

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

Setup Guide

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

Jinhui Li, Product Manager SAP Solutions on Alibaba Cloud (AliCloud)
Alex Li, Staff Solution Architect on Alibaba Cloud (AliCloud)
Chen Dong, Staff Solution Architect on Alibaba Cloud (AliCloud)
Fabian Herschel, Distinguished Architect SAP (SUSE)
Bernd Schubert, SAP Solution Architect (SUSE)

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

Setup Guide

Date: 2025-03-06

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

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

Contents

- 1 About this guide 4
- 2 Scope of this document 5
- 3 Overview 6
- 4 Infrastructure preparation 12
- 5 SAP installation 22
- 6 Implementing the cluster 30
- 7 Administration 36
- 8 Additional implementation scenarios 42
- 9 References 69
- 10 Appendix 73
- 11 Legal notice 77
- 12 GNU Free Documentation License 78

1 About this guide

1.1 Introduction

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

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

The described cluster solution can be used for SAP* SAP S/4HANA and for SAP* SAP NetWeaver. SAP NetWeaver is a common stack of middleware functionality used to support the SAP business applications. The SAP Enqueue Replication Server constitutes application level redundancy for one of the most crucial components of the SAP NetWeaver stack, the enqueue service. An optimal effect of the enqueue replication mechanism can be achieved when combining the application level redundancy with a high availability cluster solution as provided with SUSE Linux Enterprise Server for SAP Applications. The described concept has proven its maturity over several years of productive operations for customers of different sizes and branches.

1.2 Additional documentation and resources

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

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

Numerous whitepapers, a best practices guide, and other resources are provided at the SUSE Linux Enterprise Server for SAP Applications resource library: <https://www.suse.com/products/sles-for-sap/#resource>.

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

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

1.3 Feedback

Several feedback channels are available:

Bugs and Enhancement Requests

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

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

Mail

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

2 Scope of this document

This guide details how to:

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

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

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

3 Overview

This guide describes how to set up a pacemaker cluster using SUSE Linux Enterprise Server for SAP Applications 15 for the Enqueue Replication scenario on Alibaba Cloud. The goal is to match the SAP NW-HA-CLU 7.50 (former version 7.40) certification specifications and goals.

These goals include:

- Integration of the cluster with the SAP start framework *sapstartsrv* to ensure that maintenance procedures do not break the cluster stability
- Rolling Kernel Switch (RKS) awareness
- Standard SAP installation to improve support processes

The updated certification SAP NW-HA-CLU 7.50 (former version 7.40) has redefined some of the test procedures and described new expectations how the cluster should behave in special conditions. These changes allowed us to improve the cluster architecture and to design it for easier usage and setup.

Shared SAP resources are on a central NFS server.

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

3.1 Differences to previous cluster architectures

The concept is different to the old stack with the master-slave architecture. With the new certification we switch to a more simple model with primitives. This means we have on one machine the ASCS with its own resources and on the other machine the ERS with its own resources.

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

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



Note

The cluster in this guide only manages the SAP instances ASCS and ERS, because of the focus of the SAP NW-HA-CLU 7.50 (former version 7.40) certification.

If your database is SAP HANA, we recommend to set up the performance-optimized system replication scenario using our automation solution SAPHanaSR. The SAPHanaSR automation should be set up in an own two node cluster. The setup is described in a separate best practices document available at <http://documentation.suse.com/sbp/sap>. In case of using ASE database together with an HADR setup, there is an example in this document.

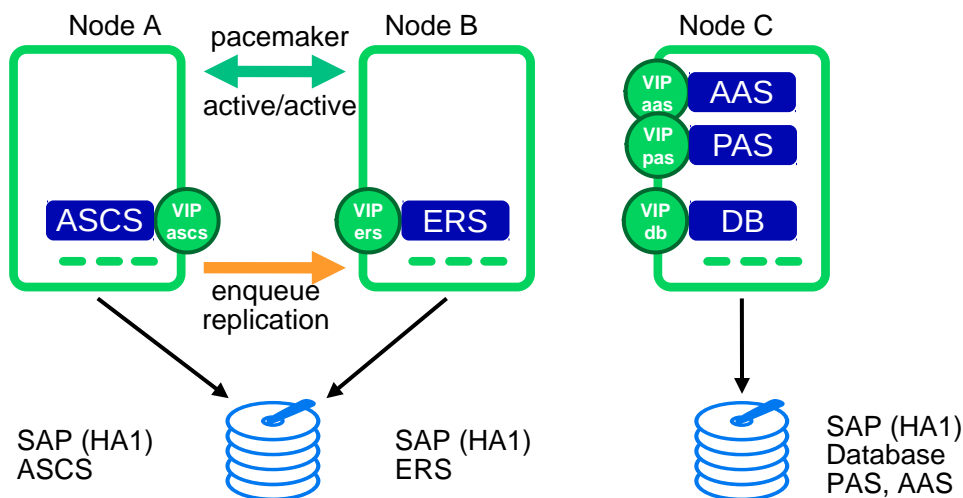


FIGURE 1: THREE SYSTEMS FOR THE CERTIFICATION SETUP

CLUSTERED MACHINES

- one machine (sapapp1) for ASCS
 - Host name: vsapascs
- one machine (sapapp2) for ERS
 - Host name: vsapers

NON-CLUSTERED MACHINE

- one machine (sapdb1) for DB and DI

3.3 High availability for the database

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

3.3.1 SAP HANA system replication

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

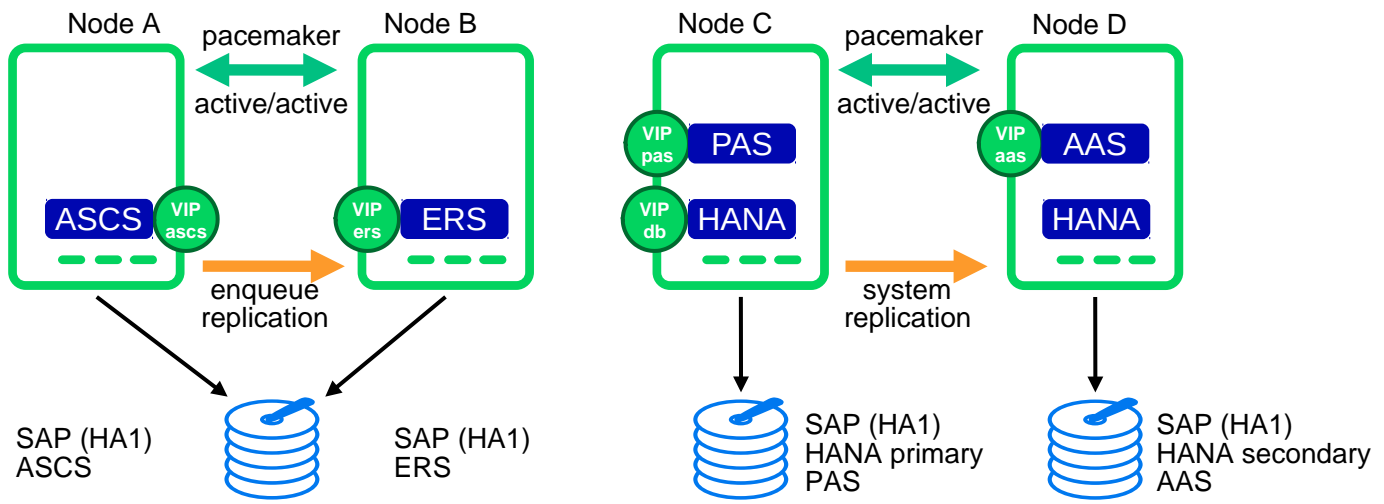


FIGURE 2: ONE CLUSTER FOR CENTRAL SERVICES, ONE FOR SAP HANA SR

The following Databases are supported in combination with this scenario:

- SAP HANA DATABASE 1.0
- SAP HANA DATABASE 2.0

3.3.2 ASE database replication

The picture below shows a solution for an ASE HADR setup. The ASE has its own HA mechanism which is managed by Fault Manager. The Fault Manager itself is a single point of failure. The implementation as integrated service or as a separate SAP instance in the Pacemaker cluster for the central services solves this weakness.




FIGURE 3: ONE CLUSTER FOR THE CENTRAL SERVICES AND THE ASE DATABASE HADR SOLUTION

The following databases are supported in combination with this scenario:

- ASE16 SP03 PL07 onwards

3.3.3 Simple stack

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





FIGURE 4: ONE CLUSTER FOR THE CENTRAL SERVICES AND ONE CLUSTER FOR THE ANY DATABASE

TABLE 1: THE FOLLOWING OS / DATABASES COMBINATION ARE EXAMPLES FOR THIS SCENARIO

SUSE Linux Enterprise Server for SAP Applications 15	
Intel X86_64	POWER LITTLE ENDIAN
SAP HANA DATABASE 1.0	
SAP HANA DATABASE 2.0	SAP HANA DATABASE 2.0
DB2 FOR LUW 11.5	
MaxDB 7.9	
ORACLE 12.2	
SAP ASE 16.0 FOR BUS. SUITE	



Note

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

3.4 Integration of SAP NetWeaver into the cluster using the Cluster Connector

The integration of the HA cluster through the SAP control framework using the `sap_suse_cluster_connector` is of special interest. The `sapstartsrv` controls SAP instances since SAP Kernel versions 6.40. One of the classical problems running SAP instances in a highly available environment is the following:

If an SAP administrator changes the status (start/stop) of an SAP instance without using the interfaces provided by the cluster software, the cluster framework will detect that as an error status. Therefore it will bring the SAP instance into the old status by either starting or stopping the SAP instance. This can result in very dangerous situations if the cluster changes the status of an SAP instance during some SAP maintenance tasks. This new updated solution enables the central component `sapstartsrv` to report state changes to the cluster software, and therefore avoids the previously described dangerous situations. (See also blog article "Using `sap_vendor_cluster_connector` for interaction between cluster framework and `sapstartsrv`") (<https://blogs.sap.com/2014/05/08/using-sapvendorclusterconnector-for-interaction-between-cluster-framework-and-sapstartsrv/comment-page-1/>).

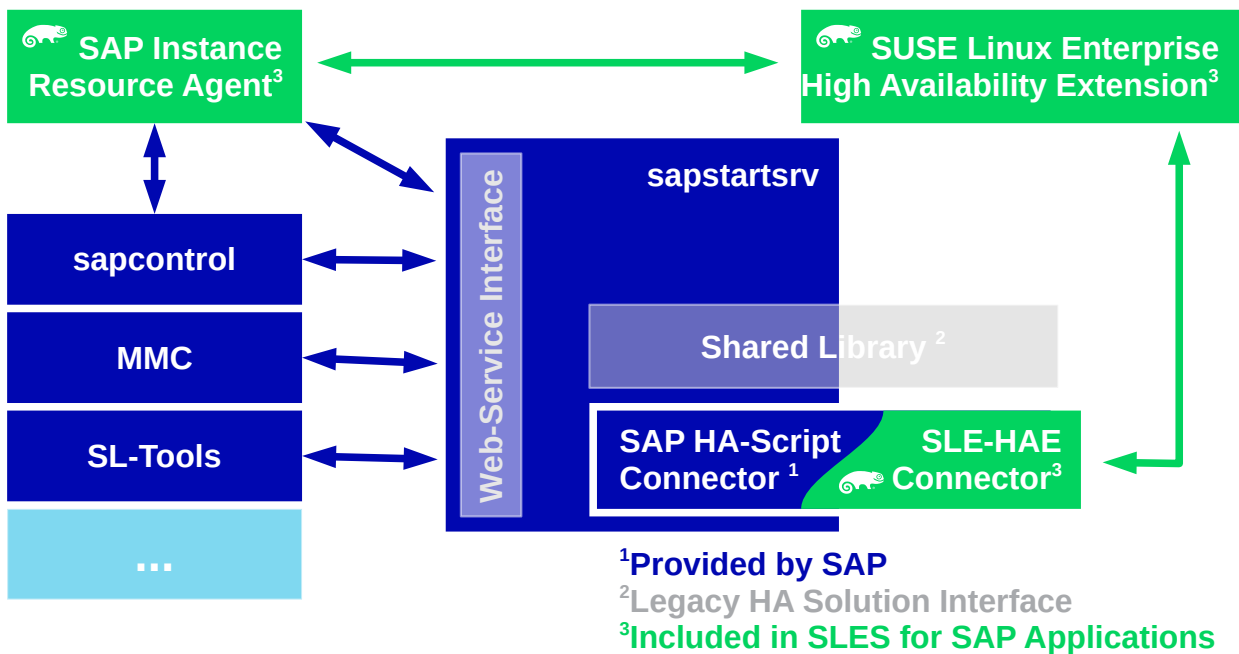


FIGURE 5: CLUSTER CONNECTOR TO INTEGRATE THE CLUSTER WITH THE SAP START FRAMEWORK



Note

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

The new version of the `sap-suse-cluster-connector` now allows to start, stop and 'move' an SAP instance. The integration between the cluster software and the `sapstartsrv` also implements the option to run checks of the HA setup using either the command line tool `sapcontrol` or the SAP management consoles (SAP MMC or SAP MC).

4 Infrastructure preparation

The next sections contain information about how to prepare your Alibaba Cloud infrastructure.


