



Best Practice for SAP

Enqueue Replication - SAP NetWeaver
High Availability on SUSE Linux Enterprise
(12)



Best Practice for SAP

Enqueue Replication - SAP NetWeaver High Availability on SUSE Linux Enterprise (12)

by Fabian Herschel (fabian.herschel@suse.com) and Lars Pinne (lars.pinne@suse.com)

Publication Date: July 07, 2016

SUSE LLC

10 Canal Park Drive

Suite 200

Cambridge MA 02141

USA

<https://www.suse.com/documentation> 

Contents

- 1 Executive Summary 1**
- 2 Overview 2**
 - 2.1 What is this guide about? 2
 - 2.2 What is not covered? 4
- I INTRODUCTION AND CONCEPT 6**
- 3 Introduction 7**
 - 3.1 SAP on Linux 7
 - 3.2 SUSE Linux Enterprise Server for SAP Applications 7
 - 3.3 SUSE Linux Enterprise High Availability Extension 9
 - Storage Scenarios and I/O Stack for SAP 11 • SAP Resource Agents 13 • The SAP SUSE cluster connector 13 • Virtualization 14
 - 3.4 SAP ASE as RDBMS for SAP 15
 - 3.5 Sizing, Deployment, and Migration Considerations 16
- 4 One Concept — Two Use Cases 18**
- II INSTALLATION OF ENQUEUE REPLICATION WITH SAP ASE 22**
- 5 Installation Overview 23**
- 6 Planning 24**
- 7 Prerequisites 25**
 - 7.1 Hardware Requirements 25
 - 7.2 Software Requirements, Connection Data, etc. 26

8 Get the needed SAP Installation Media 28

- 8.1 Download Procedure 28
- 8.2 Extract downloaded Archives 30

9 Installation of SUSE Linux Enterprise Server for SAP Applications 31

- 9.1 Installing SLES for SAP 31
- 9.2 Additional Configuration 32
- 9.3 File System Layout 33
- 9.4 Update Channels and Software Selection 33
- 9.5 Runlevel and Services 35
- 9.6 Miscellaneous 35
- 9.7 Check the Installation 36

10 Prepare Storage for data and fencing 39

11 Install SAP NetWeaver 7.50 40

- 11.1 Check operating system prerequisites 40
- 11.2 Installing the ASCS Instance on Node cl2n01 41
- 11.3 Installing SAP ASE Database 42
- 11.4 Installing Enqueue-Replication Instance on Node cl2n01 44
- 11.5 Installing Central Instance on Node cl2n01 45
- 11.6 Installing Additional Dialog Instances on Node cl2n01 47
- 11.7 Update SAP Kernel and SAPHOSTAGENT Package 48
- 11.8 SAP Profile Changes to Match the Master-Slave Scenario 48
- 11.9 Integrating the sap_suse_cluster_connector 49

12 Additional Work on the Second Node 50

- 12.1 Duplicate SAP Related Files to the Second Node 50
 - Additional files and directories for MaxDB 51

13 Configuration of SUSE Linux Enterprise High Availability Extension and SAP NetWeaver integration 52

- 13.1 Basic Cluster Configuration 52
- 13.2 Adapting the Configuration Manually 53
 - Change the corosync configuration to UCAST and two rings 53 • Setup the use of more than one sbd 55
- 13.3 Join the 2node to the Cluster 57
- 13.4 Start the Cluster for the first Time 57
- 13.5 Configure Bootstrap and STONITH Using crm Shell 59
- 13.6 Configure Resources Using crm Shell 60
- 13.7 Configure Constraints Using crm Shell 62

14 Administration and Troubleshooting 63

15 Cluster Tests 65

- 15.1 Mandatory Test Cases for SAP Enqueue Replication Server Integration 65
- 15.2 Recommended Tests for SLE High Availability Extension Basic Functionality 68

III APPENDIX 71

A Software Downloads 72

B SUSE Products Online Documentation 73

C SAP Notes 74

D Links to SLES for SAP, SAP Software, SAP ASE 76

E Licenses 78

E.1 GNU Free Documentation License 78

E.2 GNU General Public License 81

Terminology 86

1 Executive Summary

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

SAP 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 e.g. as provided by SUSE Linux Enterprise Server for SAP Applications. The described concept has proven its maturity over several years of productive operations for customers of different sizes and branches.

This guide will show you 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.

2 Overview

2.1 What is this guide about?

SUSE Linux Enterprise Server for SAP Applications is the optimal platform to run SAP applications with high availability. Together with a redundant layout of the technical infrastructure, single points of failure can be eliminated. The SAP Enqueue Replication Server constitutes application level redundancy for one of the most crucial components of the SAP NetWeaver stack, the enqueue service. "The replicated enqueue server solves the problem of the enqueue service, which is a single point-of-failure. It constitutes a fully transparent solution to this problem and enables the SAP system to continue functioning in the event that an enqueue server fails..." (SAP Help article [SAP High Availability - Standalone Enqueue Replication Service \(http://help.sap.de/saphelp_nw73/helpdata/de/45/25ebd455c56d7be10000000a114a6b/frameset.htm\)](http://help.sap.de/saphelp_nw73/helpdata/de/45/25ebd455c56d7be10000000a114a6b/frameset.htm)) An optimal effect of the enqueue replication mechanism can be achieved when combining the application level redundancy (replicated enqueues) with a high availability cluster solution e.g. as provided by SUSE Linux Enterprise Server for SAP Applications.

The integration of the HA cluster through the SAP control framework using the `sap_suse_cluster_connector` is of particular interest. "One of the classical problems running SAP instances in a highly available environment is that if a SAP administrator changes the status (start/stop) of a SAP instance without using the interfaces provided by the cluster software then the cluster framework will detect that as an error status and will bring the SAP instance into the old status by either starting or stopping the SAP instance. This can result into very dangerous situations, if the cluster changes the status of a SAP instance during some SAP maintenance tasks. The solution is that the central component `SAPSTARTSRV`, which controls SAP instances since SAP Kernel versions 6.4, will be enabled to communicate the state change (start/stop) to the cluster software." (SAP SDN article "How to Connect `SAPSTARTSRV` and Cluster Frameworks using the Components `saphascriptco.so` and `SAP_Vendor_Cluster_Connector`" <http://scn.sap.com/docs/DOC-28875>).

