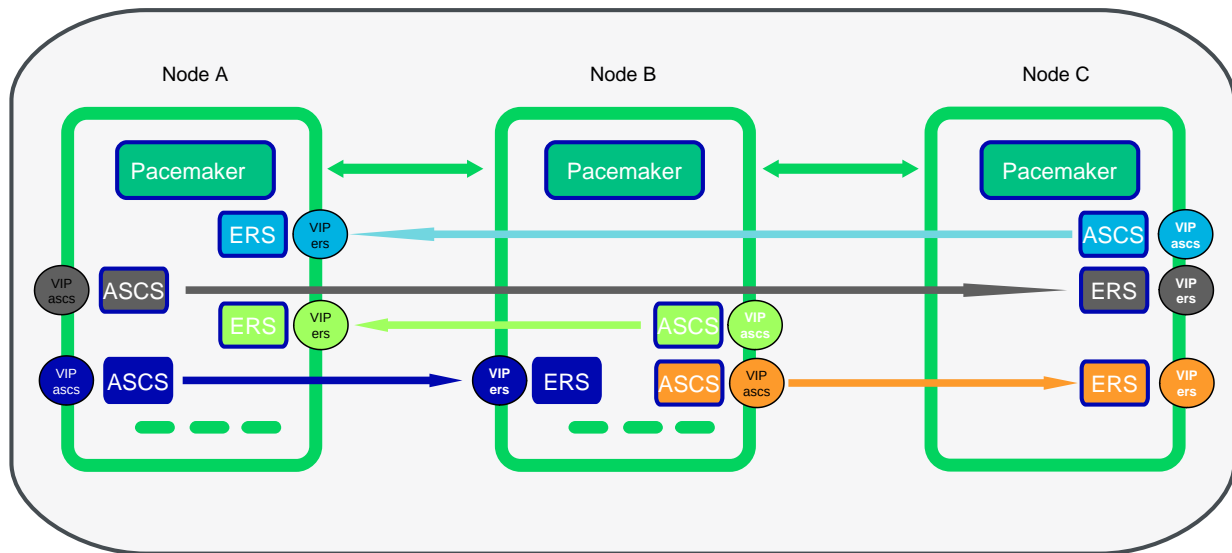


FIGURE 7: MULTIPLE SAP CENTRAL SERVICES INSTANCES RUNNING IN A MULTI-NODE CLUSTER

3.3 General Rules and Requirements



Warning

When adding more resources, be aware of the following risks:

- Higher complexity
- Higher impact if one nodes goes down
- Higher administrative effort in case of maintenance activities, like OS patching

3.3.1 Installation

The installation process is described in the **SUSE Best Practices** guides for SAP S/4HANA and SAP NetWeaver. They can be found here: <https://documentation.suse.com/sbp/sap>

Repeat the steps for preparing the hosts and infrastructure:

- Provide the storage and IP addresses as needed.
- Check the name resolution and time settings on each host.
- Provide the SAP installation sources as you will need.

- Create the directory structure similar to the first SAP instance for all new installations.
- Manually mount the file systems and assign the IP address for the virtual host name.
- Start SAP installation with virtual host name using *swpm*.

The installation process itself is similar to the first setup. The major difference is the ERS installation during the profile selection. At this point you need to choose the right directory to point to the correct ASCS installation.

3.3.2 Limitations

We recommend to limit the numbers of SAP systems (SID) running in one cluster. Each new SID will increase the complexity. The tested setup was done with up to **five** SID in one cluster.

3.4 Example ENSA1 in a Two-Node Cluster

As known from the setup based on the **SUSE Best Practices** document for SAP NetWeaver (<https://documentation.suse.com/sbp/sap> → SAP NetWeaver High Availability Cluster 7.40 - Setup Guide), the cluster integration can be done with a configuration file. Create one configuration file for each SID and load them one by one into the cluster. The description below will help to do the steps in the right direction.

The cluster is already up and running and still runs one ASCS / ERS configuration based on the group concept. The following tasks must be executed:

- Stop the new installed SAP system, which will be added to the cluster.
- Unmount the file system with *umount* and remove the IP address from the new SAP system.
- Check and/or modify the ASCS and ERS profile files and add the section for the sap-suse-cluster-connector integration, if not done already.
 - for example for VAS: `/sapmnt/VAS/profile/VAS_ASCS20_sapascs2`
 - for example for VAS: `/sapmnt/VAS/profile/VAS_ERS30_sapers2`
- Modify the *haclient* group on each cluster node and add the new <sid> adm to it, if not done already.
- Set the cluster into maintenance mode.