4.1 Infrastructure list

To set up your infrastructure, the following components are required:

- 1 Virtual Private Cloud (VPC) network; [Virtual Private Cloud \(https://www.alibabacloud.com/product/vpc\)](https://www.alibabacloud.com/product/vpc)
- 4 Elastic Compute Service (ECS) instances in different zones of the same VPC; [Elastic Compute Service \(https://www.alibabacloud.com/product/ecs\)](https://www.alibabacloud.com/product/ecs)
- Alibaba Cloud specific Virtual IP Resource Agent (<https://github.com/ClusterLabs/resource-agents/blob/master/heartbeat/aliyun-vpc-move-ip>) and STONITH device (https://github.com/ClusterLabs/fence-agents/blob/master/agents/aliyun/fence_aliyun.py);


4.2 Creating VPC

First, create a VPC via **Console**→**Virtual Private Cloud**→**VPCs**→**Create VPC**. In this example, a VPC named **ASEHADR** in the Region **EU Central 1 (Frankfurt)** has been created:



There should be at least two VSwitches (subnets) defined within the VPC network. Each VSwitch should be bound to a different Zone. In this example, the following two VSwitches (subnets) are defined:

- Switch1 "vs_a" 192.168.1.0/24 Zone A, for SAP HANA Primary Node;
- Switch2 "vs_b" 192.168.2.0/24 Zone B, for SAP HANA Secondary Node;




4.3 Creating ECS instances

Four ECS instances are created in different Zones of the same VPC via **Console**→**Elastic Compute Service ECS**→**Instances**→**Create Instance**.

Choose the "SUSE Linux Enterprise Server for SAP Applications" image from the Image Market place.

In this example, **two** ECS instances (host name: sapapp1 and sapdb1) are created in the Region **EU Central 1**, Zone A, within the VPC: **ASEHADR**, with SUSE Linux Enterprise Server for SAP Applications 15 SP1 image from the Image Market Place. In contrast, **two** ECS instances (host name: sapapp2 and sapdb2) are created in the Region **EU Central 1**, Zone B, within the VPC: **ASEHADR**, with SUSE Linux Enterprise Server for SAP Applications 15 SP1 image from the Image Market Place.



4.4 Confirming OpenApi endpoint address inside VPC

An Alibaba Cloud specific STONITH device and Virtual IP Resource Agent are mandatory for the cluster. Traditionally, these components need to access Alibaba Cloud OpenAPI through a public domain, which used to be implemented by configuring NAT Gateway and corresponding SNAT entries, or Private Zone. Nowadays, Alibaba Cloud OpenAPI is accessible inside VPC with specific endpoint address. In Region EU Central 1 (Frankfurt) the endpoint addresses are:

- ECS: ecs-vpc.eu-central-1.aliyuncs.com
- VPC: vpc-vpc.eu-central-1.aliyuncs.com

For other regions, refer to the following Alibaba Cloud documents to find the corresponding endpoint addresses:

- ECS: <https://www.alibabacloud.com/help/doc-detail/117461.htm>
- VPC: <https://www.alibabacloud.com/help/doc-detail/185725.htm>

4.5 Creating STONITH device and virtual IP ResourceAgent

To download and install the components of the following steps, the instance needs to be able to access Internet. The easiest way is to purchase an Elastic IP (<https://www.alibabacloud.com/product/eip>) and assign it to the instance and unassign it when the configuration is done. For an HA solution, a fencing device is an essential requirement. Alibaba Cloud provides its own STONITH device, which allows the servers in the HA cluster to shut down the node that is not responsive. The STONITH device leverages Alibaba Cloud OpenAPI underneath the ECS instance, which is similar to a physical reset / shutdown in an on-premise environment.

```
# curl https://raw.githubusercontent.com/ClusterLabs/fence-agents/master/agents/aliyun/fence_aliyun.py > /usr/sbin/fence_aliyun
```

```

## Add permission
# chmod 755 /usr/sbin/fence_aliyun
# chown root:root /usr/sbin/fence_aliyun
## set python
# sed -i "s|@PYTHON@|$(which python)|" /usr/sbin/fence_aliyun
## set Fence agents lib directory
# sed -i "s|@FENCEAGENTS_LIBDIR@|/usr/share/fence|" /usr/sbin/fence_aliyun
## Installation verification
# stonith_admin -I |grep fence_aliyun
## return fence_aliyun as correct
100 devices found
# fence_aliyun

```

The next component to install is Virtual IP Resource Agent (`aliyun-vpc-move-ip`). By changing the routing entries, it enables a non-overlapping, private IP addresses to be used as a virtual IP resources in an HA solution.

```

# mkdir -p /usr/lib/ocf/resource.d/aliyun
# curl https://raw.githubusercontent.com/ClusterLabs/resource-agents/master/heartbeat/aliyun-vpc-move-ip > /usr/lib/ocf/resource.d/aliyun/vpc-move-ip
# chmod 755 /usr/lib/ocf/resource.d/aliyun/vpc-move-ip
# chown root:root /usr/lib/ocf/resource.d/aliyun/vpc-move-ip

```

Install Alibaba Cloud OpenAPI SDK:

```

# pip install aliyun-python-sdk-core aliyun-python-sdk-vpc aliyun-python-sdk-ecs

```

Install Alibaba Cloud CLI:

```

# wget https://github.com/aliyun/aliyun-cli/releases/download/v3.0.65/aliyun-cli-linux-3.0.65-amd64.tgz
# tar -xvf aliyun-cli-linux-3.0.65-amd64.tgz
# mv aliyun /usr/local/bin

```

Configure Alibaba Cloud RAM role to authenticate Alibaba Cloud CLI (`aliyun`) to create OpenAPI operation call.

Login Alibaba Cloud console → Resource Access Management (RAM) → Permissions → Policies
 Create a customer policy named “SAP-HA-ROLE-POLICY” with the following content:

```

{
  "Version": "1",
  "Statement": [
    {
      "Effect": "Allow",

```

```

    "Action": [
      "ecs:StartInstance",
      "ecs:StopInstance",
      "ecs:RebootInstance",
      "ecs:Describe*"
    ],
    "Resource": [
      "*"
    ],
    "Condition": {}
  },
  {
    "Effect": "Allow",
    "Action": [
      "vpc:CreateRouteEntry",
      "vpc>DeleteRouteEntry",
      "vpc:Describe*"
    ],
    "Resource": [
      "*"
    ],
    "Condition": {}
  }
]
}

```

RAM / Policies / Create Custom Policy

← **Create Custom Policy**

* Policy Name
SAP-HA-ROLE-POLICY

Note

Configuration Mode
 Visualized
 Script

Policy Document
 Import an existing system policy

```

1  {
2    "Version": "1",
3    "Statement": [
4      {
5        "Effect": "Allow",
6        "Action": [
7          "ecs:StartInstance",
8          "ecs:StopInstance",
9          "ecs:RebootInstance",
10         "ecs:Describe*"
11        ],
12        "Resource": [
13          "*"
14        ],
15        "Condition": {}

```

Ok Return

Create a RAM role “SAP-HA-ROLE” and assign above RAM policy to it.


```
# mount 114b194b126-tkt51.eu-central-1.nas.aliyuncs.com:/ /usr/sap/SSA/ASCS00
```

Afterward, execute the below command to mount the created NAS storage to sapapp2:

```
# mkdir /usr/sap/SSA/ERS10
# mount 11e5134833f-dlk6.eu-central-1.nas.aliyuncs.com:/ /usr/sap/SSA/ERS10
```

The mount points are like this:

- sapapp1:
 - 114b194b126-tkt51.eu-central-1.nas.aliyuncs.com:/ /usr/sap/SSA/ASCS00
- sapapp2:
 - 11e5134833f-dlk6.eu-central-1.nas.aliyuncs.com:/ /usr/sap/SSA/ERS10

4.6.2 Disk for DB and dialog instances (ASE DB example)

The disk for the database and primary application server is assigned to sapdb1. In an advanced setup this disk should be shared between sapdb1 and an optional additional node building an own cluster.

TO BE MOUNTED EITHER BY OS OR AN OPTIONAL CLUSTER

- sapdb1: /dev/vdc /sybase/SSA/srsdata
- sapdb1: /dev/sdb3 /usr/sap/SSA/DVEBMGS01
- sapdb1: /dev/sdb4 /usr/sap/SSA/D02



Note

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

In our example, we use block storage for database storage. Execute the below command on node sapdb1:

```
# mkdir /sybase/SSA/srsdata
# echo /dev/vdc /sybase/SSA/srsdata ext4 acl,user_xattr,noatime 1 1 >> /etc/fstab
```

Installation Media. The installation media are normally store in a central place which can be mounted from all node which need the software. We normally mount this share or export to `/sapcd`.

4.7 IP addresses and virtual names

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

```
192.168.1.123 sapapp1
192.168.2.92 sapapp2
192.168.1.124 sapdb1
192.168.4.11 vsapascs
192.168.5.11 vsapers
```

4.8 Mount points and NFS shares

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

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

- Creating mount points and mounting NFS share on all nodes

EXAMPLE 1: SAP NETWEAVER 7.5

```
# mkdir -p /sapmnt
# mkdir -p /usr/sap/SSA/{ASCS00,D01,D02,ERS10,SYS}
# mount -t nfs :1189b8484a0-pwo18.eu-central-1.nas.aliyuncs.com:/ /sapmnt
# mount -t nfs :11c2284a122-fjk10.eu-central-1.nas.aliyuncs.com:/ /usr/sap/SSA/SYS
```

- Only ASEDB: creating mount points for the database at `sapdb1`:

```
# mkdir -p /sybase/SSA/srsdata
```

- Only HANA: creating mount points for database at `sapdb1`:

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

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

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




FIGURE 6: FILE SYSTEM LAYOUT INCLUDING NFS SHARES

We prepare the three servers for the distributed SAP installation. Server 1 (sapapp1) will be used to install the ASCS SAP instance. Server 2 (sapapp2) will be used to install the ERS SAP instance. Server 3 (sapdb1) will be used to install the dialog SAP instances and the database.

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

EXAMPLE 2: SAP NETWEAVER 7.50 ON X86_64 ARCHITECTURE WITH ASEDB

```
(ASCS  sapapp1) # mount 114b194b126-tkt51.eu-central-1.nas.aliyuncs.com:/ /usr/sap/SSA/
ASCS00
(ERS   sapapp2) # mount 11e5134833f-dlk6.eu-central-1.nas.aliyuncs.com:/ /usr/sap/SSA/
ERS10
(DB    sapdb1) # mount /dev/vdc /sybase/SSA/srsdata
(Dialog sapdb1) # mount /dev/sdb3 /usr/sap/SSA/D01
(Dialog sapdb1) # mount /dev/sdb4 /usr/sap/SSA/D02
```

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

```
# ls -la /usr/sap/SSA/
total 0
```

```

drwxr-xr-x 1 ssaadm sapsys 70 28. Mar 17:26 ./
drwxr-xr-x 1 root   sapsys 58 28. Mar 16:49 ../
drwxr-xr-x 7 ssaadm sapsys 58 28. Mar 16:49 ASCS00/
drwxr-xr-x 1 ssaadm sapsys  0 28. Mar 15:59 D02/
drwxr-xr-x 1 ssaadm sapsys  0 28. Mar 15:59 D01/
drwxr-xr-x 1 ssaadm sapsys  0 28. Mar 15:59 ERS10/
drwxr-xr-x 5 ssaadm sapsys 87 28. Mar 17:21 SYS/

```



Note

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

5 SAP installation

The overall procedure to install the distributed SAP is:

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

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




FIGURE 7: DISTRIBUTED INSTALLATION OF THE SAP SYSTEM

5.1 Linux user and group number scheme

Whenever asked by the SAP software provisioning manager (SWPM) which Linux User IDs or Group IDs to use, refer to the following table which is, of course, only an example.

```
Group sapinst      1000
Group sapsys       1001
Group sapadm       3000
Group sdba         3002

User  ssaadm       3000
User  sdb          3002
User  sqdssa       3003
User  sapadm       3004
```



Note

Adapt the value as you like. These are examples and may not fit into your company policy.

5.2 Installing ASCS on sapapp1

Temporarily we need to set the service IP address used later in the cluster as local IP, because the installer wants to resolve or use it. Make sure to use the right virtual host name for each installation step. Take care of the ASCS file systems like 114b194b126-tkt51.eu-central-1.nas.aliyuncs.com:/ and /sapcd/ (where the installation sources live) which might also need to be mounted.

```
# ip a a 192.168.4.11/24 dev eth0
# mount 114b194b126-tkt51.eu-central-1.nas.aliyuncs.com:/ /usr/sap/SSA/ASCS00
# cd /sapcd/SWPM/
# ./sapinst SAPINST_USE_HOSTNAME=vsapascs
```

- SWPM option depends on SAP NetWeaver version and architecture
 - Installing SAP NetWeaver AS for ABAP 7.52 → SAP ASE → Installation → Application Server ABAP → High-Availability System → ASCS Instance
- SID id SSA
- Use instance number 00
- Deselect using FQDN

- All passwords: use <yourSecurePwd>
- Double-check during the parameter review, if virtual name **vsapscs** is used



Note

Adapt the values, for example SID ID, instance number, virtual host name, etc. These are examples and may not fit into your company policy.