The high availability cluster for the replicated enqueue server described in this document can be installed on physical hardware as well as in virtual machines. All major virtualization platforms are supported. Information about virtualization platform support for SAP Enqueue Replication Server can be found in SAP Note 1122387 - Linux: SAP Support in virtualized environments.

Information about virtualization platform support for SAP NetWeaver and SUSE Linux Enterprise Server (SLE) could be found on <http://scn.sap.com/docs/DOC-27321> and other web pages listed in the appendix.

This guide will show you 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`.

This guide will also help you to install the following software components:

- SUSE Linux Enterprise Server for SAP Applications SP1 (including the pacemaker cluster)
- SAP NetWeaver 7.50 (other versions are supported, too).

This guide is aimed at IT professionals with skills in:

- SAP basic operating,
- Data center system concepts and configuration,
- Linux knowledge at LPI1 or CLE level.

To follow this guide you need access to the following resources:

- SUSE Linux Enterprise Server for SAP Applications 12 SP1 installation media.
- To update the systems you must have either Internet access, SUSE Manager, or a local Subscription Management Tool (SMT).
- SAP NetWeaver 7.50 Installation Media.
- Appropriate hardware (two servers, network, storage, see below).
- Connectivity to the database and NFS server.

This guide is organized into two main parts:

- Part I - Introduction and Concept

In a perfect world everyone who plans, installs, or runs SAP NetWeaver including SAP Enqueue Replication Server with high availability on SUSE Linux Enterprise Server for SAP Applications has read this chapter.

- Part II - Planning, Installing, and Testing the Cluster

Details to help perform the mentioned tasks or to learn more details about SAP Enqueue Replication Server with high availability on SLES for SAP.

2.2 What is not covered?

While this guide uses well known components, additional features of SUSE Linux Enterprise Server for SAP Applications SP1 allow for enhanced concepts, like the use of local area network-based Distributed Replicated Block Devices (DRBD) instead of storage area networks. These concepts are described in additional documents.

SLES for SAP contains SUSE Linux Enterprise Server and SUSE Linux Enterprise High Availability Extension. For platforms where SLES for SAP is not available, it is possible to set up an high availability solution using SUSE Linux Enterprise Server and SUSE Linux Enterprise High Availability Extension instead of SLES for SAP. Such a setup will lack features of SLES for SAP, like `sap_suse_cluster_connector`, page cache limit, and Expanded Service Pack Overlap Support. Nevertheless, SUSE Linux Enterprise Server and SUSE Linux Enterprise High Availability Extension would allow to run SAP Enqueue Replication Server with high availability on other hardware platforms like System z. Such platforms may need different configuration parameters that are not explained in this document.

To decide if SAP Enqueue Replication Server fits for a given situation falls into the domain of SAP application architecture, which is out of scope for this document. Also general configuration and administration of the SAP Enqueue Replication Server is not covered by this document. Finally, system sizing and performance related questions are not discussed in this document. More Best Practices for SAP on SUSE Linux Enterprise are published at <https://www.suse.com/products/sles-for-sap/resource-library/sap-best-practices.html>.

Selection, planning, and installing of the needed database is not described here. We refer to the database planning and installation where it helps to accomplish our procedures. More information on planning and installing a Linux high availability platform for databases is available at the web pages listed in the appendix.

Setting up a multipath connection to the SAN storage is not described here in detail. We refer to the multipath configuration where necessary for our tasks. More information about Linux multipath configuration can be found in the product manuals and at the web pages mentioned in the appendix. Information on connecting SUSE Linux Enterprise Server to a given SAN storage can be found in the storage manufacturer's documentation.

I Introduction and Concept

3 Introduction 7

4 One Concept — Two Use Cases 18

3 Introduction

3.1 SAP on Linux

SUSE has multiple dedicated resources working at SAP headquarters and the SAP LinuxLab to ensure maximum interoperability between our products with SAP software and technologies.

SAP has built SAP LinuxLab to assist with the release of SAP software on Linux. LinuxLab supports other SAP departments in development of the Linux platform, processes Linux-specific support problems and acts as an information hub to all the SAP partners in the Linux ecosystem.

Where possible, SAP endorsed open standards and technologies. This allowed SAP to support a very wide range of operating systems and hardware platforms. Open-source-based Linux provides the maximum in “openness,” so it was only natural for SAP to start supporting it in 1999. SAP tries to be “operating system agnostic” and act neutral on the customer's chosen operating systems. Unlike other software vendors, SAP has clearly stated its policies toward open source and Linux. For instance, the usage of binary only (closed source) device drivers (kernel modules) is not supported. This helps the Linux and open source communities, since hardware vendors are encouraged to either publish the specifications and APIs of their hardware so the Linux community can write drivers, or make driver source code that can be included in the Linux kernel available (see SAP note 784391).

3.2 SUSE Linux Enterprise Server for SAP Applications

SLES for SAP is based on the SUSE Linux Enterprise technology, a highly reliable, scalable, secure and optimized server operating system that is built to power both physical and virtual mission-critical workloads. SUSE Linux Enterprise Server for SAP Applications is optimized for all mission-critical SAP software solutions. SUSE Linux Enterprise Server is validated and certi-

fied by SAP and is, therefore, well positioned to support SAP software solutions including appliances. The graphic shows the components SUSE Linux Enterprise Server for SAP Applications consists of:

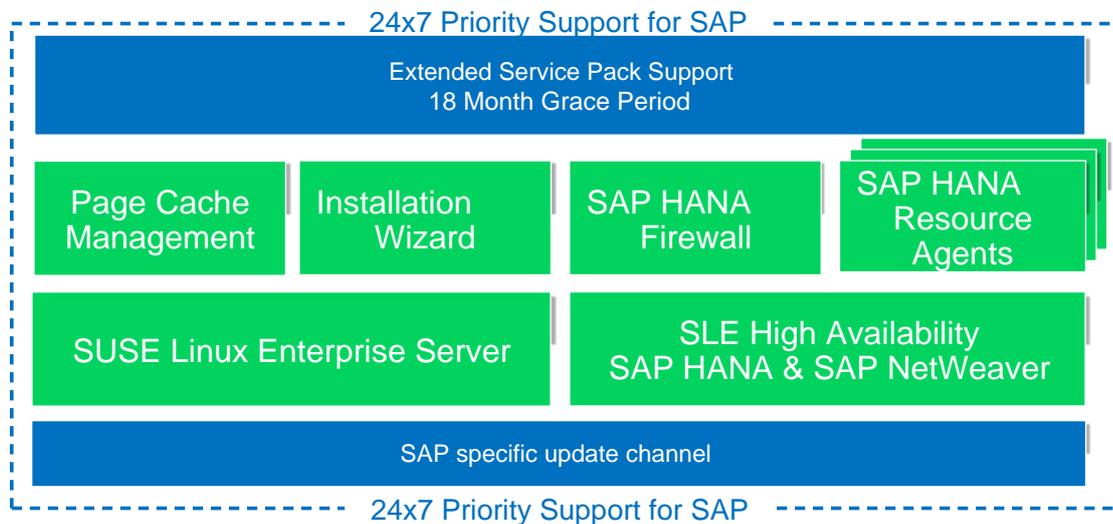


FIGURE 3.1: COMPONENTS OF SUSE LINUX ENTERPRISE SERVER FOR SAP APPLICATIONS

System parameters and components are prepared for the SAP workload:

- specific Linux software selection pattern,
- kernel settings,
- page cache limit,
- file permissions and security settings,
- runlevel services configuration,
- UUID generation,
- SAP SUSE cluster connector.

YaST, the main installation and administration tool of SUSE Linux Enterprise Server, provides a selection of software patterns for SAP that install the prerequisites needed for running SAP software. Installing a high available cluster using SUSE Linux Enterprise Server for SAP Applications is more convenient because all needed packages including the cluster and SAP-related packages are already included in one single product.

SLES for SAP includes SUSE Linux Enterprise Server Priority Support for SAP Applications. This support offering provides customers with a single support point of entry, from the operating system through the application (<https://www.suse.com/products/sles-for-sap/support/>). SLES for SAP also includes Extended Service Pack Overlap Support. Find more details about SLES for SAP at <https://www.suse.com/products/sles-for-sap/>.

In this document we use SUSE Linux Enterprise Server for SAP Applications 12 including the latest updates from SUSE Customer Center.

The following SUSE Linux Enterprise Server version 12 products including all service packs (SP) are certified as 64-bit version for SAP:

- SUSE Linux Enterprise Server for SAP Applications 12 for x86_64
- SUSE Linux Enterprise Server 12 for x86_64
- SUSE Linux Enterprise Server 12 for IBM System z (s390x)

SAP and SUSE are working together to ensure that SUSE Linux Enterprise Server service packs always match the certification of the respective product. In fact, SAP recommends always using the latest available service pack.

SUSE will provide at least ten years of general support for platform and operating system products, including its revisions, starting on the date of a product's general availability. When general support ends, SUSE will offer extended support for a minimum of three years. This gives SAP customers a long installation run-time, ensuring a low TCO.

3.3 SUSE Linux Enterprise High Availability Extension

Your data is the most valuable asset that you have — it is what your business depends on. Robust, scalable and manageable storage is a top priority for your IT department. High availability storage, applications, and services are critical for your business to be competitive. The SLE High Availability Extension comes with SLES for SAP and satisfies these needs. It includes high availability service and application clustering, file systems/clustered file systems, network file systems (NAS), volume managers, networked storage systems and drivers (SAN), and the management of all these components working together.

Unlike proprietary solutions, SLE High Availability Extension keeps costs low by integrating open source, enterprise-class components. The key components of the extension are:

- Pacemaker, a high availability cluster manager that supports multinode failover.
- Resource Agents to monitor and manage resources, compliant to the Open Cluster Framework (OCF).
- Oracle Cluster File System 2 (OCFS2), a parallel cluster file system that offers scalability.
- Cluster Logical Volume Manager (cLVM2), a logical volume manager for the Linux kernel, which provides a method of allocating space on mass storage devices that is more flexible than conventional partitioning schemes.
- Distributed Replicated Block Devices (DRBD) provides fast data resynchronization capabilities over LAN, replicated storage area network (SAN) semantics, allowing cluster-aware file systems to be used without additional SANs.
- High Availability Web Konsole (HAWK) and various command line tools.

Availability is a result of the interaction of cluster software with application services at the front and the operating system and hardware resources at the backend. Following this basic idea, cluster software like Pacemaker uses a lot of modules, such as services, resource agents, a messaging layer, network and file system availability, and a stable Linux kernel designed and configured for productive server systems in data centers.

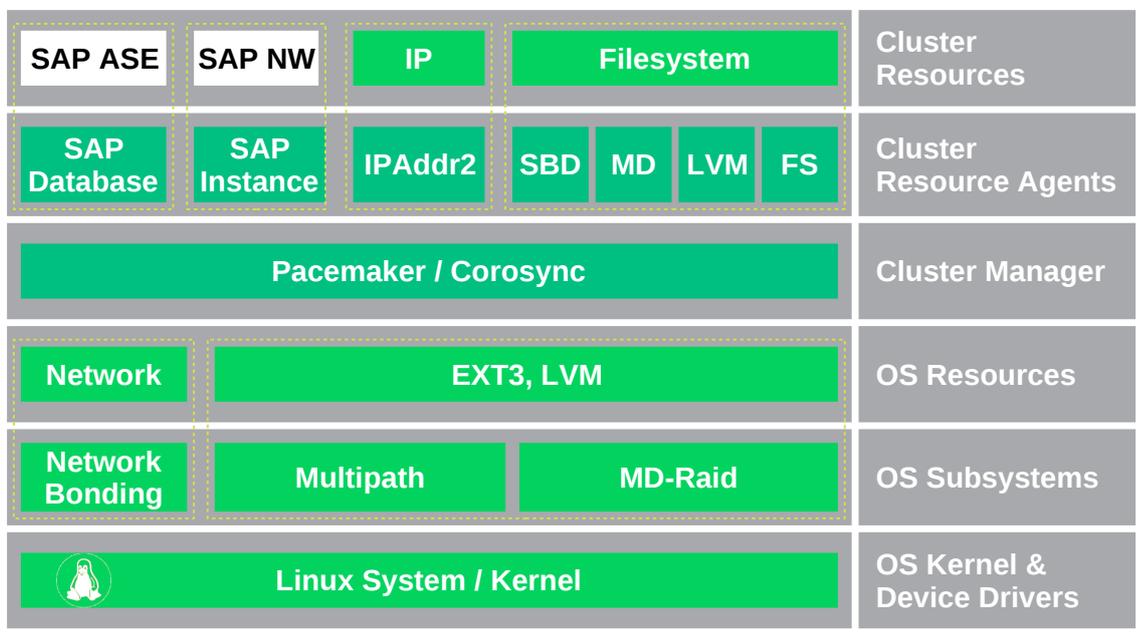


FIGURE 3.2: MODULES OF A HIGH AVAILABILITY SAP CLUSTER

The central application of our cluster is the SAP system itself. We need to provide the SAP database and the central SAP instance with high availability (white boxes). Operating system (light colored boxes) and cluster software (dark colored boxes) together give us the needed functionality.