- Load the new configuration with *crm configure load update <filename>*.
- Release the cluster from maintenance mode.
- Check the cluster and SAP instances with *crm_mon -lrfn* and *sapcontrol -nr*

EXAMPLE 1: CONFIGURATION FILE AND ADAPTING CLUSTER

As user **root** prepare a file for the SAP system VAS

```
# cat ensal-2nd-vas.txt

primitive rsc_fs_VAS_ASCS20 Filesystem \
  params device="172.17.0.1:/srv/install/sapascs2" directory="/usr/sap/VAS/ASCs20"
  fstype=nfs options="vers=3" \
  op start timeout=60s interval=0 \
  op stop timeout=60s interval=0 \
  op monitor interval=20s timeout=300s
primitive rsc_fs_VAS_ERS30 Filesystem \
  params device="172.17.0.1:/srv/install/sapers2" directory="/usr/sap/VAS/ERS30"
  fstype=nfs options="vers=3" \
  op start timeout=60s interval=0 \
  op stop timeout=60s interval=0 \
  op monitor interval=20s timeout=300s
primitive rsc_ip_VAS_ASCS20 IPAddr2 \
  params ip=172.17.1.17 \
  op monitor interval=10s timeout=20s
primitive rsc_ip_VAS_ERS30 IPAddr2 \
  params ip=172.17.1.18 \
  op monitor interval=10s timeout=20s
primitive rsc_sap_VAS_ASCS20 SAPInstance \
  operations $id=rsc_sap_VAS_ASCS20-operations \
  op monitor interval=11 timeout=60 on-fail=restart \
  params InstanceName=VAS_ASCS20_sapascs2 START_PROFILE="/sapmnt/VAS/profile/
VAS_ASCS20_sapascs2" AUTOMATIC_RECOVER=false \
  meta resource-stickiness=5000 failure-timeout=60 \
    migration-threshold=1 priority=10
primitive rsc_sap_VAS_ERS30 SAPInstance \
  operations $id=rsc_sap_VAS_ERS30-operations \
  op monitor interval=11 timeout=60 on-fail=restart \
  params InstanceName=VAS_ERS30_sapers2 START_PROFILE="/sapmnt/VAS/profile/
VAS_ERS30_sapers2" AUTOMATIC_RECOVER=false IS_ERS=true \
  meta priority=1000
group grp_VAS_ASCS20 rsc_ip_VAS_ASCS20 rsc_fs_VAS_ASCS20 rsc_sap_VAS_ASCS20 \
  meta resource-stickiness=3000
group grp_VAS_ERS30 rsc_ip_VAS_ERS30 rsc_fs_VAS_ERS30 rsc_sap_VAS_ERS30
colocation col_sap_VAS_no_both -5000: grp_VAS_ERS30 grp_VAS_ASCS20
location loc_sap_VAS_failover_to_ers rsc_sap_VAS_ASCS20 \
```

```

rule 2000: runs_ers_VAS eq 1
order ord_sap_VAS_first_start_ascscs Optional: rsc_sap_VAS_ASCS20:start
rsc_sap_VAS_ERS30:stop symmetrical=false