5.3 Installing ERS on sapapp2

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

```
# ip a a 192.168.5.11/24 dev eth0
# mount 11e5134833f-dlk6.eu-central-1.nas.aliyuncs.com:/ /usr/sap/SSA/ERS10
# cd /sapcd/SWPM/
# ./sapinst SAPINST_USE_HOSTNAME=vsapers
```

- SWPM option depends on SAP NetWeaver version and architecture
 - Installing SAP NetWeaver AS for ABAP 7.52 → SAP ASE → Installation → Application Server ABAP → High-Availability System → Enqueue Replication Server Instance
- Use instance number 10
- Deselect using FQDN
- Double-check during the parameter review if virtual name **vsapers** is used
- If you get an error during the installation about permissions, change the ownership of the ERS directory



Note

Adapt the values, for example SID ID, instance number, virtual host name, etc. These are examples and may not fit into your company policy.

```
# chown -R ssaadm:sapsys /usr/sap/SSA/ERS10
```


- If you get a prompt to manually stop/start the ASCS instance, log in to sapapp1 as user ssaadm and call `sapcontrol`.

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

5.4 Post-steps for ASCS and ERS

5.4.1 Stopping ASCS and ERS

On sapapp1

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

On sapapp2

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

5.4.2 Maintaining sapservices

Ensure `/usr/sap/sapservices` hold both entries (ASCS+ERS) on both cluster nodes. This allows the `sapstartsrv` clients to start the service like (do not execute this at this point in time).

As user ssaadm, do:

```
# sapcontrol -nr 10 -function StartService SSA
```

The `/usr/sap/sapservices` looks like (typically one line per instance):

```
#!/bin/sh
LD_LIBRARY_PATH=/usr/sap/SSA/ASCS00/exe:$LD_LIBRARY_PATH; export LD_LIBRARY_PATH; /usr/
sap/SSA/ASCS00/exe/sapstartsrv pf=/usr/sap/SSA/SYS/profile/SSA_ASCS00_vsapascs -D -u
ssaadm
LD_LIBRARY_PATH=/usr/sap/SSA/ERS10/exe:$LD_LIBRARY_PATH; export LD_LIBRARY_PATH; /usr/
sap/SSA/ERS10/exe/sapstartsrv pf=/usr/sap/SSA/ERS10/profile/SSA_ERS10_vsapers -D -u
ssaadm
```

5.4.3 Integrating the cluster framework using sap-suse-cluster-connector

Install the package **sap-suse-cluster-connector** version 3.1.x from our repositories on **both** cluster nodes:

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



Note

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

For the ERS and ASCS instance edit the instance profiles `SSA_ASCS00_vsapasc` and `SSA_ERS10_vsapers` in the profile directory `/usr/sap/SSA/SYS/profile/`.

You need to tell the `sapstartsrv` to load the HA script connector library and to use the `sap-suse-cluster-connector`.

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

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

```
# usermod -aG haclient ssaadm
```

5.4.4 Adapting SAP profiles to match the SAP NW-HA-CLU 7.50 (former version 7.40) certification

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

EXAMPLE 3: FILE `/USR/SAP/SSA/SYS/PROFILE/SSA_ASCS00_VSAPASCS`

```
Start_Program_01 = local $_EN pf=$_PF
```

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

For the ERS change instance the start command from *Restart_Program_xx* to *Start_Program_xx* for the enqueue replication server (enrepserver).

EXAMPLE 4: FILE /USR/SAP/SSA/SYS/PROFILE/SSA_ERS10_VSAPERS

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

5.4.5 Starting ASCS and ERS

On sapapp1

```
# su - ssaadm  
# sapcontrol -nr 00 -function StartService SSA  
# sapcontrol -nr 00 -function Start
```

On sapapp2

```
# su - ssaadm  
# sapcontrol -nr 10 -function StartService SSA  
# sapcontrol -nr 10 -function Start
```

5.5 Installing DB on sapdb1 (example MaxDB)

The MaxDB needs min.40 GB. We use /dev/vdc and mount the partition to /*sapdb*.

A detailed description can be found [here \(https://documentation.suse.com/sbp/all/single-html/SAP_NW740_SLE12_SetupGuide/#_installing_db_on_hacert03_example_maxdb\)](https://documentation.suse.com/sbp/all/single-html/SAP_NW740_SLE12_SetupGuide/#_installing_db_on_hacert03_example_maxdb)

5.6 Installing DB on sapdb1 (example SAP HANA)

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

A detailed description can be found at [Example HANA DB \(https://documentation.suse.com/sbp/all/single-html/SAP_NW740_SLE12_SetupGuide/#_installing_db_on_hacert03_example_sap_hana\)](https://documentation.suse.com/sbp/all/single-html/SAP_NW740_SLE12_SetupGuide/#_installing_db_on_hacert03_example_sap_hana).

5.7 Installing DB on sapdb1 (example ASE DB)

The storage sizing depends on many indicators. Check the sizing recommendations for the planned use case.

```
# cd /<path to the SWPM>/  
# ./sapinst
```

- We are installing SAP NetWeaver AS for ABAP 7.52 → SAP ASE → Installation → Application Server → ABAP → High Availability System → Database Instance
- Profile directory /sapmnt/SSA/profile
- Master Password: enter your own value
- SAP System Administrator: enter the password from the ASCS / ERS installation
- General SAP System Parameters: Unicode
- Deselect using FQDN
- Operating System User for SAP Database Administration: specify a UID if needed.
- SAP System Administrator: use the correct UID for *sapadm*, must be the same as for ASCS / ERS
- SAP ASE Database System Parameters
 - physical Memory since in MB
- Double-check all values during the parameter review

5.8 Installing the Primary Application Server (PAS) on sapdb1

```
# ip a a 192.168.1.118/24 dev eth0  
# mount /dev/sdb3 /usr/sap/SSA/D01  
# cd /sapcd/SWPM/  
# ./sapinst SAPINST_USE_HOSTNAME=sapssad1
```

- SWPM option depends on SAP NetWeaver version and architecture

- Installing SAP NetWeaver AS for ABAP 7.52 → SAP ASE → Installation → Application Server ABAP → High-Availability System → Primary Application Server Instance (PAS)
- Use instance number 01
- Deselect using FQDN
- For our hands-on setup use a default secure store key
- Do not install Diagnostic Agent
- No SLD
- Double-check during the parameter review if virtual name **sapssad1** is used

5.9 Installing an Additional Application Server (AAS) on sapdb1

```
# ip a a 192.168.1.119/24 dev eth0
# mount /dev/sdb4 /usr/sap/SSA/D02
# cd /sapcd/SWPM/
# ./sapinst SAPINST_USE_HOSTNAME=sapssad2
```

- SWPM option depends on SAP NetWeaver version and architecture
 - Installing SAP NetWeaver AS for ABAP 7.52 → SAP ASE → Installation → Application Server ABAP → High-Availability System → Additional Application Server Instance (AAS)
- Use instance number 02
- Deselect using FQDN
- Do not install Diagnostic Agent
- Double-check during the parameter review if virtual name **sapssad2** is used

6 Implementing the cluster

The main procedure to implement the cluster is as follows:

- Install the cluster software if not already done during the installation of the operating system
- Configure the cluster communication framework corosync
- Configure the cluster resource manager
- Configure the cluster resources



Note

The SBD device/partition need to be created in beforehand. In this setup guide we do not use the SBD device.

TASKS

1. Setup NTP (best with yast2) and enable it
2. Install pattern ha_sles on both cluster nodes

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

6.1 Configuring the cluster base

TASKS

- Install and configure the cluster stack at first machine

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

```
# ha-cluster-init -y -i eth0 -u
```

- Join the second node

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

```
# ha-cluster-join -y -c 192.168.1.123 -i eth0
```

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

```
Last updated: Thu Nov 21 14:25:53 2019  Last change: Thu Nov 21 14:23:21 2019 by ssaadm
via crm_resource on sapapp1
Stack: corosync
Current DC: sapapp1 (version 1.1.19-20181105.ccd6b5b10) - partition with quorum
2 nodes configured

Online: [ sapapp1 sapapp2 ]
```

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

```
# crm configure show
...
property cib-bootstrap-options: \
    have-watchdog=false \
    dc-version="2.0.1+20190417.13d370ca9-3.9.1-2.0.1+20190417.13d370ca9" \
    cluster-infrastructure=corosync \
    cluster-name=hacluster \
    stonith-enabled=true \
    last-lrm-refresh=1494346532
rsc_defaults rsc-options: \
    resource-stickiness=1 \
    migration-threshold=3
op_defaults op-options: \
    timeout=600 \
    record-pending=true
```

6.2 Configuring cluster resources

We need a changed SAPInstance resource agent for SAP NetWeaver to **not** use the master-slave construct anymore. This also implies a move to a more cluster-like construct to start and stop the ASCS and the ERS themselves and **not** only the complete master-slave.

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




FIGURE 8: RESOURCES AND CONSTRAINTS

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

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

6.2.1 Preparing the cluster for adding the resources

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

As user root, do:

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

6.2.2 Configuring the Stonith resources for an Alibaba Cloud infrastructure

Alibaba Cloud provides its own STONITH device, which allows the servers in the HA cluster to shut down the other which is not responsible. The STONITH device leverage Alibaba Cloud OpenAPI underneath the ECS instance, which is similar to a physical reset / shutdown on a on-premise environment."

EXAMPLE 5: ALIBABA CLOUD FENCING AGENT

```
primitive res_ALIYUN_STONITH_1 stonith:fence_aliyun \
```



```

op monitor interval=120 timeout=60 \
params plug=i-gw87xi82sj2dy2ysaw19 ram_role=SAP-HA-ROLE region=eu-central-1 \
meta target-role=Started
primitive res_ALIYUN_STONITH_2 stonith:fence_aliyun \
op monitor interval=120 timeout=60 \
params plug=i-gw86pnh1jy1dw0vfer3w ram_role=SAP-HA-ROLE region=eu-central-1 \
meta target-role=Started

```

EXAMPLE 6: STONITH LOCATION RULES

```

location loc_sapapp1_stonith_not_on_sapapp1 res_ALIYUN_STONITH_1 -inf: sapapp1
location loc_sapapp2_stonith_not_on_sapapp2 res_ALIYUN_STONITH_2 -inf: sapapp2

```

Create a txt file (like `crm_stonith.txt`) with your preferred text editor, enter both examples (primitives and location rules) to that file and load the configuration to the cluster manager configuration.

As user root, do:

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

6.2.3 Configuring the resources for the ASCS

First we configure the resources for the file system, IP address and the SAP instance. Of course you need to adapt the parameters to your environment.

EXAMPLE 7: ASCS PRIMITIVE

```

primitive rsc_fs_SSA_ASCS00 Filesystem \
  params device="114b194b126-tkt51.eu-central-1.nas.aliyuncs.com:/" directory="/usr/sap/SSA/ASC00" \
  fstype=nfs \
  op start timeout=60s interval=0 \
  op stop timeout=60s interval=0 \
  op monitor interval=20s timeout=40s
primitive rsc_ip_SSA_ASCS00 ocf:aliyun:vpc-move-ip \
  params ip=192.168.4.11 routing_table=vtb-gw8irrnvm8vd29iji5ufk endpoint=vpc-vpc.eu-central-1.aliyuncs.com interface=eth0 \
  op monitor interval=10s timeout=20s
primitive rsc_sap_SSA_ASCS00 SAPIInstance \
  operations $id=rsc_sap_SSA_ASCS00-operations \
  op monitor interval=11 timeout=60 on-fail=restart \
  params InstanceName=SSA_ASCS00_vsapascs \
  START_PROFILE="/sapmnt/SSA/profile/SSA_ASCS00_vsapascs" \
  AUTOMATIC_RECOVER=false \

```

```
meta resource-stickiness=5000 failure-timeout=60 \  
migration-threshold=1 priority=10
```

EXAMPLE 8: ASCS GROUP

```
group grp_SSA_ASCS00 \  
rsc_ip_SSA_ASCS00 rsc_fs_SSA_ASCS00 rsc_sap_SSA_ASCS00 \  
meta resource-stickiness=3000
```

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

As user root, do:

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

6.2.4 Configuring the resources for the ERS

Second, we configure the resources for the file system, IP address and the SAP instance. Of course you need to adapt the parameters to your environment.