In this document, SUSE Linux Enterprise Server for SAP Applications 12 with updates from SUSE Customer Center is used.

3.3.1 Storage Scenarios and I/O Stack for SAP

The file systems required by the operating system (i.e., mounted on `/` or `/var`) are usually stored on local hard drives using RAID-1 or RAID-5 disk arrays. Sometimes LVM is also used to gain more flexibility, i.e. for online file system extensions. This implies that these file systems can only be accessed by the local OS or applications running locally on this server. Even if these file systems could be placed on LUNs in a storage area network (SAN) outside the server, they are treated as local.

Besides the usual OS file systems, SAP and the SAP databases require their own file systems. These file systems are not stored locally. Instead they are provided by NFS file servers or on LUNs exported by storage units in storage area networks. Please be aware that the storage which the database resides on is subject to specifications and recommendations of the database vendor. Some file systems have to be shared across all cluster nodes such that every cluster node can access the file system at the same time. The recommended way of providing this service is the usage of a network file system (NFS). The NFS service also has to be highly available. In this document we use a dedicated NFS. This server could be either a third-party product or a SUSE Linux Enterprise High Availability Extension cluster that runs NFS resources. Under specific circumstances, it might be possible to integrate the NFS server into the SAP cluster. Both the separate NFS cluster as well as the integrated NFS server are described in separate documents.

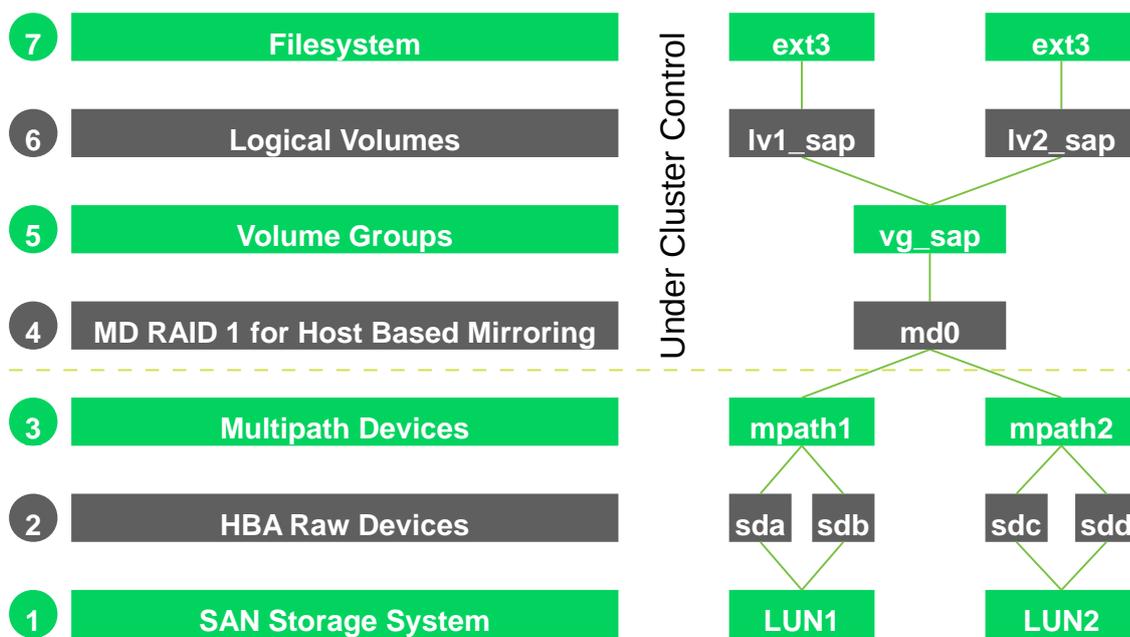


FIGURE 3.3: SAN IO LAYERS OF HIGH AVAILABILITY STACK FOR SAP

Other file systems, mounted to i.e. `/usr/sap/<SID>` or `/oracle/<SID>`, only have to be available on one cluster node at the same time. However, each cluster node must be able to access these file systems if the cluster manager decides to use them.

In our current set of Best Practices for SAP NetWeaver, we use LVM2 on top of MD RAID, which has shown its road capability for years. This storage stack can be used with or without cluster. The particular Best Practice document on replicated enqueue server does not need any MD RAID under control of the cluster.

The UNIX file system is the highest layer of a whole I/O stack consisting of multiple I/O layers. Each layer provides a certain kind of functionality. For all I/O critical tasks, we have configured an I/O stack that supports the following functions:

- Low latency: high I/O throughput and fast response times
- Host-based mirroring for storing data simultaneously on two separate storage units in a SAN
- Logical Volume Manager for a flexible management of file systems

- Multipath I/O for an additional level of redundancy for file systems, stored on LUNs in the SAN
- Online resizing (extending) of file systems, snapshots of file systems using LVM snapshots, moving or copying file system

While this guide focuses on well-known components, additional features of SUSE Linux Enterprise High Availability Extension allow for enhanced concepts, like the use of local area network-based Distributed Replicated Block Devices (DRBD) instead of storage area networks. Those concepts are described in additional documents.

3.3.2 SAP Resource Agents

The SAP Instance resource agent is responsible for starting, stopping, and monitoring the services in a SAP instance. The following services are monitored in the resource agent: `disp + work`, `msg_server`, `enserver`, `enrepservice`, `jcontrol`, and `jstart`.

The SAP Instance resource agent can be used to manage the following SAP instances:

- SAP NetWeaver AS ABAP Release 6.20 – 7.50
- SAP NetWeaver AS Java Release 6.40 - 7.50 (for 6.40 please also read SAP Note 995116)
- SAP NetWeaver AS ABAP + Java Add-In Release 6.20 - 7.50 (Java is not monitored by the cluster)

The resource agent is part of the SLE High Availability Extension and SLES for SAP. Another resource agent, the SAP Database resource agent, is also included. The purpose of the latter RA is to control the database instance of an SAP system. This resource agent expects a standard SAP installation and therefore needs fewer parameters to configure. The monitor operation of the resource agent can test the availability of the database by using SAP tools (`R3trans` or `jdbconnect`). This ensures that the database is really accessible for the SAP system.