```

EXAMPLE 2: CONFIGURATION EXAMPLE FOR A 3RD SID, FOR EXAMPLE WAS

As user **root** prepare a file for the SAP system WAS

```

# cat ensal-3nd-was.txt

primitive rsc_fs_WAS_ASCS31 Filesystem \
params device="172.17.0.1:/srv/install/sapascs3" directory="/usr/sap/WAS/ASC31"
fstype=nfs options="vers=3" \
op start timeout=60s interval=0 \
op stop timeout=60s interval=0 \
op monitor interval=20s timeout=300s
primitive rsc_fs_WAS_ERS41 Filesystem \
params device="172.17.0.1:/srv/install/sapers3" directory="/usr/sap/WAS/ERS41"
fstype=nfs options="vers=3" \
op start timeout=60s interval=0 \
op stop timeout=60s interval=0 \
op monitor interval=20s timeout=300s
primitive rsc_ip_WAS_ASCS31 IPAddr2 \
params ip=172.17.1.27 \
op monitor interval=10s timeout=20s
primitive rsc_ip_WAS_ERS41 IPAddr2 \
params ip=172.17.1.28 \
op monitor interval=10s timeout=20s
primitive rsc_sap_WAS_ASCS31 SAPInstance \
operations $id=rsc_sap_WAS_ASCS31-operations \
op monitor interval=11 timeout=60 on-fail=restart \
params InstanceName=WAS_ASCS31_sapascs3 START_PROFILE="/sapmnt/WAS/profile/
WAS_ASCS31_sapascs3" AUTOMATIC_RECOVER=false \
meta resource-stickiness=5000 failure-timeout=60 \
migration-threshold=1 priority=10
primitive rsc_sap_WAS_ERS41 SAPInstance \
operations $id=rsc_sap_WAS_ERS41-operations \
op monitor interval=11 timeout=60 on-fail=restart \
params InstanceName=WAS_ERS41_sapers3 START_PROFILE="/sapmnt/WAS/profile/
WAS_ERS41_sapers3" AUTOMATIC_RECOVER=false IS_ERS=true \
meta priority=1000
group grp_WAS_ASCS31 rsc_ip_WAS_ASCS31 rsc_fs_WAS_ASCS31 rsc_sap_WAS_ASCS31 \
meta resource-stickiness=3000
group grp_WAS_ERS41 rsc_ip_WAS_ERS41 rsc_fs_WAS_ERS41 rsc_sap_WAS_ERS41
colocation col_sap_WAS_no_both -5000: grp_WAS_ERS41 grp_WAS_ASCS31
location loc_sap_WAS_failover_to_ers rsc_sap_WAS_ASCS31 \
rule 2000: runs_ers_WAS eq 1

```



```
order ord_sap_WAS_first_start_ascs Optional: rsc_sap_WAS_ASCS31:start
rsc_sap_WAS_ERS41:stop symmetrical=false
```

EXAMPLE 3: CONFIGURATION EXAMPLE FOR A 4TH SID AND FOLLOWING, FOR EXAMPLE XAS

This is similar to the SID examples before. You must adapt the:

- File system sources which will be mounted
- IP address for virtual host name of ASCS and ERS
- Instance number for ASCS and ERS
- Local mount point for ASCS and ERS
- Profile path for ASCS and ERS

After the configuration has been prepared in a file, this can be loaded into the cluster.

EXAMPLE 4: ADD THE 2ND AND FURTHER SID INTO THE CLUSTER

The cluster is already up and running and still runs one ASCS / ERS configuration based on the group concept. The following procedure is recommended:

- Stop the new installed SAP system, which will be added into the cluster.
- Unmount the file system with *umount* and remove the IP address from the new SAP system.
- Check and/or modify the ASCS and ERS profile files and add the section for the sap-suse-cluster-connector integration, if not done already.
 - for example for VAS: */sapmnt/VAS/profile/VAS_ASCS20_sapascs2*
 - for example for VAS: */sapmnt/VAS/profile/VAS_ERS30_sapers2*
- Modify the *haclient* group on each cluster node and add the new <sid> adm to it, if not done already.
- Set the cluster into maintenance mode.
- Load the new configuration with *crm configure load update <filename>*.
- Pre-check the cluster with *crm status*.
- Release the cluster from maintenance mode.
- Check the cluster and SAP instances with *crm_mon* and *sapcontrol -nr*

Log in to one of the cluster nodes and run the commands from there.

As user **root**, set the cluster to maintenance and load the configuration into the cluster with *crm configure load update <filename>*.

```
# crm configure property maintenance-mode=true
# crm configure load update ensa1-2nd-vas.txt
# crm configure load update ensa1-3nd-was.txt
# crm configure load update ensa1-4nd-xas.txt
...
# crm status
```

You should get back the prompt only. No further messages should be shown. If a message is displayed, there might be something wrong in the configuration file.



Note

If a wrong path is used, for example for the profile, this will not be detected during the configuration load. However, it will be shown during the cluster start action.