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

EXAMPLE 9: ERS PRIMITIVE

```
primitive rsc_fs_SSA_ERS10 Filesystem \  
params device="11e5134833f-dlk6.eu-central-1.nas.aliyuncs.com:/" directory="/usr/sap/  
SSA/ERS10" fstype=nfs \  
op start timeout=60s interval=0 \  
op stop timeout=60s interval=0 \  
op monitor interval=20s timeout=40s  
primitive rsc_ip_SSA_ERS10 ocf:aliyun:vpc-move-ip \  
params ip=192.168.5.11 routing_table=vtb-gw8irrnvm8vd29iji5ufk endpoint=vpc-vpc.eu-  
central-1.aliyuncs.com interface=eth0 \  
op monitor interval=10s timeout=20s  
primitive rsc_sap_SSA_ERS10 SAPIInstance \  
operations $id=rsc_sap_SSA_ERS10-operations \  
op monitor interval=11 timeout=60 on-fail=restart \  
params InstanceName=SSA_ERS10_vsapers \  
START_PROFILE="/sapmnt/SSA/profile/SSA_ERS10_vsapers" \  
AUTOMATIC_RECOVER=false IS_ERS=true \  
meta priority=1000
```

EXAMPLE 10: ERS GROUP

```
group grp_SSA_ERS10 \  
rsc_ip_SSA_ERS10 rsc_fs_SSA_ERS10 rsc_sap_SSA_ERS10 \  
meta resource-stickiness=3000
```

```
rsc_ip_SSA_ERS10 rsc_fs_SSA_ERS10 rsc_sap_SSA_ERS10
```

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

As user root, do:

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

6.2.5 Configuring the colocation constraints between ASCS and ERS

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

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

EXAMPLE 11: LOCATION CONSTRAINT

```
colocation col_sap_SSA_no_both -5000: grp_SSA_ERS10 grp_SSA_ASCS00
location loc_sap_SSA_fail-over_to_ers rsc_sap_SSA_ASCS00 \
    rule 2000: runs_ers_SSA eq 1
order ord_sap_SSA_first_start_ascs Optional: rsc_sap_SSA_ASCS00:start \
    rsc_sap_SSA_ERS10:stop symmetrical=false
```

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

As user root, do:

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

6.2.6 Activating the cluster

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

As user root, do:

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

7 Administration

7.1 Dos and Don'ts

7.1.1 Never stop the ASCS instance

For normal operation **do not stop** the ASCS SAP instance with any tool such as cluster tools or SAP tools. The stop of the ASCS instance might lead to a loss of enqueue locks. Because following the new SAP NW-HA-CLU 7.50 (former version 7.40) certification the cluster must allow local restarts of the ASCS. This feature is needed to allow rolling kernel switch (RKS) updates without reconfiguring the cluster.



Warning

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

7.1.2 Moving ASCS

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

As *ssaadm*, do:

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

7.1.3 Never block resources

With SAP NW-HA-CLU 7.50 (former version 7.40) it is **not longer allowed to block resources** from being controlled manually. This means using the variable *BLOCK_RESOURCES* in */etc/sysconfig/sap_suse_cluster_connector* is not allowed anymore.

7.1.4 Always use unique instance numbers

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

7.1.5 Setting the cluster in maintenance mode

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

As user root, do:

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

As user ssaadm (the full path is needed), do:

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

7.1.6 Ending the cluster maintenance

As user root, do:

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

7.1.7 Cleaning up resources

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

As user root , do:

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



Warning

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

7.2 Testing the cluster

We strongly recommend that you at least process the following tests before you plan going into production with your cluster:

7.2.1 Checking product names with `HAGetFailoverConfig`

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

As user `ssaadm`, do:

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

7.2.2 Starting SAP checks using `HACheckConfig` and `HAGetFailoverConfig`

Check if the HA configuration tests are showing no errors.

As user `ssaadm`, do:

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

7.2.3 Manually moving ASCS

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

As user `root`, do:

```
# crm resource move rsc_sap_SSA_ASCS00 force  
## wait until the ASCS is been moved to the ERS host  
# crm resource clear rsc_sap_SSA_ASCS00
```

7.2.4 Migrating ASCS using `HAfailoverToNode`

Check if moving the ASCS instance using SAP tools like `sapcontrol` does work properly

As user `ssaadm`, do:

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

7.2.5 Testing ASCS migration after failure

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

As user *root*, do:

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

7.2.6 Restarting ASCS inplace using Stop and Start

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



Warning

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

As user *ssaadm*, do:

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

7.2.7 Restarting the ASCS instance automatically (simulating rolling kernel switch)

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

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

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

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

As user *ssaadm*, 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 sapssaas -nr 00 -function StopWait 120 1
  sleep 1
  sapcontrol -host sapssaas -nr 00 -function StartWait 120 1
  sleep 1
done
```

```
$ bash ascs_restart.sh
```

7.2.8 Rolling kernel switch procedure

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

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

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

As user *ssaadm*, type the following commands:

```
## sapcontrol [-user <sidadm psw>] -host <host> -nr <INSTANCE_NR> -function UpdateSystem  
120 300 1  
# sapcontrol -user ssaadm <yourSecurePwd> -host vsapascs -nr 00 -function UpdateSystem  
120 300 1  
# sapcontrol -nr 00 -function GetSystemUpdateList -host vsapascs \  
-user ssaadm <yourSecurePwd>  
# sapcontrol -nr 00 -function GetVersionInfo -host vsapascs \  
-user ssaadm <yourSecurePwd>  
# sapcontrol -nr 10 -function GetVersionInfo -host vsapers \  
-user ssaadm <yourSecurePwd>  
# sapcontrol -nr 01 -function GetVersionInfo -host sapssad1 \  
-user ssaadm <yourSecurePwd>  
# sapcontrol -nr 02 -function GetVersionInfo -host sapssad2 \  
-user ssaadm <yourSecurePwd>
```

7.2.9 Additionally recommended tests

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

8 Additional implementation scenarios

8.1 Adaptive server enterprise replication fail-over automation integration

8.1.1 FM integration with a SUSE Linux Enterprise High Availability Extension cluster

The standard SAP on Alibaba Cloud for an HA setup is as follows: * Multi-AZ deployment with ASCS, * Primary DB running in one AZ, * and their counterpart ERS and Secondary DB running in the second AZ of the same region. The Primary Application Server & Additional Application servers based on the load can be distributed in both AZ's as well to provide resiliency.

Considering a scenario where SAP NetWeaver or Business Suite system is running on SAP Sybase ASE: The completely automated HA for the ABAP Stack (ASCS) is provided by the SUSE Linux Enterprise High Availability Extension cluster. For the Sybase ASE DB, the HA feature is provided with the *Always On* configuration. The fail-over orchestration is done by the Fault Manager (FM) utility which traditionally was installed on a third host (other than the Primary & Secondary DB). In an SAP world, the FM utility comes along with an SAP DB dependent kernel and gets installed in the ASCS Work directory `/usr/sap/<SID>/ASCS<instnr>/exe/`. The fail-over of the ASCS instance along with the associated directories (provided they are installed on a shared file system using NFS) is being taken care of by the SUSE Linux Enterprise High Availability Extension cluster.

8.1.2 Using Sybase ASE *Always On*

SAP Sybase ASE comes with an *Always On* feature which provides native HA & DR capability. The *Always on* option is a high availability and disaster recovery (HADR) system that consists of two SAP ASE servers: One is designated as the primary server on which all transaction processing takes place. The other acts as a warm standby (called "standby server" in DR mode, and as a "companion" in HA mode) for the primary server, and contains copies of designated databases from the primary server. The fail-over orchestration is carried out by ASE provided utility called Fault Manager. The Fault Manager monitors the various components of the HADR environment – Replication Management Agent (RMA), ASE, Replication Server, applications, databases, and the operating system. Its primary job is to ensure the high

availability (zero data loss during fail-over) of the ASE cluster by initiating automatic fail-over with minimal manual intervention. In an SAP Stack, the Fault Manager utility (`syb-dbfm`) comes as part of the DB (Sybase ASE) dependent SAP kernel. Refer to the SAP Standard ASE HA-DR guide (<https://help.sap.com/viewer/efe56ad3cad0467d837c8ff1ac6ba75c/16.0.3.6/en-US/a6645e28bc2b1014b54b8815a64b87ba.html>) for setting up the Sybase ASE DB in HA mode.

! Important

In the following section we use sometimes examples and sometimes general examples. In the general are terms like `<SID>; <instance nr>`. They must be adapted to your environment. As an example, `su - <sid>adm` means:

```
su - ssaadm
```

or in capital letters `cd /usr/sap/<SID>/ASCS<instance nr>/work` means:

```
cd /usr/sap/SSA/ASCS00/work
```

8.1.3 Preparing the database host

! Important

This guide does not duplicate the official HADR documentation. The following procedure describes the key points which you need to take care of.

EXAMPLE 12: INSTALLATION OF A 32-BIT ENVIRONMENT

```
# zypper install glibc-32bit libgcc_s1-32bit
```

For the example this software stack is used:

- SL TOOLSET 1.0 — SWPM → 1.0 SP25 for NW higher than 7.0x
- saphostagent → 7.21 patch 41
- SAP Kernel → 7.53 PL421
- SAP Installation Export → (51051806_1)
- Sybase RDBMS → ASE 16.0.03.06 RDBMS (51053561_1)



Note

It is very useful to refer to the table of installation information which helps to be prepared for the next steps: **SAP Adaptive Server Enterprise - Installation Worksheet** <https://help.sap.com/viewer/efe56ad3cad0467d837c8ff1ac6ba75c/16.0.3.6/en-US/3fe35550f3814b2bb411d5494976e25a.html>



Important

The Fault Manager is enhanced to work in this setup. The minimal versions which support this scenario are * SAP Kernel 749 PL632 * SAP Kernel 753 PL421

8.1.4 Installing the database for replication scenario

The installation can be done with the SWPM which is provided by SAP.

INSTALLING THE PRIMARY DATABASE WITH SWPM:

- SWPM option depends on SAP NetWeaver version and architecture
 - Software Provisioning Manager 1.0 SP 25 → SAP NetWeaver AS for ABAP 7.52 → SAP ASE → Installation → Application Server ABAP → High-Availability System → Database Instance

The following information is requested from the wizard:

- Master Password <secure>
- SAP System Code Page: Unicode (default)
- Uncheck: → Set FQDN for SAP system
- Sybase database Administrator UID: 2003
- In our **demo setup** we have deselect → Use separate devices for sybmgmtdb database (consider different settings for productive environments)

After the basis installation is finished the primary database must be prepared for the replication. First the user **sa** must be unlocked.

```
# su - syb<sid>
# isql -Usapss0 -P <secure password> -S<SID> -X -w1900
# 1> go
# 1> exec sp_locklogin sa, 'unlock'
# 2> go
# Account unlocked.
# (return status = 0)
# 1> quit
```

In the next step, install the SRS software with a response file and enter the following command as user *syb<sid>*: Consult the HADR guide for an example for such a response file. <https://help.sap.com/viewer/efe56ad3cad0467d837c8ff1ac6ba75c/16.0.3.8/en-US/47d295cd825f4e878e493afc0ead77a4.html?q=srs%20response%20file>

```
# /sapcd/ase-16.0.03.06/BD_SYBASE_ASE_16.0.03.06_RDBMS_for_BS_/SYBASE_LINUX_X86_64/
setup.bin -f /sybase/SSA/srs-setup.txt -i silent
```

Activate HADR on primary node with a response file and enter the following command as user *syb<sid>*:

```
# setuphadr /sybase/SSA/SSA_primary_lin.rs.txt
```



Note

If the installation stops with an error message as displayed here, perform the steps explained below:

```
Clean up environment.
Environment cleaned up.
Error: Fail to connect to "PRIM" site SAP ASE at "<hostname>:4901".
```

Check if the host name and port number are correct and the database server is up and running. If everything is correct and network connection should be available, it might help to modify the *interface* file. Try to add a new line in the */sybase/<SID>/interfaces* file for the *<SID>* section with the IP address of the corresponding host name.

```
# vi /sybase/<SID>/interfaces
...
master tcp ether <hostname> 4901
master tcp ether 172.17.1.21 4901
...
```

Create a secure store key entry for the database:

```
# /usr/sap/hostctrl/exe/saphostctrl -user sapadm <secure password> -function
LiveDatabaseUpdate -dbname <SID> -dbtype syb -dbuser DR_admin -dbpass <Secure password>
-updatemethod Execute -updateoption TASK=SET_USER_PASSWORD -updateoption USER=DR_ADMIN
```

INSTALLING THE COMPANION DATABASE WITH SWPM:

- SWPM option depends on SAP NetWeaver version and architecture
 - Software Provisioning Manager 1.0 SP 25 → SAP NetWeaver AS for ABAP 7.52 → SAP ASE → Database Replication → Setup of Replication Environment

The following information is requested from the wizard:

- Replication System Parameters → SID, Master Password, check Set up a secondary database instance
- Primary Database server → host name or virt. name
- Primary Database server port → default is 4901, depends on the setup of your primary server

After the basis installation is finished the companion database must be prepared for the replication. First the user **sa** must be unlocked.