3.3.3 The SAP SUSE cluster connector

The integration of the HA cluster through the SAP control framework using the `sap_suse_cluster_connector` is of special interest. "One of the classical problems running SAP instances in a highly available environment is that if a SAP administrator changes the status (start/

stop) of a SAP instance without using the interfaces provided by the cluster software than the cluster framework will detect that as an error status and will bring the SAP instance into the old status by either starting or stopping the SAP instance. This can result in very dangerous situations, if the cluster changes the status of a SAP instance during some SAP maintenance tasks. The solution is that the central component SAPSTARTSRV, which controls SAP instances since SAP Kernel versions 6.4, will be enabled to communicate the state change (start/stop) to the cluster software." (SAP SDN article "How to Connect SAPSTARTSRV and Cluster Frameworks using the Components saphascriptco.so and SAP_Vendor_Cluster_Connector" <http://scn.sap.com/docs/DOC-28875>).

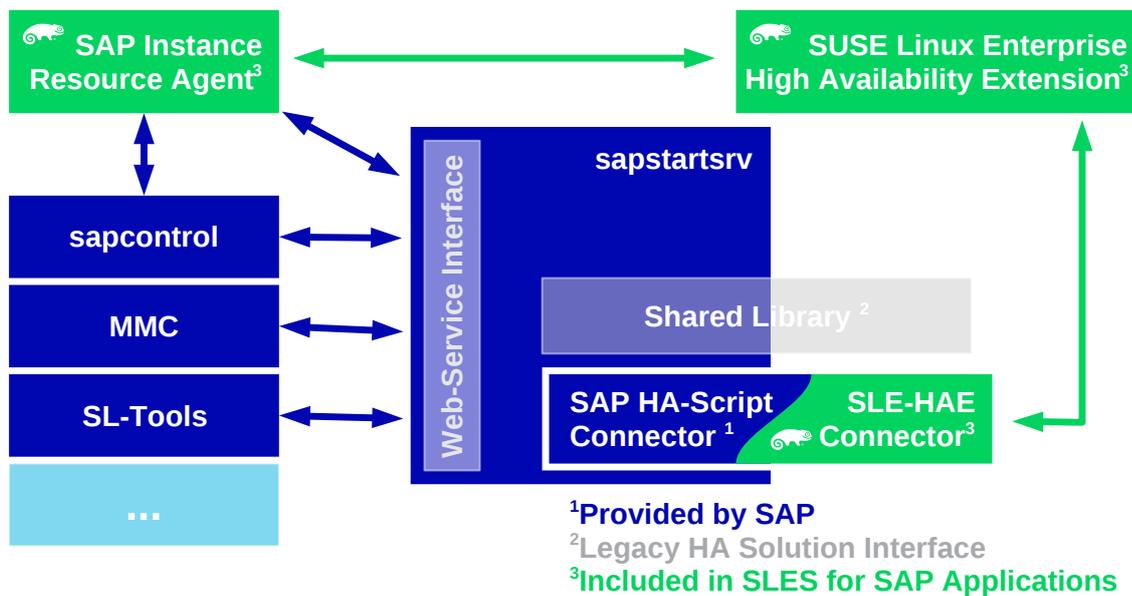


FIGURE 3.4: CLUSTER INTEGRATION WITH SAP CONTROL FRAMEWORK

3.3.4 Virtualization

The SLE High Availability Extension used in this document is supported on all major virtualization platforms, particularly current releases of VMware ESX. Detailed information about virtualization platform support for SLE High Availability Extension could be found on web pages listed in the appendix.

3.4 SAP ASE as RDBMS for SAP

Five major RDBMS (Relational Database Management System) are supported with SAP on SLES, on x86_64 either with or without the SUSE Linux Enterprise High Availability Extension. In addition, for SAP deployments with each of these databases, SAP priority support is available from SUSE:

- DB2 for LUW 10.5
- MaxDB 7.9
- Oracle 12.1
- SAP ASE 16.0
- SAP HANA DATABASE 1.0

Supported hardware platforms for all RDBMS on SLES for SAP 12 is x86_64 (also known as “x86 64-bit”, “AMD64”, “EM64T”).

SAP ASE was chosen as RDBMS for this document. It was mainly chosen to shorten the planning and installation procedure. Since the document is targeting a basic installation of SAP in a high availability cluster, the use cases are covered with SAP ASE. This installation example does not imply any decision about how one of the database management systems fits for a given workload.

SAP ASE (ASE) is a high-performance relational database management system for mission-critical, data-intensive environments. ASE key features include data encryption to protect from internal and external breaches; partitioning technology for better performance and easier maintenance; and, virtualization and clustering capabilities for continuous availability and efficient use of resources. In-memory database technology provides significant improvements in response time and throughput for high data volume and high concurrent user organizations. And the addition of data compression allows enterprises to reduce storage costs and increase performance on large and growing data sets.

SAP ASE product options are:

- ASE In-Memory Databases Option: fully integrated within ASE, the In-Memory Databases Option equips applications with instant responsiveness and very high throughput. The in-memory databases have zero-disk footprint and reside completely in memory.
- ASE Partitions Option: allows users to manage large tables and indexes by dividing them into smaller, more manageable pieces.

- ASE Encrypted Column Option: Allows data to be natively and selectively encrypted and stored with ASE.
- ASE Security and Directory Services Package: ensures data privacy through row-based access controls, the encryption of in-transit data, and support for LDAP, Active Directory and Pluggable Authentication Modules (PAM) services.
- ASE Active Messaging Option: provides a way to capture transactions in an Adaptive Server database and deliver them as events to external applications using either: JMS message bus, EAServer and Sonic Software; or Message Queue Interface (MQI) by WebSphere MQ.
- ASE High Availability Option: provides near continuous database access for critical business applications in the event of unexpected system failures.

Sybase's Linux commitment began in 1999 with the delivery of an enterprise-class DBMS on Linux and was the first to deliver a free enterprise-class commercial database that includes both pilot development and deployment. More information is available at: <http://go.sap.com/product/data-mgmt/sybase-ase.html>. As integrated database for the SAP NetWeaver technology platform, the respective certification applies for SLES on x86-64. For SAP systems, the appropriate product certification matrix should be applied (<https://support.sap.com/pam>). For support with SLE, see also SAP note 1554717 - SYB: Planning information for SAP on Sybase ASE.