As user **root** verify the new inactive loaded configuration

```
# crm status
freki:~ # crm status
Stack: corosync
Current DC: freki (version 1.1.18+20180430.b12c320f5-3.3.1-b12c320f5) - partition
with quorum
Last updated: Tue Oct 15 16:57:07 2019
Last change: Tue Oct 15 16:56:56 2019 by hacluster via crmd on freki

2 nodes configured
31 resources configured

*** Resource management is DISABLED ***
The cluster will not attempt to start, stop or recover services

Online: [ freki geri ]

Full list of resources:

Resource Group: grp_YAS_ASCS00
rsc_ip_YAS_ASCS00 (ocf::heartbeat:IPaddr2): Started freki
rsc_fs_YAS_ASCS00 (ocf::heartbeat:Filesystem): Started freki
rsc_sap_YAS_ASCS00 (ocf::heartbeat:SAPInstance): Started freki
Resource Group: grp_YAS_ERS10
rsc_ip_YAS_ERS10 (ocf::heartbeat:IPaddr2): Started geri
rsc_fs_YAS_ERS10 (ocf::heartbeat:Filesystem): Started geri
```

```

rsc_sap_YAS_ERS10 (ocf::heartbeat:SAPInstance): Started geri
rsc_ip_hawk (ocf::heartbeat:IPAddr2): Started freki
Resource Group: grp_XAS_ASCS42
rsc_ip_XAS_ASCS42 (ocf::heartbeat:IPAddr2): Stopped (unmanaged)
rsc_fs_XAS_ASCS42 (ocf::heartbeat:Filesystem): Stopped (unmanaged)
rsc_sap_XAS_ASCS42 (ocf::heartbeat:SAPInstance): Stopped (unmanaged)
Resource Group: grp_XAS_ERS52
rsc_ip_XAS_ERS52 (ocf::heartbeat:IPAddr2): Stopped (unmanaged)
rsc_fs_XAS_ERS52 (ocf::heartbeat:Filesystem): Stopped (unmanaged)
rsc_sap_XAS_ERS52 (ocf::heartbeat:SAPInstance): Stopped (unmanaged)
Resource Group: grp_WAS_ASCS31
rsc_ip_WAS_ASCS31 (ocf::heartbeat:IPAddr2): Stopped (unmanaged)
rsc_fs_WAS_ASCS31 (ocf::heartbeat:Filesystem): Stopped (unmanaged)
rsc_sap_WAS_ASCS31 (ocf::heartbeat:SAPInstance): Stopped (unmanaged)
Resource Group: grp_WAS_ERS41
rsc_ip_WAS_ERS41 (ocf::heartbeat:IPAddr2): Stopped (unmanaged)
rsc_fs_WAS_ERS41 (ocf::heartbeat:Filesystem): Stopped (unmanaged)
rsc_sap_WAS_ERS41 (ocf::heartbeat:SAPInstance): Stopped (unmanaged)
Resource Group: grp_ZAS_ASCS53
rsc_ip_ZAS_ASCS53 (ocf::heartbeat:IPAddr2): Stopped (unmanaged)
rsc_fs_ZAS_ASCS53 (ocf::heartbeat:Filesystem): Stopped (unmanaged)
rsc_sap_ZAS_ASCS53 (ocf::heartbeat:SAPInstance): Stopped (unmanaged)
Resource Group: grp_ZAS_ERS63
rsc_ip_ZAS_ERS63 (ocf::heartbeat:IPAddr2): Stopped (unmanaged)
rsc_fs_ZAS_ERS63 (ocf::heartbeat:Filesystem): Stopped (unmanaged)
rsc_sap_ZAS_ERS63 (ocf::heartbeat:SAPInstance): Stopped (unmanaged)
Resource Group: grp_VAS_ASCS20
rsc_ip_VAS_ASCS20 (ocf::heartbeat:IPAddr2): Stopped (unmanaged)
rsc_fs_VAS_ASCS20 (ocf::heartbeat:Filesystem): Stopped (unmanaged)
rsc_sap_VAS_ASCS20 (ocf::heartbeat:SAPInstance): Stopped (unmanaged)
Resource Group: grp_VAS_ERS30
rsc_ip_VAS_ERS30 (ocf::heartbeat:IPAddr2): Stopped (unmanaged)
rsc_fs_VAS_ERS30 (ocf::heartbeat:Filesystem): Stopped (unmanaged)
rsc_sap_VAS_ERS30 (ocf::heartbeat:SAPInstance): Stopped (unmanaged)

```

As **user root**, release the cluster from maintenance and check that the new cluster resources became active.

```

# crm configure property maintenance-mode=false
# crm_mon -rfn

```

As **user vasadm**, for example on ASCS host, check the SAP system.

```

# su - vasadm
# sapcontrol -nr 20 -function GetSystemInstanceList
# sapcontrol -nr 20 -function GetProcessList -host sapascs2
# sapcontrol -nr 30 -function GetProcessList -host sapers2

```



Note

Repeat these steps for each SID you have installed.