```
# su - syb<sid>
# isql -Usapssso -P <secure password> -S<SID> -X -w1900
# 1> go
# 1> exec sp_locklogin sa, 'unlock'
# 2> go
# Account unlocked.
# (return status = 0)
# 1> quit
```

Next step installing the SRS software with a response file on the companion site and enter the following command as user **syb<sid>**:

```
# /sapcd/ase-16.0.03.06/BD_SYBASE_ASE_16.0.03.06_RDBMS_for_BS_/SYBASE_LINUX_X86_64/
setup.bin -f /sybase/SSA/srs-setup.txt -i silent
```

Activate HADR on companion node with a response file and enter the following command as user **syb<sid>**:

```
# setuphadr /sybase/SSA/SSA_companion_lin.rs.txt
```



Note

In certain circumstances the installation is not successful. It could help to set up the primary system again and install the companion afterward.



Note

If the system is reinstalled and the companion system reports **Missing read/write permissions** for this directory */tmp/.SQLAnywhere*, check the permission on both node. In case the ownership must be changed run the setup again on both nodes. Start with the **Primary**.

Creating a secure store key entry for the database:

```
# /usr/sap/hostctrl/exe/saphostctrl -user sapadm <secure password> -function  
LiveDatabaseUpdate -dbname <SID> -dbtype syb -dbuser DR_admin -dbpass <Secure password>  
-updatemethod Execute -updateoption TASK=SET_USER_PASSWORD -updateoption USER=DR_ADMIN
```

8.1.5 Installing Fault Manager



Note

In this scenario, the FM will be integrated into a cluster who takes already care of the ASCS and ERS of an SAP system. The goal is to make the FM highly available itself and to reuse existing resources.

The Fault Manager is configured on the ASCS host. The benefit from this setup is that the syb-bfm service can be monitored and tracked with the existing pacemaker for the ASCS / ERS replication.

Option one:

- Installation of FM service as part of ASCS (Business Suite)

Option two:

- Stand-alone installation of FM (non Business Suite)

- To make this type of installation ready for a pacemaker implementation, additional requirements needs to be fulfilled
- A file system ~ 2GB which can be moved between all cluster nodes
- Virtual host name for FM instance
- An unused instance number of the SAP system which is already implemented in the cluster (ASCS/ERS pair) [Section 5.2, "Installing ASCS on sapapp1"](#)
- A virtual IP address which can be moved between all cluster nodes



Note

Depending on a later integration of the Fault Manager into the pacemaker cluster, additional storage and IP resources are required. Check [Section 8.2, "Integrating the Fault Manager into the cluster"](#) before you start the installation.

EXAMPLE 13: FAULT MANAGER INSTALLATION AS PART OF THE ASCS INSTANCE

```
# su - <sid>adm
# cd /usr/sap/<SID>/ASCS<instance number>/exe/
# sybdbfm install
```

This is an example of the installation process:

```
replication manager agent user DR_admin and password set in Secure Store.
Keep existing values (yes/no)? (yes)
SAPHostAgent connect user sapadm and password set in Secure Store.
Keep existing values (yes/no)? (yes)
Enter value for primary database host: (sapdb1)
sapdb1
Enter value for primary database name: (SSA)
Enter value for primary database port: (4901)
Enter value for primary site name: (FRA1)
Enter value for primary database heart beat port: (13777)
Enter value for standby database host: (sapdb2)
sapdb1
Enter value for standby database name: (SSA)
Enter value for standby database port: (4901)
Enter value for standby site name : (FRA2)
Enter value for standby database heart beat port: (13787)
Enter value for fault manager host: (vsapfm)
Enter value for heart beat to heart beat port: (13797)
```



```
Enter value for support for floating database ip: (no)
Enter value for use SAP ASE Cockpit if it is installed and running: (no)
```

Update the values as per your environment for the Primary DB & companion DB host name, SID & Site Name. Make sure to use the virtual host name for the ASCS host. When the Fault Manager is installed, profile for it will be created in the `/sapmnt/<SID>/profile` by the name `SYBHA.PFL` and will have the configuration details. Restart the ASCS Instance which will also start the Fault Manager that has been added to the start profile as below:

EXAMPLE 14: ASCS PROFILE AFTER FM INSTALLATION AS INTEGRATED SERVICE

```
# cat /sapmnt/<SID>/profile/<SID>_ASCS<instance number>_<virt. ASCS hostname>
....
#-----
# copy sybdbfm and dependent
#-----
_CP_SYBDBFM_ARG1 = list:$(DIR_CT_RUN)/instancedb.lst
Execute_06 = immediate $(DIR_CT_RUN)/sapcpe$(FT_EXE) pf=$(PF) $_CP_SYBDBFM_ARG1)
_CP_SYBDBFM_ARG2 = list:$(DIR_GLOBAL)/syb/linuxx86_64/cpe_sybodbc.lst
_CP_SYBDBFM_ARG3 = source:$(DIR_GLOBAL)/syb/linuxx86_64/sybodbc
Execute_07 = immediate $(DIR_CT_RUN)/sapcpe$(FT_EXE) pf=$(PF) $_CP_SYBDBFM_ARG2)
$_CP_SYBDBFM_ARG3)
#-----
# Start sybha
#-----
_SYBHAD = sybdbfm.sap$(SAPSYSTEMNAME)_$(INSTANCE_NAME)
_SYBHA_PF = $(DIR_PROFILE)/SYBHA.PFL
Execute_08 = local rm -f $_SYBHAD)
Execute_09 = local ln -s -f $(DIR_EXECUTABLE)/sybdbfm$(FT_EXE) $_SYBHAD)
Restart_Program_02 = local $_SYBHAD) hadm pf=$(SYBHA_PF)
#-----
....
```



Note

In case of a re-installation it might be better to overwrite the existing user name and password in the secure store for the `sapadm` and `DR_admin` if the old values are not 100% known.

EXAMPLE 15: FAULT MANAGER INSTALLATION AS STAND-ALONE SERVICE

The following preparation steps are needed to make the Fault Manager service as highly available and flexible as possible.

- Creating new mount point on **all** nodes where FM should run later
- Mounting shared file system (iSCSI, FC-LUN, NFS)
- Manually adding vIP address for FM instance
- Adapting the `/etc/hosts` with vIP and host name of FM on **all** nodes where FM should run later

For the example below we used this values: SID: SSA (the same SID of the SAP system where the DB is connected too simplifies the integration) instance number: 42 (new) instance name: FM (new) virtual IP: 192.168.6.11 (new), (overlay IP address) virtual host name: vsapfm (new) storage: 1134554a661-wss81.eu-central-1.nas.aliyuncs.com:/ (new), (NFS4 cloud storage)

The official SAP documentation can be found here: <https://help.sap.com/viewer/efe56ad3cad0467d837c8ff1ac6ba75c/16.0.3.6/en-US/e0b6940a381343a8a7c36e90e4e74ae7.html>

As `ssaadm` install the Fault Manager

```
# su - <sid>adm
# cd /usr/sap/SSA/FM42
# FaultManager/setup.bin -f <fault_manager_responses.txt>
```

The Fault Manager installer response file is automatically generated when you complete the HADR configuration on the companion node. The response file is located in `$SYBASE/log/fault_manager_responses.txt`.

A few parameters that need to be updated in the `SYBHA.PFL` to make the fail-over working.

- *[Option-1]* : ASCS integration
- *[Option-2]* : independent integration

EXAMPLE 16: FOR OPTION 1 THE `SYBHA.PFL` FILE IN CASE OF ASCS INTEGRATION.

```
ha/syb/support_cluster = 1
ha/syb/fail-over_if_unresponsive = 1
ha/syb/allow_restart_companion = 1
ha/syb/set_standby_available_after_fail-over = 1
ha/syb/chk_restart_repserver = 1
ha/syb/cluster_fmhost1 = Hostname for Node 1 of the ASCS HA Setup
ha/syb/cluster_fmhost2 = Hostname for Node 2 of the ASCS HA Setup
```

```
ha/syb/use_boot_file_always = 1
ha/syb/dbfmhost = virtual hostname of ASCS instance
```

EXAMPLE 17: FOR OPTION 2 THE *SYBHA.PFL* FILE IN CASE OF INDEPENDENT INTEGRATION.

```
ha/syb/support_cluster = 1
ha/syb/fail-over_if_unresponsive = 1
ha/syb/allow_restart_companion = 1
ha/syb/set_standby_available_after_fail-over = 1
ha/syb/chk_restart_repserver = 1
ha/syb/cluster_fmhost1 = Hostname for Node 1 of the HA Setup
ha/syb/cluster_fmhost2 = Hostname for Node 2 of the HA Setup
ha/syb/use_boot_file_always = 1
ha/syb/dbfmhost = virtual hostname of FM instance
```

Details of all the Fault Manager parameters can be found in the **SAP ASE HA DR User Guide**. Those highlighted in bold are of interest for the setup. Since the Fault Manager is installed with the ASCS which can fail-over from Node 1 to Node 2, the parameters *ha/syb/cluster_fmhost1* and *ha/syb/cluster_fmhost2* provide the physical host names of both nodes where the Fault Manager can potentially run.

EXAMPLE 18: *AS USER SSAADM*, CHECK IF THE *sybdbfm* PROCESS IS SHOWN.

```
ssaadm> sapcontrol -nr 00 -function GetProcessList

23.02.2020 22:11:52
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
msg_server, MessageServer, GREEN, Running, 2020 04 22 18:28:31, 27:43:21, 17731
enserver, EnqueueServer, GREEN, Running, 2020 04 22 18:28:31, 27:43:21, 17732
sybdbfm, , GREEN, Running, 2020 04 22 18:28:31, 27:43:21, 17733
```

The example above shows the Fault Manager integration as part of the ASCS instance.

In a scenario where the complete Availability Zone (AZ1), where the ASCS and Primary database are running, goes down, the DB fail-over is not triggered until the ASCS fail-over is complete and the Fault Manager is up and running in the 2nd Availability Zone (AZ2). The FM then needs to read the boot file to get the prior state of the ASE DB. This is mandatory to ensure that the Fault Manager can trigger the fail-over correctly. The parameter *ha/syb/use_boot_file_always = 1* makes sure that the Fault Manager always reads from the boot file which is part of the work directory (the same for ASCS and FM) and fail-over along with the Fault Manager.

EXAMPLE 19: FM STATUS CHECK AND DB REPLICATION INFORMATION

The status of the FM can be checked as below. Navigate to the ASCS work directory and then run `sybdbfm.sap.<SID>_ASCS<instance number> status` :

As user `ssaadm` for ASCS integration, do:

```
# cd /usr/sap/<SID>/ASCS<instance number>/work
# ./sybdbfm.sap.<SID>_ASCS<instance number> status

fault manager running, pid = 4118, fault manager overall status = OK, currently
executing in mode PAUSING
* sanity check report (65405)*.
node 1: server sapdb1, site FRA1.
db host status: OK.
db status OK hadr status PRIMARY.
node 2: server sapdb2, site FRA2.
db host status: OK.
db status OK hadr status STANDBY.
replication status: SYNC_OK.
failover prerequisites fulfilled: YES.
```

As user `ssaadm` for stand-alone integration, do:

```
# cd /usr/sap/SSA/FM42/work
# ./sybdbfm status

fault manager running, pid = 4118, fault manager overall status = OK, currently
executing in mode PAUSING
* sanity check report (65405)*.
node 1: server sapdb1, site FRA1.
db host status: OK.
db status OK hadr status PRIMARY.
node 2: server sapdb2, site FRA2.
db host status: OK.
db status OK hadr status STANDBY.
replication status: SYNC_OK.
failover prerequisites fulfilled: YES.
```

Checking the log file is also a suitable method to validate the status.

As user `ssaadm`, do:

```
# cd /usr/sap/<SID>/ASCS<instance number>/work
# tail -f dev_sybdbfm
# ...

2020 02/28 15:34:30.523 (23234) ----- Log messages -----
```

```

2020 02/28 15:34:30.523 (23234) Info: saphostcontrol: Executing LiveDatabaseUpdate

2020 02/28 15:34:30.523 (23234) Info: saphostcontrol: LiveDatabaseUpdate
successfully executed

2020 02/28 15:34:30.524 (23234) call is running.
2020 02/28 15:34:30.534 (23234) call exited (exit code 0).
2020 02/28 15:34:30.534 (23234) db status is:
DB_OK.
2020 02/28 15:34:42.561 (23234) * sanity check report (136).
2020 02/28 15:34:42.562 (23234) node 1: server <DB server1>, site <site name one>.
2020 02/28 15:34:42.562 (23234) db host status: OK.
2020 02/28 15:34:42.562 (23234) db status OK hadr status PRIMARY.
2020 02/28 15:34:42.562 (23234) node 2: server <DB server2>, site <site name two>.
2020 02/28 15:34:42.562 (23234) db host status: OK.
2020 02/28 15:34:42.562 (23234) db status OK hadr status STANDBY.
2020 02/28 15:34:42.562 (23234) replication status: SYNC_OK.
2020 02/28 15:34:57.688 (23234) sanity check report (137).
2020 02/28 15:34:57.688 (23234) node 1: server <DB server1>, site <site name one>.
2020 02/28 15:34:57.688 (23234) db host status: OK.
2020 02/28 15:34:57.688 (23234) db status OK hadr status PRIMARY.
2020 02/28 15:34:57.688 (23234) node 2: server <DB server2>, site <site name two>.
2020 02/28 15:34:57.688 (23234) db host status: OK.
2020 02/28 15:34:57.688 (23234) db status OK hadr status STANDBY.
2020 02/28 15:34:57.688 (23234) replication status: SYNC_OK.
2020 02/28 15:35:12.827 (23234) sanity check report (138)*.
2020 02/28 15:35:12.827 (23234) node 1: server <DB server1>, site <site name one>.
2020 02/28 15:35:12.827 (23234) db host status: OK.
2020 02/28 15:35:12.827 (23234) db status OK hadr status PRIMARY.
2020 02/28 15:35:12.827 (23234) node 2: server <DB server2>, site <site name two>.
2020 02/28 15:35:12.827 (23234) db host status: OK.
2020 02/28 15:35:12.827 (23234) db status OK hadr status STANDBY.
2020 02/28 15:35:12.827 (23234) replication status: SYNC_OK.
# ...


```

8.2 Integrating the Fault Manager into the cluster

We have **two options** to implement the FM in the pacemaker environment.


Fault Manager is part of the ASCS instance

- This setup is typically use for SAP Business Suite. (HADR users guide: Installing HADR for Business Suite → Using the Fault Manager with Business Suite)
- The Fault Manager instance is monitored and maintained by pacemaker as sub-instance of the ASCS primitive. That means the Fault Manager is started and stopped and moved along with the ASCS instance.



Fault Manager is running as single instance (own SAP instance and cluster resource)

- Additional configuration steps and resources are required (storage and IP).
- This setup is typically use for SAP none Business Suite. (HADR users guide: Installing HADR for Custom Application → Installing The Fault Manager)
- The FM is totally independent from any other cluster resource. This could be a benefit during maintenance procedures.



Fault Manager is integrated as included service along with the ASCS

EXAMPLE 20: OPTION ONE:

The cluster configuration for the *primitive rsc_sap_<SID>_ASCS<instance number>* needs to be modified. In the example, we use the following values:

- <SID> ⇒ SSA
- <instance number> ⇒ 00
- virtual host name ⇒ vsapascs