SAP ASE installation media can be obtained from the SAP portal along with SAP NetWeaver. The installation is seamlessly integrated into the SAP installer. SAP offers several services around SAP ASE for SAP applications. More information can be found on the web pages listed in the appendix.

3.5 Sizing, Deployment, and Migration Considerations

Running SAP workloads in SUSE Linux Enterprise Server can reduce the TCO considerably. Through the homogenization of the data center (and thus the consolidation of resources), cost savings can be achieved even for highly specialized SAP installations. SUSE Linux Enterprise Server is very well prepared for automated installation and deployment, making it easy to deploy many servers. SUSE provides tools for managing updates and releases to ease day-to-day administration.

SUSE Linux Enterprise Server is very capable of providing the base for small or large systems. Customers run smaller central instances or larger distributed systems all with the same system base. It is quite possible to run multiple SAP instances in parallel on one system even when using high availability clusters. An increasing number of installations is done in virtualized environments. SUSE Linux Enterprise Server fits very well into such scenarios. SLES, SLE High Availability Extension, and SLES for SAP are supported on all major virtualization platforms. For information about virtualization platform support for SAP NetWeaver, please refer to SAP Note 1122387.

SAP requests that system sizing is done by the hardware vendor. SUSE has good relationships with many hardware vendors to make sure SUSE Linux Enterprise Server runs smoothly on a broad range of enterprise servers fit to run SAP workloads. SUSE and its partners are providing customers with solutions to their specific needs when it comes to Linux deployment. SUSE consulting has been developing best practices for high availability SAP installations and provides this information to customers and partners. Basic hardware requirements for a Proof of Concept on setting up SAP Enqueue Replication Server with SLES for SAP are given in a later chapter.

Thousands of successful SAP migrations to Linux have been made. The results regarding cost savings, performance and reliability have exceeded expectations in many instances. Since most data centers have adopted a Linux strategy, the know-how for deploying and administrating Linux systems is often in place and available. SAP-specific configurations and administration experience is available through SUSE consultation and partners. This makes the operating system side of the migration less risky and a ROI can be seen within the first six months of migration. SAP provides check lists and guidelines for the OS and database migration (see SAP Note 82478 - SAP system OS/DB migration).

4 One Concept — Two Use Cases

SAP NetWeaver installation can be adapted to several architectures for the entire system. You can plan and implement an “enqueue replication” SAP system with two dialog instances (i.e. PAS01 and D02) and the database on a different machine or cluster. After the SAP installation we add high availability components and bring the SAP system into a high availability scenario. There might also be reasons to install a SAP system with additional application servers. In all previously mentioned cases, SAP is treated in a two-tier manner, that means all SAP instances *could run* on the same system.

! Important

Throughout this document, all instance names (like ASCS00) or instance numbers (like 02) are examples only. Of course you can use the instance names and numbers of your SAP installation. The only rule is that your instances names and numbers match the SAP conventions, recommendations, and restrictions.

Installing an "enqueue replication" SAP system, we typically have the following types of instances in our cluster:

- ASCS00 (SCS20) - central service instance with message server and enqueue server,
- ERS10 - enqueue *replication* server,
- PAS01 - primary application server (ABAP or Dual-Stack only),
- D02, (D03, ...) - Dialog instance (ABAP or Dual-Stack only),
- J21, (J22, ...) - Java application server instance (Java only).

This document focuses on the ABAP application instances only.

Figure 4.1, “Clustered Enqueue Replication with External Database” shows a typical two node cluster running SAP NetWeaver including SAP Enqueue Replication Server. It is most important that the central service instance (ASCS00) and the enqueue replication instance (ERS10) are running on different nodes. At least one Dialog instance (PAS00 or D02) should run on a different node than ASCS00. Otherwise all dialog instances would fail together with the central service instance should the node crash. This would make the enqueue replication mechanism quite useless.

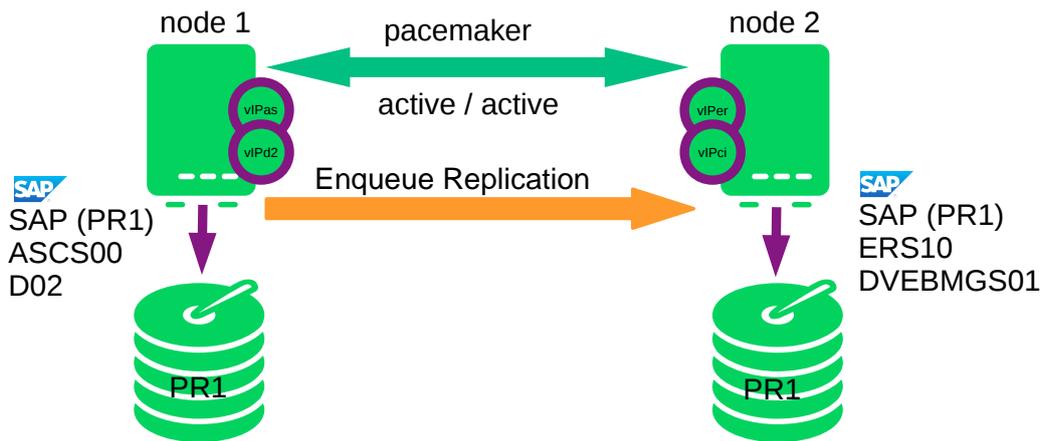


FIGURE 4.1: CLUSTERED ENQUEUE REPLICATION WITH EXTERNAL DATABASE

Another common use case is to run SAP Enqueue Replication Server and the database on the same cluster. In this case, high availability for the SAP Enqueue Replication Server is covered by the described concept as well, but you need to add the database to the cluster resources.

This document covers the following two use cases:

- Enqueue replication high availability external database,
- Enqueue replication high availability integrated database.

The installation of the first use case (Enqueue replication high availability external database) will be shown step-by-step in this document.