EXAMPLE 5: EXAMPLE OUTPUT OF FIVE RUNNING SID IN A TWO-NODE CLUSTER

As user **root**, use *crm_mon* or *crm status*.

```
# crm_mon -lrfn

Stack: corosync
Current DC: freki (version 1.1.18+20180430.b12c320f5-3.3.1-b12c320f5) - partition
with quorum
Last updated: Tue Oct 15 16:26:57 2019
Last change: Tue Oct 15 16:25:55 2019 by root via cibadmin on freki

2 nodes configured
31 resources configured

Node freki: online
    rsc_ip_WAS_ASCS31      (ocf::heartbeat:IPaddr2):      Started
    rsc_fs_WAS_ASCS31      (ocf::heartbeat:Filesystem):    Started
    rsc_sap_WAS_ASCS31     (ocf::heartbeat:SAPInstance):   Started
    rsc_ip_VAS_ERS30       (ocf::heartbeat:IPaddr2):      Started
    rsc_fs_VAS_ERS30       (ocf::heartbeat:Filesystem):    Started
    rsc_sap_VAS_ERS30      (ocf::heartbeat:SAPInstance):   Started
    rsc_ip_YAS_ERS10       (ocf::heartbeat:IPaddr2):      Started
    rsc_fs_YAS_ERS10       (ocf::heartbeat:Filesystem):    Started
    rsc_sap_YAS_ERS10      (ocf::heartbeat:SAPInstance):   Started
    rsc_ip_ZAS_ERS63       (ocf::heartbeat:IPaddr2):      Started
    rsc_fs_ZAS_ERS63       (ocf::heartbeat:Filesystem):    Started
    rsc_sap_ZAS_ERS63      (ocf::heartbeat:SAPInstance):   Started
    rsc_ip_XAS_ERS52       (ocf::heartbeat:IPaddr2):      Started
    rsc_fs_XAS_ERS52       (ocf::heartbeat:Filesystem):    Started
    rsc_sap_XAS_ERS52      (ocf::heartbeat:SAPInstance):   Started
Node geri: online
    rsc_ip_WAS_ERS41       (ocf::heartbeat:IPaddr2):      Started
    rsc_fs_WAS_ERS41       (ocf::heartbeat:Filesystem):    Started
    rsc_sap_WAS_ERS41      (ocf::heartbeat:SAPInstance):   Started
    rsc_ip_YAS_ASCS00      (ocf::heartbeat:IPaddr2):      Started
    rsc_fs_YAS_ASCS00      (ocf::heartbeat:Filesystem):    Started
    rsc_sap_YAS_ASCS00     (ocf::heartbeat:SAPInstance):   Started
    rsc_ip_XAS_ASCS42      (ocf::heartbeat:IPaddr2):      Started
    rsc_fs_XAS_ASCS42      (ocf::heartbeat:Filesystem):    Started
    rsc_sap_XAS_ASCS42     (ocf::heartbeat:SAPInstance):   Started
    rsc_ip_VAS_ASCS20      (ocf::heartbeat:IPaddr2):      Started
    rsc_fs_VAS_ASCS20      (ocf::heartbeat:Filesystem):    Started
```

```

rsc_sap_VAS_ASCS20      (ocf::heartbeat:SAPInstance): Started
rsc_ip_ZAS_ASCS53      (ocf::heartbeat:IPAddr2):      Started
rsc_fs_ZAS_ASCS53      (ocf::heartbeat:Filesystem):    Started
rsc_sap_ZAS_ASCS53      (ocf::heartbeat:SAPInstance):  Started
rsc_ip_hawk             (ocf::heartbeat:IPAddr2):      Started

```

No inactive resources

Migration Summary:

```

* Node freki:
* Node geri:

```

Each group consists of three resources, *rsc_ip*, *rsc_fs* and *rsc_sap*. The resource *rsc_ip_hawk* is the optional virtual IP for the HAWK Web interface.