```
# crm configure edit rsc_sap_SSA_ASCS00
```

```
primitive rsc_sap_SSA_ASCS00 SAPIInstance \
    operations $id=rsc_sap_SSA_ASCS00-operations \
    op monitor interval=11 timeout=60 on-fail=restart \
    params InstanceName=SSA_ASCS00_vsapascs \
    START_PROFILE="/sapmnt/SSA/profile/SSA_ASCS00_vsapascs" \
    AUTOMATIC_RECOVER=false MONITOR_SERVICES="sybdbfm|msg_server|enserver" \
    meta resource-stickiness=5000 failure-timeout=60 migration-threshold=1
priority=10
```

The Fault Manager service is not part of the default observed SAP instance services. If we specify the **MONITOR_SERVICES** all default settings are overwritten by the named services. That means we have to count all services which are shown as a result of the `sapcontrol -nr 00 -function GetProcessList` command. The example above is for an ENSA1 configuration.



Note

The cluster configuration is different for ENSA1 and ENS2 installation. The names for the **MONITOR_SERVICES** differ between this two versions.

Fault Manager is running as single instance

The next steps may differ depending how the Fault Manager was installed before. In case the Fault Manager was installed as integrated service with the ASCS, you must separate them first.

FOR THE EXAMPLE BELOW WE USED THESE VALUES:

- SID: SSA
 - (the same SID of the SAP system where the DB is connected, too, simplifies the integration)
- instance number: 42 (new)
- instance name: FM (new)
- virtual IP(overlay IP address): 192.168.6.11 (new)
- virtual host name: vsapfm (new)
- storage(NFS4 cloud storage): 1134554a661-wss81.eu-central-1.nas.aliyuncs.com:/ (new)

Fault Manager separation procedure

- Create mount points on all cluster nodes
- Maintain DNS (/etc/hosts)
- Deactivate FM in the ASCS profile
- Create a new profile for FM

- Update the `/usr/sap/sapservices`
- Copy the basic files for the initial start
- Check if the Fault Manager is able to start

EXAMPLE 21: **OPTION TWO:**

Host preparation on all nodes which are cluster members.

```
# mkdir -p /usr/sap/SSA/FM42
## adding the vIP and the host name of FM instance
# vi /etc/hosts
```

You must execute the separation steps only on one cluster node. If the Fault Manager is already running, it needs to be stopped first. The Fault Manager configuration must be uncommented in the ASCS profile. Edit the file `/usr/sap/SSA/SYS/profile/SSA_ASCS00_vs-apasc` and uncomment the Fault Manager sections.

```
#-----
# copy sybdbfm and dependent
#-----
# CP_SYBDBFM_ARG1 = list:$(DIR_CT_RUN)/instancedb.lst
# Execute_00 = immediate $(DIR_CT_RUN)/sapcpe$(FT_EXE) pf=$(PF) $(CP_SYBDBFM_ARG1)
# CP_SYBDBFM_ARG2 = list:$(DIR_GLOBAL)/syb/linuxx86_64/cpe_sybodbc.lst
# CP_SYBDBFM_ARG3 = source:$(DIR_GLOBAL)/syb/linuxx86_64/sybodbc
# Execute_01 = immediate $(DIR_CT_RUN)/sapcpe$(FT_EXE) pf=$(PF) $(CP_SYBDBFM_ARG2)
$(CP_SYBDBFM_ARG3)
# CPARG1 = list:$(DIR_EXECUTABLE)/sapcrypto.lst
# Execute_02 = immediate $(DIR_CT_RUN)/sapcpe$(FT_EXE) pf=$(PF) $(CPARG1)
#-----
# Start sybha
#-----
# _SYBHAD = sybdbfm.sap$(SAPSYSTEMNAME)$(INSTANCE_NAME)
# _SYBHA_PF = $(DIR_PROFILE)/SYBHA.PFL
# Execute_03 = local rm -f $(SYBHAD)
# Execute_04 = local ln -s -f $(DIR_EXECUTABLE)/sybdbfm$(FT_EXE) $(SYBHAD)
# Restart_Program_02 = local $(SYBHAD) hadm pf=$(SYBHA_PF)
```

Now you need a **new** instance profile for the Fault Manager. You can take a **copy** of the ASCS profile and adapt it carefully.

The result should look like this:

```
# cat /usr/sap/SSA/SYS/profile/SSA_FM42_vsapfm
SAPSYSTEMNAME = SSA
SAPSYSTEM = 42
INSTANCE_NAME = FM42
DIR_CT_RUN = $(DIR_EXE_ROOT)$(DIR_SEP)$(OS_UNICODE)$(DIR_SEP)linuxx86_64
```

```

DIR_EXECUTABLE = $(DIR_INSTANCE)/exe
SAPLOCALHOST = vsapfm
DIR_PROFILE = $(DIR_INSTALL)$(DIR_SEP)profile
PF = $(DIR_PROFILE)/SSA_FM42_vsapfm
SETENV_00 = DIR_LIBRARY=$(DIR_LIBRARY)
SETENV_01 = LD_LIBRARY_PATH=$(DIR_LIBRARY):%(LD_LIBRARY_PATH)
SETENV_02 = SHLIB_PATH=$(DIR_LIBRARY):%(SHLIB_PATH)
SETENV_03 = LIBPATH=$(DIR_LIBRARY):%(LIBPATH)
SETENV_04 = PATH=$(DIR_EXECUTABLE):%(PATH)
SETENV_05 = SECUDIR=$(DIR_INSTANCE)/sec
#-----
# copy sybdbfm and dependent
#-----
_CP_SYBDBFM_ARG1 = list:$(DIR_CT_RUN)/instancedb.lst
Execute_00 = immediate $(DIR_CT_RUN)/sapcpe$(FT_EXE) pf=$(PF) $( _CP_SYBDBFM_ARG1)
_CP_SYBDBFM_ARG2 = list:$(DIR_GLOBAL)/syb/linuxx86_64/cpe_sybodbc.lst
_CP_SYBDBFM_ARG3 = source:$(DIR_GLOBAL)/syb/linuxx86_64/sybodbc
Execute_01 = immediate $(DIR_CT_RUN)/sapcpe$(FT_EXE) pf=$(PF) $( _CP_SYBDBFM_ARG2)
$( _CP_SYBDBFM_ARG3)
_CPARG1 = list:$(DIR_CT_RUN)/sapcrypto.lst
Execute_02 = immediate $(DIR_CT_RUN)/sapcpe$(FT_EXE) pf=$(PF) $( _CPARG1)
#-----
# Start sybha
#-----
_SYBHAD = sybdbfm.sap$(SAPSYSTEMNAME)$(INSTANCE_NAME)
_SYBHA_PF = $(DIR_PROFILE)/SYBHA.PFL
Execute_03 = local rm -f $( _SYBHAD)
Execute_04 = local ln -s -f $(DIR_EXECUTABLE)/sybdbfm$(FT_EXE) $( _SYBHAD)
Restart_Program_02 = local $( _SYBHAD) hadm pf=$( _SYBHA_PF)
#suse cluster connector integration
service/halib = $(DIR_EXECUTABLE)/saphascriptco.so
service/halib_cluster_connector = /usr/bin/sap_suse_cluster_connector

```

The SAP *sapstartsrv* needs an entry in the */usr/sap/sapservices* for the Fault Manager. This must be done on **all** cluster nodes. The ASCS entry can be used as template for the Fault Manager. The ERS entry is different and cannot be used as a template.

```

# cat /usr/sap/sapservices
...
LD_LIBRARY_PATH=/usr/sap/SSA/FM42/exe:$LD_LIBRARY_PATH; export LD_LIBRARY_PATH; /
usr/sap/SSA/FM42/exe/sapstartsrv pf=/usr/sap/SSA/SYS/profile/SSA_FM42_vsapfm -D -u
ssaadm

```

Before you can test if the Fault Manager is able to start as single instance, you need some files.

```

# ip a a 192.168.6.11 dev eth0

```

```
# mount 1134554a661-wss81.eu-central-1.nas.aliyuncs.com:/ /usr/sap/SSA/FM42
# mkdir -p /usr/sap/SSA/FM42/{exe,work}
# chown -R ssaadm.sapsys /usr/sap/SSA/FM42
# cp -p /usr/sap/SSA/{ASCS00,FM42}/exe/sapstartsrv
# cp -p /usr/sap/SSA/{ASCS00,FM42}/exe/sapstart
# cp -p /usr/sap/SSA/{ASCS00,FM42}/exe/libsapnwrfc.so
# cp -p /usr/sap/SSA/ASCS00/exe/libicu* /usr/sap/SSA/FM42/exe/
```

The new configuration can be tested as shown. Use CTRL+c to stop it.

```
# LD_LIBRARY_PATH=/usr/sap/SSA/FM42/exe:$LD_LIBRARY_PATH; export LD_LIBRARY_PATH;
# /usr/sap/SSA/FM42/exe/sapstartsrv pf=/usr/sap/SSA/SYS/profile/SSA_FM42_vsapfm -u
  ssaadm

..
SAP Service SAPSSA_42 successfully started.
```

If the result **successfully started** is shown, use Ctrl+c and interrupt the process. Now do the live test with the sapstart framework. In any other cases check your log files, for example */usr/sap/SSA/FM42/work*.

```
# sapcontrol -nr 42 -function StartService SSA
# sapcontrol -nr 42 -function Start
```

As user *ssaadm*, check if the *sybdbfm* process is shown:

```
# sapcontrol -nr 42 -function GetProcessList

23.02.2020 22:11:52
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
sybdbfm, , GREEN, Running, 2020 04 22 18:28:31, 27:43:21, 17733
```

EXAMPLE 22: CLUSTER INTEGRATION AS AN INDEPENDENT INSTANCE

Prepare a file which contains the resource for the Fault Manager. We are using the same method of three primitives (IP, file system, SAP Instance) as used for the ASCS or ERS. The values must be adapted to your infrastructure.

```
# vi crm-fm.txt
primitive rsc_fs_SSA_FM42 Filesystem \
    params device="1134554a661-wss81.eu-central-1.nas.aliyuncs.com:/" \
    directory="/usr/sap/SSA/FM42" \
    fstype=nfs
options="vers=4,minorversion=0,rsize=1048576,wsize=1048576,hard,timeo=600,retrans=2,_netdev,nores
\
    op start timeout=60s interval=0 \
    op stop timeout=60s interval=0 \
```

```

op monitor interval=20s timeout=300s \
meta target-role=Started
primitive rsc_ip_SSA_FM42 ocf:aliyun:vpc-move-ip \
params address=192.168.6.11 routing_table=vtb-gw8irrnvm8vd29iji5ufk
interface=eth0 \
op monitor interval=50s timeout=60s \
meta target-role=Started
primitive rsc_sap_SSA_FM42 SAPIInstance \
operations $id=rsc_sap_SSA_FM42-operations \
op monitor interval=11 timeout=60 on-fail=restart \
params InstanceName=SSA_FM42_myVipNFM \
START_PROFILE="/sapmnt/SSA/profile/SSA_FM42_myVipNFM" \
AUTOMATIC_RECOVER=false MONITOR_SERVICES="sybdbfm" \
meta priority=100 failure-timeout=60 migration-threshold=3 target-role=Started
group grp_SSA_FM42 rsc_ip_SSA_FM42 rsc_fs_SSA_FM42 rsc_sap_SSA_FM42

```

Upload the configuration to the cluster and check the cluster

```

# crm configure load update crm-fm.txt
# crm status

```

8.3 Operating a Pacemaker-controlled and FM-monitored ASE replication setup

An ASE DB replication setup controlled by the Fault Manager needs some special rules which must be followed. First of all, it is important to understand how the status of the replication and the Fault Manager itself can be checked. The following chapter will also give some guidance on how to improve the takeover time and how to control such an environment.

8.3.1 Checking the status

EXAMPLE 23: CHECKING THE STATUS OF THE DATABASE SITUATION WHEN FM IS RUNNING TOGETHER WITH ASCS

Check the status and locate the actual primary DB host.

As user *ssaadm* on the ASCS host, do:

```

# cd /usr/sap/<SID>/ASCS<instance nr>/work
# ./sybdbfm.sap<SID>_ASCS<instance nr> status

```

Check the log file *dev_sybdbfm*

```

2020 03/28 19:38:52.200 (3290) *** sanity check report (2)***.
2020 03/28 19:38:52.200 (3290) node 1: server sapdb1, site FRA1.

```

```

2020 03/28 19:38:52.200 (3290) db host status: OK.
2020 03/28 19:38:52.200 (3290) db status OK hadr status STANDBY.
2020 03/28 19:38:52.200 (3290) node 2: server sapdb2, site FRA2.
2020 03/28 19:38:52.201 (3290) db host status: OK.
220 03/28 19:38:52.201 (3290) db status OK hadr status PRIMARY.
2020 03/28 19:38:52.201 (3290) replication status: SYNC_OK.

```

EXAMPLE 24: CHECKING THE STATUS OF THE DATABASE SITUATION WHEN FM IS RUNNING AS STAND-ALONE INSTANCE

Check the status and locate the actual primary DB host.

As user *ssaadm* on the host where the Fault Manager is running, do:

```

# ssh ssaadm@vsapfm
# cd /usr/sap/SSA/FM42/work
# ./sybdbfm.sapSSA_FM42 status

```

As user *root* on the database host, do:

```

# /usr/sap/hostctrl/exe/saphostctrl -user sapadm <secure password> -dbname SSA -
dbtype syb -function GetDatabaseSystemStatus
# /usr/sap/hostctrl/exe/saphostctrl -user sapadm <secure password> -dbname SSA -
dbtype syb -function GetDatabaseStatus
# /usr/sap/hostctrl/exe/saphostctrl -user sapadm <secure password> -dbname SSA
-dbtype syb -function LiveDatabaseUpdate -updatemethod Check -updateoption
TASK=REPLICATION_STATUS

```

As user *syb* <sid> on the database host, do:

```

# isql -UDR_admin -P <secure password> -S<db host>:4909 -X -w 1000
1> sap_status active_path
2> go

```

8.3.2 Modifying the operating system for DB failover

The application server (PAS and AAS) environment must be adapted for the DB fail-over situation (takeover). On each host which is providing a dialog server (PAS; AAS) the *.dbenv.sh* and/or *.dbenv.csh* file needs to be extended.

EXAMPLE 25: MODIFY THE DB ENVIRONMENT SETTINGS ON THE DIALOG SERVER

Add the missing value and extend the settings as shown below on each host who runs a dialog application server. The names **server1** and **server2** specify the host name of the DB host's where the DB can be run in active mode.

As user *ssaadm*, do:

```

# vi .dbenv.csh

```

```
...
setenv dbs_syb_server <server1:server2>
setenv dbs_syb_ha 1
...
```

As user *ssaadm*, do:

```
# vi .dbenv.sh
...
dbs_syb_server=<server1:server2>
export dbs_syb_server
dbs_syb_ha=1
export dbs_syb_ha
...
```



Important

The instance must be restarted to activate the changes.

EXAMPLE 26: OS SETTINGS FOR FASTER REACTION TIME AFTER PRIMARY DB HOST IS DOWN

The default *tcp_retries* value is too high and causes a very long takeover time. With ASE16 PL7 the behavior is modified. Up to this patch the change below improves the takeover time.

As user *root*, do:

```
# echo 3 >/proc/sys/net/ipv4/tcp_retries2
## makes the changes online
# vi /etc/sysctl.conf
...
net.ipv4.tcp_retries2 = 3
...
## makes the changes reboot persistent
```

8.3.3 Start and Stop procedures

EXAMPLE 27: STARTING AND STOPPING THE SAP SYSTEM AND DATABASES IN REPLICATION MODE

If the Fault Manager is monitoring the Primary and Companion database and the Fault Manager is monitored by Pacemaker, there is a special procedure needed to start and stop the system.

IN GENERAL THESE STEPS ARE IMPORTANT TO START THE SYSTEM:

- Start companion database + replication server

- Start primary database + replication server
- Change cluster maintenance mode to false
 - Start ASCS with FM (automatic)
 - Start ERS (automatic)
- Start PAS and AAS instances
- Optional: release cluster maintenance mode, if the SAP system was started manually
 - File system must be mounted and IP must be set manually
 - As user `<sid> adm` with `sapcontrol -nr <instance number> -function StartSystem`

As user `root` on companion database host, do:

```
# /usr/sap/hostctrl/exe/saphostctrl -function StartDatabase -dbname <SID> -dbtype syb
# /usr/sap/hostctrl/exe/saphostctrl -function StartDatabase -dbname <SID>_REP -dbtype syb
```

As user `root` on primary database host, do:

```
# /usr/sap/hostctrl/exe/saphostctrl -function StartDatabase -dbname <SID> -dbtype syb
# /usr/sap/hostctrl/exe/saphostctrl -function StartDatabase -dbname <SID>_REP -dbtype syb
```

As user `root` on one of the Pacemaker host for ASCS and ERS, do:

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

As user `<sid> adm` on the host for PAS or AAS, do:

```
# sapcontrol -nr <instance number> -function StartSystem
```



Note

If the system should start one by one, use the command `sapcontrol -nr <instance number> -function StartSystem`. The sequence must be: ASCS; ERS; PAS; AAS.

IN GENERAL THESE STEPS ARE IMPORTANT TO **STOP** THE SYSTEM:

- Set cluster maintenance mode to `true`

- Stop PAS and AAS instances
- Stop ASCS with FM
- Stop ERS
- Stop primary database + replication server
- Stop companion database + replication server

As *user root*, do:

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

As *user <sid> adm* on one of the Pacemaker host for ASCS and ERS or PAS / AAS, do:

```
# sapcontrol -nr <instance number> -function StopSystem
```



Note

If the system should stop one by one, use the command `sapcontrol -nr <instance number> -function Stop` on each instance host. The process must be: AAS; PAS; ASCS; ERS.

As *user root* on primary database host, do:

```
# /usr/sap/hostctrl/exe/saphostctrl -function StopDatabase -dbname <SID> -dbtype syb
# /usr/sap/hostctrl/exe/saphostctrl -function StopDatabase -dbname <SID>_REP -dbtype syb
```

As *user root* on companion database host, do:

```
# /usr/sap/hostctrl/exe/saphostctrl -function StopDatabase -dbname <SID> -dbtype syb
# /usr/sap/hostctrl/exe/saphostctrl -function StopDatabase -dbname <SID>_REP -dbtype syb
```



Important

The Pacemaker-controlled server must be stopped in a proper way, too. Depending on the stonith method which is implemented, different procedures are available.

As *user root* on one cluster node, do:

```
# crm cluster run "crm cluster stop"
```


As user *root* on each node, do:

```
# reboot
## or
# poweroff
```

8.3.4 Testing the replication and Fault Manager cluster integration

Important for each high availability solution is an extensive testing procedure. That ensures that the solution is working as expected in case of a failure.

EXAMPLE 28: TRIGGERING A DATABASE FAIL-OVER AND MONITORING IF FM IS WORKING

Check the status and locate the primary site. As user *ssaadm* on the ASCS host, do:

```
# cd /usr/sap/<SID>/ASCS<instance nr>/work
# ./sybdbfm.sap<SID>_ASCS<instance nr> status
```

Check the log file *dev_sybdbfm*

```
2020 03/28 19:38:52.200 (3290) *** sanity check report (2)***.
2020 03/28 19:38:52.200 (3290) node 1: server sapdb1, site FRA1.
2020 03/28 19:38:52.200 (3290) db host status: OK.
2020 03/28 19:38:52.200 (3290) db status OK hadr status STANDBY.
2020 03/28 19:38:52.200 (3290) node 2: server sapdb2, site FRA2.
2020 03/28 19:38:52.201 (3290) db host status: OK.
2020 03/28 19:38:52.201 (3290) db status OK hadr status PRIMARY.
2020 03/28 19:38:52.201 (3290) replication status: SYNC_OK.
```

- Now destroy the primary database server.
- Monitor the takeover process with the FM.

As user *ssaadm* on the ASCS host (FM running as integrated ASCS service), do:

```
# cd /usr/sap/<SID>/ASCS<instance nr>/work
# tail -f dev_sybdbfm
```

EXAMPLE 29: SELECTED OUTPUT FROM THE TAKEOVER PROCESS.

```
...
2020 03/2711:08:38.301 (3290) * sanity check report (270)* .
2020 03/2711:08:38.301 (3290) node 1: server sapdb1, site FRA1.
2020 03/2711:08:38.301 (3290) db host status: OK.
2020 03/2711:08:38.301 (3290) db status OK hadr status STANDBY.
2020 03/2711:08:38.301 (3290) node 2: server sapdb2, site FRA2.
2020 03/2711:08:38.301 (3290) db host status: OK.
```

```

2020 03/2711:08:38.301 (3290) db status OK hadr status PRIMARY.
2020 03/2711:08:38.301 (3290) replication status: SYNC_OK.
2020 03/2711:08:50.416 (3290) ERROR in function SimpleFetch (1832)
(SQLExecDirect failed): (30046) [08S01] [SAP][ASE ODBC Driver]Connection to the
server has been lost. Unresponsive Connection was disconnected during command
timeout. Check the server to determine the status of any open transactions.
2020 03/2711:08:50.416 (3290) ERROR in function SimpleFetch (1832)
(SQLExecDirect failed): (30149) [HYT00] [SAP][ASE ODBC Driver]The command has timed
out.
2020 03/2711:08:50.416 (3290) execution of statement master..sp_hadr_admin
get_request, '1' failed.
2020 03/2711:08:50.416 (3290) ERROR in function SimpleFetch (1824) (SQLAllocStmt
failed): (30102) [HY010] [SAP][ASE ODBC Driver]Function sequence error
2020 03/2711:08:50.416 (3290) execution of statement select top 1
convert( varchar(10), @@hadr_mode ) || ' ' || convert( varchar(10), @@hadr_state )
from sysobjects failed.
2020 03/2711:08:50.416 (3290) disconnect connection
2020 03/2711:09:22.505 (3290) ERROR in function SQLConnectWithRetry (1341)
(SQLConnectWithRetry failed): (30293) [HY000] [SAP][ASE ODBC Driver]The socket
failed to connect within the timeout specified.
2020 03/2711:09:22.505 (3290) ERROR in function SQLConnectWithRetry (1341)
(SQLConnectWithRetry failed): (30012) [08001] [SAP][ASE ODBC Driver]Client unable
to establish a connection
2020 03/2711:09:22.505 (3290) connected with warnings (555E69805100)
2020 03/2711:09:22.505 (3290) ERROR in function SimpleFetch (1824) (SQLAllocStmt
failed): (30293) [HY000] [SAP][ASE ODBC Driver]The socket failed to connect within
the timeout specified.
2020 03/2711:09:22.505 (3290) ERROR in function SimpleFetch (1824) (SQLAllocStmt
failed): (30012) [08001] [SAP][ASE ODBC Driver]Client unable to establish a
connection
2020 03/2711:09:22.505 (3290) execution of statement select top 1
convert( varchar(10), @@hadr_mode ) || ' ' || convert( varchar(10), @@hadr_state )
from sysobjects failed.
2020 03/2711:09:22.505 (3290) disconnect connection
2020 03/2711:09:22.505 (3290) primary site unusable.
...
2020 03/2711:09:22.984 (3290) primary site unusable.
2020 03/2711:09:22.984 (3290) * sanity check report (271)* .
2020 03/2711:09:22.984 (3290) node 1: server sapdb1, site FRA1.
2020 03/2711:09:22.984 (3290) db host status: OK.
2020 03/2711:09:22.984 (3290) db status OK hadr status STANDBY.
2020 03/2711:09:22.984 (3290) node 2: server sapdb2, site FRA2.
2020 03/2711:09:22.984 (3290) db host status: UNUSABLE.
2020 03/2711:09:22.984 (3290) db status DB INDOUBT hadr status UNREACHABLE.
2020 03/2711:09:22.984 (3290) replication status: SYNC_OK.
2020 03/2711:09:23.047 (3290) doAction: Primary database is declared dead or
unusable.

```