File systems and places for the SAP instances:

TABLE 4.1: FILE SYSTEMS AND PLACES

Mounted on	Installed on	Controlled by
/usr/sap/trans	NFS	autofs
/sapmnt/HA0	NFS	autofs
/usr/sap/HA0/SYS	NFS	autofs
/usr/sap/HA0/ASCS00	local on both nodes	OS

Mounted on	Installed on	Controlled by
/usr/sap/HA0/ERS10	local on both nodes	OS
/usr/sap/HA0/PAS01	NFS	cluster
/usr/sap/HA0/D02	NFS	cluster

Figure 4.2, "Clustered Enqueue Replication with External Database" shows the cluster status after node1 has been damaged completely. The cluster runs a list of actions:

- Missing node 1 is fenced via sbd. After timeout of the fencing action, node1 is assumed to be fenced.
- Node2 has to takeover resources formerly placed on node1.
- The slave instance of the master/slave construct is *promoted* to be the Master. This results in the ASCS00 instance running on node2. The cluster now has no enqueue replicator server anymore.
- The SAP Dialog Instance D02 is also *started* on node2.
- The cluster is in status idle again.

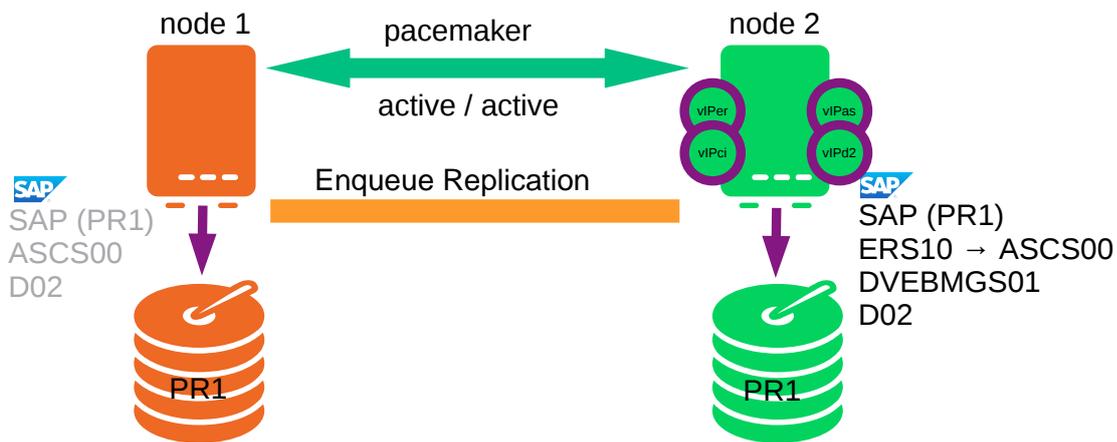


FIGURE 4.2: CLUSTERED ENQUEUE REPLICATION WITH EXTERNAL DATABASE

Figure 4.3, "Clustered Enqueue Replication with External Database", illustrates the status after node1 is rebooted and the cluster software is started by the administrator. The cluster has to run the following actions:

- Re-join of node1 into the cluster.
- Depending on scores, resources are fallback to node1 or not. In our configuration resources are not falling back.
- ERS10 (the enqueue replication server) is started as slave of the master/slave construct.
- The ASCS00 instance must stay on node2 as long as possible. It should not fallback to node1 without a new failure or only if told to do so by migration rules.

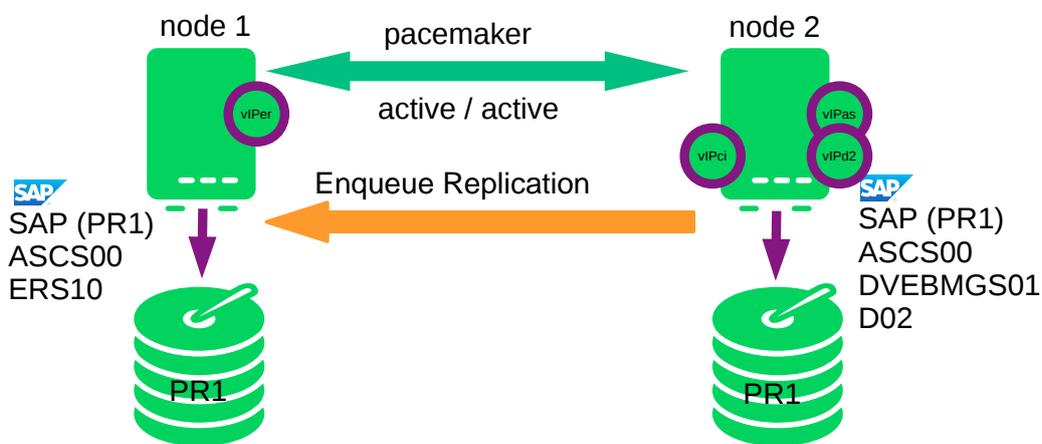


FIGURE 4.3: CLUSTERED ENQUEUE REPLICATION WITH EXTERNAL DATABASE

II Installation of Enqueue Replication with SAP ASE

- 5 Installation Overview **23**
- 6 Planning **24**
- 7 Prerequisites **25**
- 8 Get the needed SAP Installation Media **28**
- 9 Installation of SUSE Linux Enterprise Server for SAP Applications **31**
- 10 Prepare Storage for data and fencing **39**
- 11 Install SAP NetWeaver 7.50 **40**
- 12 Additional Work on the Second Node **50**
- 13 Configuration of SUSE Linux Enterprise High Availability Extension and SAP NetWeaver integration **52**
- 14 Administration and Troubleshooting **63**
- 15 Cluster Tests **65**

5 Installation Overview

This part describes the installation of a SAP NetWeaver “Enqueue Replication High Availability”, based on SAP ASE SUSE Linux Enterprise Server 12, for a proof of concept. The procedure is divided into the following steps:

- Planning: *Chapter 6, Planning.*
- Check prerequisites: *Chapter 7, Prerequisites.*
- Get the needed SAP NetWeaver installation media: *Chapter 8, Get the needed SAP Installation Media.*
- Install SUSE Linux Enterprise Server for SAP Applications 12: *Chapter 9, Installation of SUSE Linux Enterprise Server for SAP Applications.*
- Prepare SAN storage for SBD: *Chapter 10, Prepare Storage for data and fencing.*
- Install SAP NetWeaver 7.50 based on SAP ASE: *Chapter 11, Install SAP NetWeaver 7.50.*
- Preparation for SAP NetWeaver on second node *Chapter 12, Additional Work on the Second Node.*
- Configure SLE High Availability Extension and SAP NetWeaver integration: *Chapter 13, Configuration of SUSE Linux Enterprise High Availability Extension and SAP NetWeaver integration.*
- Check final result: *Chapter 14, Administration and Troubleshooting.*
- Recommended cluster tests: *Chapter 15, Cluster Tests.*

6 Planning

Proper planning is essential for a well performing SAP system. For planning and support of your SAP Installation, visit <http://support.sap.com> to download installation guides, review installation media lists, and to browse through the SAP Notes. This section focuses on aspects of planning a SAP installation.