3.5 Example ENSA2 in a Two-Node Cluster

As known from the setup based on the **SUSE Best Practice** document for SAP S/4HANA (<https://documentation.suse.com/sbp/sap> → SAP S/4HANA - Enqueue Replication 2 High Availability Cluster - Setup Guide), the cluster integration can be done with a configuration file. Create one configuration file for each SID and load them one by one into the cluster. The following description will help to do the steps in the right direction.

The cluster is already up and running and still runs one ASCS / ERS configuration based on the group concept. The following tasks must be executed:

- Stop the new installed SAP system, which will be added to the cluster.
- Unmount the file system with *umount* and remove the IP address from the new SAP system.
- Check and/or modify the ASCS and ERS profile files and add the section for the sap-suse-cluster-connector integration, if not done already.
 - for example for VAS: */sapmnt/VAS/profile/VAS_ASCS20_sapascs2*
 - for example for VAS: */sapmnt/VAS/profile/VAS_ERS30_sapers2*
- Modify the *haclient* group on each cluster node and add the new *<sid> adm* to it, if not done already.
- Set the cluster into maintenance mode.
- Load the new configuration with *crm configure load update <filename>*.

- Pre-check the cluster with *crm status*.
- Release the cluster from maintenance mode.
- Check the cluster and SAP instances with *crm_mon* and *sapcontrol -nr*

EXAMPLE 6: PREPARE A CONFIGURATION FILE AND EXTEND THE CLUSTER

As user **root** prepare a file for the SAP system VAS

```
# cat ensa2-2nd-vas.txt

primitive rsc_fs_VAS_ASCS20 Filesystem \
  params device="172.17.0.1:/srv/install/sapascs2" directory="/usr/sap/VAS/ASCs20"
  fstype=nfs options="vers=3" \
  op start timeout=60s interval=0 \
  op stop timeout=60s interval=0 \
  op monitor interval=20s timeout=300s
primitive rsc_fs_VAS_ERS30 Filesystem \
  params device="172.17.0.1:/srv/install/sapers2" directory="/usr/sap/VAS/ERS30"
  fstype=nfs options="vers=3" \
  op start timeout=60s interval=0 \
  op stop timeout=60s interval=0 \
  op monitor interval=20s timeout=300s
primitive rsc_ip_VAS_ASCS20 IPAddr2 \
  params ip=172.17.1.17 \
  op monitor interval=10s timeout=20s
primitive rsc_ip_VAS_ERS30 IPAddr2 \
  params ip=172.17.1.18 \
  op monitor interval=10s timeout=20s
primitive rsc_sap_VAS_ASCS20 SAPIInstance \
  operations $id=rsc_sap_VAS_ASCS20-operations \
  op monitor interval=11 timeout=60 on-fail=restart \
  params InstanceName=VAS_ASCS20_sapascs2 START_PROFILE="/sapmnt/VAS/profile/
VAS_ASCS20_sapascs2" AUTOMATIC_RECOVER=false \
  meta resource-stickiness=5000
primitive rsc_sap_VAS_ERS30 SAPIInstance \
  operations $id=rsc_sap_VAS_ERS30-operations \
  op monitor interval=11 timeout=60 on-fail=restart \
  params InstanceName=VAS_ERS30_sapers2 START_PROFILE="/sapmnt/VAS/profile/
VAS_ERS30_sapers2" AUTOMATIC_RECOVER=false IS_ERS=true
group grp_VAS_ASCS20 rsc_ip_VAS_ASCS20 rsc_fs_VAS_ASCS20 rsc_sap_VAS_ASCS20 \
  meta resource-stickiness=3000
group grp_VAS_ERS30 rsc_ip_VAS_ERS30 rsc_fs_VAS_ERS30 rsc_sap_VAS_ERS30
colocation col_sap_VAS_no_both -5000: grp_VAS_ERS30 grp_VAS_ASCS20
order ord_sap_VAS_first_start_ascs Optional: rsc_sap_VAS_ASCS20:start
rsc_sap_VAS_ERS30:stop symmetrical=false
```