```

2020 03/2711:09:23.047 (3290) disconnect connection
2020 03/2711:09:23.047 (3290) database host cannot be reached.
2020 03/2711:09:23.047 (3290) doAction: fail-over.
...
2020 03/2711:11:55.497 (3290) * sanity check report (273)* .
2020 03/2711:11:55.497 (3290) node 1: server sapdb1, site FRA1.
2020 03/2711:11:55.497 (3290) db host status: OK.
2020 03/2711:11:55.497 (3290) db status OK hadr status PRIMARY.
2020 03/2711:11:55.497 (3290) node 2: server sapdb2, site FRA2.
2020 03/2711:11:55.497 (3290) db host status: UNUSABLE.
2020 03/2711:11:55.498 (3290) db status DB INDOUBT hadr status UNREACHABLE.
2020 03/2711:11:55.498 (3290) replication status: UNKNOWN.
2020 03/2711:11:55.555 (3290) doAction: Standby database is declared dead or
unusable.
2020 03/2711:11:55.555 (3290) disconnect connection
2020 03/2711:11:55.555 (3290) doAction: Companion db host is declared unusable.
2020 03/2711:11:55.555 (3290) doAction: no action defined.
2020 03/2711:11:58.568 (3290) Error: NIECONN_REFUSED (No route to host),
NiRawConnect failed in plugin_fopen()
...
host is coming back online ##
2020 03/2711:18:45.579 (3290) call is running.
2020 03/2711:18:45.589 (3290) call exited (exit code 0).
2020 03/2711:18:45.589 (3290) db status is: DB_OK.
2020 03/2711:18:45.589 (3290) doAction: Standby database is declared dead or
unusable.
2020 03/2711:18:45.589 (3290) disconnect connection
2020 03/2711:18:45.589 (3290) doAction: Companion db host is declared ok.
2020 03/2711:18:45.589 (3290) doAction: restart database.
2020 03/2711:18:45.805 (3290) Webmethod returned successfully
...
2020 03/2711:22:43.677 (3290) * sanity check report (286)* .
2020 03/2711:22:43.677 (3290) node 1: server sapdb1, site FRA1.
2020 03/2711:22:43.677 (3290) db host status: OK.
2020 03/2711:22:43.677 (3290) db status OK hadr status PRIMARY.
2020 03/2711:22:43.677 (3290) node 2: server sapdb2, site FRA2.
2020 03/2711:22:43.677 (3290) db host status: OK.
2020 03/2711:22:43.677 (3290) db status OK hadr status STANDBY.
2020 03/2711:22:43.677 (3290) replication status: SYNC_OK.
...

```

As user root, do:

```

# /usr/sap/hostctrl/exe/saphostctrl -user sapadm <secure password> -dbname <SID>
-dbtype syb -function LiveDatabaseUpdate -updatemethod Check -updateoption
TASK=REPLICATION_STATUS

Webmethod returned successfully

```

```
Operation ID: 5254001F87CB1C75B5C34755C991EDFA
```

```
----- Response data -----
```

```
TASK_NAME=REPLICATION_STATUS
```

```
REPLICATION_STATUS=active
```

```
PRIMARY_SITE=<site1>
```

```
STANDBY_SITE=<site2>
```

```
REPLICATION_MODE=sync
```

```
ASE transaction log backlog (MB)=0
```

```
Replication queue backlog (MB)=0
```

```
TASK_STATUS=OK
```

```
----- Log messages -----
```

```
Info: saphostcontrol: Executing LiveDatabaseUpdate
```

```
Info: saphostcontrol: LiveDatabaseUpdate successfully executed
```

EXAMPLE 30: TRIGGERING AN FM FAILURE

Killing the Fault Manager process more than five times will bring pacemaker in action. Up to five times the `saphostagent` will take care of the SAP process. If this fail-count is reached in a specific time window, the service will not be restarted.

As user ssaadm, do:

```
# pkill -9 sybdbfm
## check that the PID has changed
# sapcontrol -nr 42 -function GetProcessList
# pkill -9 sybdbfm
...
# sapcontrol -nr 42 -function GetProcessList
...
sybdbfm, , GRAY, Stopped, , , 11154
...
```

Now pacemaker will restart the Fault Manager instance locally first. *As user root, do:*

```
# crm_mon -lrfn
...
Migration Summary:
* Node <hostname>:
rsc_sap_SSA_FM42: migration-threshold=3 fail-count=1 last-failure='Fri Mar 27
13:46:39 2020
...
```



Note

If the **fail-count** reaches the defined threshold, the Fault Manager instance is moved away from that host. If the Fault Manager is integrated as part of the ASCS, both will be moved away.

9 References

For more information, see the documents listed below.

9.1 Pacemaker

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

9.2 Related Manual Pages

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

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

9.3 Related SUSE TID

- SUSE SAP Best Practice Guide Errata (<https://www.suse.com/support/kb/doc?id=7023713>)
- SAP S/4 HANA - Enqueue Replication 2 High Availability Cluster - Setup Guide (<https://www.suse.com/support/kb/doc?id=7023714>)
- Integration of sap-suse-cluster-connector does not work as expected (<https://www.suse.com/support/kb/doc?id=000019244>)
- Long Client hang to Cluster after failover of ERS Instance (<https://www.suse.com/support/kb/doc?id=000019293>)
- sap_suse_cluster_connector stuck as HAActive: FALSE (<https://www.suse.com/support/kb/doc?id=000019924>)
- Use of Filesystem resource for ASCS/ERS HA setup not possible (<https://www.suse.com/support/kb/doc?id=000019944>)

9.4 Related SAP Notes

- 405827 - Linux: Recommended file systems (<https://launchpad.support.sap.com/#/notes/405827>)
- 953653 - Rolling Kernel Switch (<https://launchpad.support.sap.com/#/notes/953653/E>)

- 1056161 - SUSE Priority Support for SAP applications (<https://launchpad.support.sap.com/#/notes/1056161>)
- 1092448 - IBM XL C/C++ runtime environment for Linux on system p (<https://launchpad.support.sap.com/#/notes/1092448/E>)
- 1153713 - Problems with SAP Management Console (Java) (<https://launchpad.support.sap.com/#/notes/1153713/E>)
- 1275776 - Linux: Preparing SLES for SAP environments (<https://launchpad.support.sap.com/#/notes/1275776>)
- 1763512 - Support details for SUSE Linux Enterprise for SAP Applications (<https://launchpad.support.sap.com/#/notes/1763512/E>)
- 1944799 - SAP HANA Guidelines for SLES Operating System Installation (<https://launchpad.support.sap.com/#/notes/1944799>)
- 1984787 - SUSE LINUX Enterprise Server 12: Installation notes (<https://launchpad.support.sap.com/#/notes/1984787/E>)
- 2077934 - Rolling kernel switch in HA environments (<https://launchpad.support.sap.com/#/notes/2077934/E>)
- 2205917 - SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12 (<https://launchpad.support.sap.com/#/notes/2205917>)
- 2235581 - SAP HANA: Supported Operating Systems (<https://launchpad.support.sap.com/#/notes/2235581/E>)
- 2254173 - Linux: Rolling Kernel Switch in Pacemaker based NetWeaver HA environments (<https://launchpad.support.sap.com/#/notes/2254173/E>)
- 2308598 - Error "Ifconfig is obsolete" happens when you run startsap to start SAP system (<https://launchpad.support.sap.com/#/notes/2308598>)
- 2369910 - SAP Software on Linux: General information (<https://launchpad.support.sap.com/#/notes/2369910/E>)
- 2578899 - SUSE Linux Enterprise Server 15: Installation Note (<https://launchpad.support.sap.com/#/notes/2578899>)
- 2684254 - SAP HANA DB: Recommended OS settings for SLES 15 / SLES for SAP Applications 15 (<https://launchpad.support.sap.com/#/notes/2684254>)

- 2855499 - FAIL: RKS Warning(s): Unsupported SCS instance with additional gateway found (<https://launchpad.support.sap.com/#/notes/2855499>)
- 2992028 - SYB: Fault Manager: install Fault Manager in separated instance (<https://launchpad.support.sap.com/#/notes/2992028>)
- 3075829 - New sapstarstrv Web service method "ABAPSetServerInactive" (<https://launchpad.support.sap.com/#/notes/3075829>)
- 3091152 - sapstartsrv - improved deregistration for UNIX/Linux (<https://launchpad.support.sap.com/#/notes/3091152>)
- 3115889 - SAP Web Dispatcher embedded deployment in an ASCS/SCS instance (<https://launchpad.support.sap.com/#/notes/3115889>)
- 3139184 - Linux: systemd integration for sapstartsrv and SAP Hostagent (<https://launchpad.support.sap.com/#/notes/3139184>)
- 3145200 - SAP Host Agent 7.22 PL57 (<https://launchpad.support.sap.com/#/notes/3145200>)

10 Appendix

10.1 CRM configuration

The complete crm configuration for SAP system SSA looks as follows:

```
## nodes

node sapapp1
node sapapp2

## aliyun_fence

primitive res_ALIYUN_STONITH_1 stonith:fence_aliyun \
  op monitor interval=120 timeout=60 \
  params plug=i-gw87xi82sj2dy2ysaw19 ram_role=SAP-HA-ROLE region=eu-central-1 \
  meta target-role=Started
primitive res_ALIYUN_STONITH_2 stonith:fence_aliyun \
  op monitor interval=120 timeout=60 \
  params plug=i-gw86pnh1jy1dw0vfer3w ram_role=SAP-HA-ROLE region=eu-central-1 \
  meta target-role=Started

## primitives for ASCS and ERS

primitive rsc_fs_SSA_ASCS00 Filesystem \
  params device="114b194b126-tkt51.eu-central-1.nas.aliyuncs.com:/" directory="/usr/sap/
SSA/ASC00" fstype=nfs \
  op start timeout=60s interval=0 \
  op stop timeout=60s interval=0 \
  op monitor interval=20s timeout=40s
primitive rsc_fs_SSA_ERS10 Filesystem \
  params device="11e5134833f-dlk6.eu-central-1.nas.aliyuncs.com:/" directory="/usr/sap/
SSA/ERS10" fstype=nfs \
  op start timeout=60s interval=0 \
  op stop timeout=60s interval=0 \
  op monitor interval=20s timeout=40s
primitive rsc_ip_SSA_ASCS00 ocf:aliyun:vpc-move-ip \
  params ip=192.168.4.11 routing_table=vrb-2zeqrgjv9pv2m85oqvhw endpoint=vpc-vpc.eu-
central-1.aliyuncs.com interface=eth0 \
  op monitor interval=10s timeout=20s
primitive rsc_ip_SSA_ERS10 ocf:aliyun:vpc-move-ip \
  params ip=192.168.5.11 routing_table=vrb-2zeqrgjv9pv2m85oqvhw endpoint=vpc-vpc.eu-
central-1.aliyuncs.com interface=eth0 \
  op monitor interval=10s timeout=20s
primitive rsc_sap_SSA_ASCS00 SAPIInstance \
```

```

operations $id=rsc_sap_SSA_ASCS00-operations \
op monitor interval=11 timeout=60 on-fail=restart \
params InstanceName=SSA_ASCS00_vsapascs \
  START_PROFILE="/sapmnt/SSA/profile/SSA_ASCS00_vsapascs" \
  AUTOMATIC_RECOVER=false \
meta resource-stickiness=5000 failure-timeout=60 migration-threshold=1 \
  priority=10
primitive rsc_sap_SSA_ERS10 SAPIInstance \
operations $id=rsc_sap_SSA_ERS10-operations \
op monitor interval=11 timeout=60 on-fail=restart \
params InstanceName=SSA_ERS10_vsapers \
  START_PROFILE="/sapmnt/SSA/profile/SSA_ERS10_vsapers" \
  AUTOMATIC_RECOVER=false IS_ERS=true \
meta priority=1000
primitive stonith-sbd stonith:external/sbd \
params pcmk_delay_max=30s

## group definitions for ASCS and ERS

group grp_SSA_ASCS00 rsc_ip_SSA_ASCS00 rsc_fs_SSA_ASCS00 rsc_sap_SSA_ASCS00 \
  meta resource-stickiness=3000
group grp_SSA_ERS10 rsc_ip_SSA_ERS10 rsc_fs_SSA_ERS10 rsc_sap_SSA_ERS10

## constraints between ASCS and ERS

colocation col_sap_SSA_not_both -5000: grp_SSA_ERS10 grp_SSA_ASCS00
location loc_sap_SSA_fail-over_to_ers rsc_sap_SSA_ASCS00 \
  rule 2000: runs_ers_SSA eq 1
order ord_sap_SSA_first_ascs Optional: rsc_sap_SSA_ASCS00:start rsc_sap_SSA_ERS10:stop
  symmetrical=false

## constraints between node and stonith resources

location loc_sapapp1_stonith_not_on_sapapp1 res_ALIYUN_STONITH_1 -inf: sapapp1
location loc_sapapp2_stonith_not_on_sapapp2 res_ALIYUN_STONITH_2 -inf: sapapp2

## crm properties and more

property cib-bootstrap-options: \
  have-watchdog=false \
    dc-version="2.0.1+20190417.13d370ca9-3.9.1-2.0.1+20190417.13d370ca9" \
  cluster-infrastructure=corosync \
  cluster-name=hacluster \
  stonith-enabled=true \
  last-lrm-refresh=1494346532
rsc_defaults rsc-options: \
  resource-stickiness=1 \

```

```
migration-threshold=3
op_defaults op-options: \
timeout=600 \
record-pending=true
```

10.2 Corosync configuration of the two-node cluster

Find below the corosync configuration including a secondary heartbeat ring.

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

    transport: udpu
}

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

```
}  
  
nodelist {  
  node {  
    ring0_addr: 192.168.1.123  
    nodeid: 1  
  }  
  
  node {  
    ring0_addr: 192.168.2.92  
    nodeid: 2  
  }  
}  
  
quorum {  
  
  # Enable and configure quorum subsystem (default: off)  
  # see also corosync.conf.5 and votequorum.5  
  provider: corosync_votequorum  
  expected_votes: 2  
  two_node: 1  
}
```

11 Legal notice

Copyright © 2006–2025 SUSE LLC and contributors. All rights reserved.

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

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

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

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

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

12 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.