As user **root** prepare a file for the SAP system WAS

```
# cat ensa2-3nd-was.txt

primitive rsc_fs_WAS_ASCS31 Filesystem \
  params device="172.17.0.1:/srv/install/sapascs3" directory="/usr/sap/WAS/ASCS31"
  fstype=nfs options="vers=3" \
  op start timeout=60s interval=0 \
  op stop timeout=60s interval=0 \
  op monitor interval=20s timeout=300s
primitive rsc_fs_WAS_ERS41 Filesystem \
  params device="172.17.0.1:/srv/install/sapers3" directory="/usr/sap/WAS/ERS41"
  fstype=nfs options="vers=3" \
  op start timeout=60s interval=0 \
  op stop timeout=60s interval=0 \
  op monitor interval=20s timeout=300s
primitive rsc_ip_WAS_ASCS31 IPAddr2 \
  params ip=172.17.1.27 \
  op monitor interval=10s timeout=20s
primitive rsc_ip_WAS_ERS41 IPAddr2 \
  params ip=172.17.1.28 \
  op monitor interval=10s timeout=20s
primitive rsc_sap_WAS_ASCS31 SAPIInstance \
  operations $id=rsc_sap_WAS_ASCS31-operations \
  op monitor interval=11 timeout=60 on-fail=restart \
  params InstanceName=WAS_ASCS31_sapascs3 START_PROFILE="/sapmnt/WAS/profile/
WAS_ASCS31_sapascs3" AUTOMATIC_RECOVER=false \
  meta resource-stickiness=5000
primitive rsc_sap_WAS_ERS41 SAPIInstance \
  operations $id=rsc_sap_WAS_ERS41-operations \
  op monitor interval=11 timeout=60 on-fail=restart \
  params InstanceName=WAS_ERS41_sapers3 START_PROFILE="/sapmnt/WAS/profile/
WAS_ERS41_sapers3" AUTOMATIC_RECOVER=false IS_ERS=true
group grp_WAS_ASCS31 rsc_ip_WAS_ASCS31 rsc_fs_WAS_ASCS31 rsc_sap_WAS_ASCS31 \
  meta resource-stickiness=3000
group grp_WAS_ERS41 rsc_ip_WAS_ERS41 rsc_fs_WAS_ERS41 rsc_sap_WAS_ERS41
colocation col_sap_WAS_no_both -5000: grp_WAS_ERS41 grp_WAS_ASCS31
order ord_sap_WAS_first_start_ascs Optional: rsc_sap_WAS_ASCS31:start
rsc_sap_WAS_ERS41:stop symmetrical=false
```

EXAMPLE 7: CONFIGURATION EXAMPLE FOR SID XAS AND FOLLOWING

This is similar to the SID examples before. You must adapt the:

- File system sources which will be mounted
- IP address for virtual host name of ASCS and ERS
- Instance number for ASCS and ERS

- Local mount point for ASCS and ERS
- Profile path for ASCS and ERS

EXAMPLE 8: ADD THE 2ND AND FURTHER SID INTO THE CLUSTER

The cluster is already up and running and still run one ASCS / ERS configuration based on the group concept. The following tasks must be executed:

- Stop the new installed SAP system, which will be added to the cluster.
- Unmount the file system with *umount* and remove the IP address from the new SAP system.
- Check and/or modify the ASCS and ERS profile files and add the section for the sap-suse-cluster-connector integration, if not done already.
 - for example for VAS: */sapmnt/VAS/profile/VAS_ASCS20_sapascs2*
 - for example for VAS: */sapmnt/VAS/profile/VAS_ERS30_sapers2*
- Modify the *haclient* group on each cluster node and add the new <sid> adm to it, if not done already.
- Set the cluster into maintenance mode.
- Load the new configuration with *crm configure load update <filename>*.
- Pre-check the cluster with *crm status*.
- Release the cluster from maintenance mode.
- Check the cluster and SAP instances with *crm_mon -lrf* and *sapcontrol -nr*

Login to one of the cluster nodes and run the commands from there.

As user **root** set cluster from maintenance and load the configuration into the cluster

```
# crm configure property maintenance-mode=true
# crm configure load update ensa2-2nd-vas.txt
# crm configure load update ensa2-3nd-was.txt
# crm configure load update ensa2-4nd-xas.txt
...
# crm status
```

You should get back the prompt only. No further messages should be shown. If a message is displayed, there might be something wrong in the configuration file.



Note

If a wrong path is used, for example for the profile, this will not be detected during the configuration load. However, it will be shown during the cluster start action.

As user **root** release cluster from maintenance and check the new cluster resources

```
# crm configure property maintenance-mode=false
# crm_mon -rfn
```

As user **vasadm**, for example on ASCS host, check the SAP system

```
# su - vasadm
# sapcontrol -nr 20 -function GetSystemInstanceList
# sapcontrol -nr 20 -function GetProcessList -host sapascs2
# sapcontrol -nr 30 -function GetProcessList -host sapers2
```



Note

Repeat the steps for each SID you have installed.

3.6 Example ENSA1 in a Multi-Node Cluster

This base setup is already documented in the **SUSE Best Practices** document for SAP NetWeaver (<https://documentation.suse.com/sbp/sap> → SAP S/4HANA - SAP NetWeaver High Availability Cluster 7.40 - Setup Guide - Best Practice Guide). The OS and node preparation must be done for each future cluster member:

- patch level
- settings for date and time, DNS
- access to the same fencing device (method)

The major difference compared to the two-node setup is: run the `ha-cluster-join` command multiple time until all nodes are member of the cluster.

The cluster integration can be done with a configuration file. Create one configuration file for each SID and load them one by one into the cluster. The steps are equal to the steps described in [Section 3.4, “Example ENSA1 in a Two-Node Cluster”](#).

3.7 Example ENSA2 in a Multi-Node Cluster

This base setup is already documented in the **SUSE Best Practices** document for SAP S/4HANA (<https://documentation.suse.com/sbp/sap> → SAP S/4HANA - Enqueue Replication 2 High Availability Cluster - Best Practice Guide). The OS and node preparation must be done for each future cluster member:

- patch level
- settings for date and time, DNS
- access to the same fencing device (method)

The major difference compared to the two-node setup is: run the *ha-cluster-join* command multiple times until all nodes are member of the cluster.

The cluster integration can be done with a configuration file. Create one configuration file for each SID and load them one by one into the cluster. The steps are equal to the steps described in [Section 3.5, "Example ENSA2 in a Two-Node Cluster"](#).

3.8 Example ENSA1 and ENSA2 in a Two-Node Cluster

As known from the setup based on the **SUSE Best Practices** document for SAP NetWeaver and SAP S/4HANA (<https://documentation.suse.com/sbp/sap> → SAP NetWeaver High Availability Cluster 7.40 - Setup Guide) (<https://documentation.suse.com/sbp/sap> → SAP S/4HANA - Enqueue Replication 2 High Availability Cluster - Setup Guide), the cluster integration can be done with a configuration file. Create one configuration file for each SID and load them one by one into the cluster.

We expect the cluster is already up and running and still runs one ASCS / ERS configuration based on the group concept.



Note

Follow and combine the chapters "Example ENSA1 in a Two-Node Cluster" and "Example ENSA2 in a Two-Node Cluster". Be careful and keep in mind each SAP system must be independent from each other.

[Section 3.4, "Example ENSA1 in a Two-Node Cluster"](#)

[Section 3.5, "Example ENSA2 in a Two-Node Cluster"](#)

3.9 Example ENSA1 and ENSA2 in a Multi-Node Cluster

The setup here is similar to the one described in the chapter [Section 3.8, “Example ENSA1 and ENSA2 in a Two-Node Cluster”](#) but only with more nodes in the cluster.

After the basic cluster part is done, the procedure is the same as for a two-node cluster setup. Check the following chapter for further instructions: [Section 3.8, “Example ENSA1 and ENSA2 in a Two-Node Cluster”](#).

3.10 Workload Balancing of Central Services



Note

Under evaluation.

3.10.1 Automatic Load Balancing by Pacemaker

There is no additional task necessary.

3.10.2 Manual Load Balancing by Pacemaker



Note

Under evaluation.

3.11 Testing the New Configuration of Multiple SID in One Cluster



Important

A well-defined and overall test of the new configuration is extremely recommended. Read the requirements for each setup carefully. The test procedure will demonstrate if the cluster works as expected. We already have described multiple test scenarios in our base Best Practices documentation for SAP NetWeaver and SAP S/4HANA.

4 Multi-SID - Database Services




Note

This is currently under evaluation. Send an e-mail to saphana@suse.com (<mailto:saphana@suse.com>) [↗](#) if this might be from interest for you.

5 Legal Notice

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

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

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

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

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

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

6 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.