The first major step is to size your SAP system and then derive the hardware sizing to be used for implementation. Use the SAP benchmarks (<http://www.sap.com/solutions/benchmark/index.epx>) to estimate sizing. It is also strongly advised to consult your hardware vendor. If you plan to migrate an existing SAP system, you should first obtain or estimate the system characteristics of the “old” SAP system. The key values of these characteristics include:

- SAPS of the old SAP system,
- memory (RAM) size and usage of the old hardware,
- disk size, performance and usage of the old SAP system,
- network performance and utilization of the old hardware,
- language support (including Unicode).

If you have valid key values, you can adapt these to the characteristics of your “new” SAP system. If you plan a new installation instead of a migration, you might need to adapt experiences with other SAP installations or use some of the published benchmarks as mentioned above.

- Estimate the SAPS in the new SAP system. This includes planning additional capacities, if needed. The calculation should also include estimated growth calculations, such as a boost of SAPS per year. Typical SAP growth is between 10% and 20% per year.
- Choose RAM size as well as disk size and performance for the Linux system. Also include a boost of the usage. Depending on the usage, the disk space may grow 30% per year. The disk size must also include the export and r3trans areas if they are not provided by other servers.
- Check if Unicode support is necessary for the new system.

7 Prerequisites

This chapter describes what hardware and software is needed for a proof of concept. It also outlines how to gather all information that is necessary to succeed.

7.1 Hardware Requirements

The final hardware sizing for production systems is always the responsibility of the hardware vendor. At the time this document was written, a base installation of SAP NetWeaver on top of SLES for SAP 12 for a proof of concept could be done on the following hardware equipment:

- 2* CPU x86-64 2GHz (AMD64 or EM64T),
- 8GB RAM,
- 40GB disk space for OS and SAP application,
- some hundred GB up to some TB additional disk space is needed for RDBMS and test data,
- 3* Gbit/s ethernet,
- 2 small (we only need some MB) LUNs in SAN for SBD. This LUNs must not be mirrored by the host (like MD-Raid, LVM, DRBD).

Although the cluster offers High Availability across two servers, each single server should already have the maximum possible availability by using redundant devices:

- self correcting ECC RAM,
- two or three redundant power supplies (connected to two circuits),
- several redundant cooling fans,
- two or more internal disks with RAID(1/5/6/10) controller,
- redundant LAN network controllers,
- redundant LAN network links (connected to two switches),
- redundant SAN host bus controllers,
- redundant SAN FC links (connected to two switches).

Make sure to use certified hardware. Information about certified hardware can be found in the SUSE YES database (<http://www.suse.com/yesssearch/Search.jsp>), in the SAP Notes, and on the hardware manufacturer's pages. Use certification notes from the SUSE YES database and the hardware manufacturer to select appropriate hardware components.

7.2 Software Requirements, Connection Data, etc.

Before you can start with the installation, make sure you meet the prerequisites.

- SAP documentation: Use the latest version of the Master Guide. Check SAP Service Marketplace immediately before starting the installation. The Master Guide is regularly updated at: <http://service.sap.com/instguidesnw75>.
- Sizing data. After planning the installation, you should have this data.
- Configuration data:
 - IP addresses, DNS names, Gateway,
 - SAP SIDs,
 - Linux User names (like <sid>adm), UIDs.
- Connection data:
 - SAN LUNs (names, lun numbers) and multipath configuration parameters. There are some special parameter settings for multipath and SAN-HBA kernel modules, depending on the hardware setup (SAN storage model and SAN setup). Check if SAN storages require partition alignment for performance reasons. Refer to the installation and configuration guides from SUSE and hardware vendors.
 - Access to the system management boards.
 - In addition to the network that connects the SAP servers to the clients, we recommend two additional dedicated network links between the two servers for cluster intercommunication. At least one additional dedicated network link is mandatory.

- Infrastructure such as DNS server, NTP server and a ping-able highly available network node. This network node could be the gateway between the SAP system and the clients who need to access the service. If the gateway is no longer available, the service is not available. The cluster can determine which cluster node has (ping) connection to the ping node and can migrate a service if needed.
- SAP Installation media (for details see the table in the next section). The SAP installation media can either be ordered as a physical CD/DVD or downloaded from <http://service.sap.com/swdc>. The next section describes the procedure for downloading the SAP media.
- SAP S-User (partner user) to download the media and installation guides and to browse through the SAP notes system. To download media, the S-User must have permission to download the installation media. Ask your company's SAP partner manager to create an S-User and to grant the proper rights.
- To download the SAP installation media, you will need the SAP download manager.
- To run the download manager you need a matching Java version. In former PoCs, IBM Java 1.7.1 (package `java-1_7_1-ibm-1.7.1_sr3.20-18.1.x86_64`) worked very well. Have a look at the installation notes presented during the procedure to download the SAP download manager.
- An up-to-date patch level of the SUSE Linux Enterprise Server 12 installation. You will need:
 - a SUSE Customer Center account,
 - "SUSE Linux Enterprise Server for SAP Applications 12" installation media for x86-64 hardware platform,
 - possibly some additional hardware-specific driver updates,
 - a software management software such as the subscription management tool (optional).
- To test the SAP system you either need to have a previously installed SAP client (**guilogon**, **guistart**) or you need to install this software on at least one of your workstations.

8 Get the needed SAP Installation Media

This section describes in a general way how to download the SAP installation media. As an example, we download the media for SAP NetWeaver 7.50 with database engine SAP ASE 16.0.02.01 RDBMS for SUSE Linux Enterprise Server 12 x86-64.

TABLE 8.1: SAP NETWEAVER INSTALLATION SOURCES

Title	Download Object	Type
NW 7.5 Installation Export	51050042_3	ZIP
SAP Kernel 7.45 Linux on x86_64 64bit	51050082_3	ZIP
SWPM 1.0 SP09 for NW higher than 7.0x	SWPM10SP09_6-20009701.SAR	SAR
SAP ASE 16.0.02.01 RDBMS Linux on x86_64 64bit	51050162_1	ZIP
SAPCAR_0-80000935.EXE	SAPCAR	EXE

We recommend using a workstation with SUSE Linux Enterprise Server 12 or SUSE Linux Enterprise Desktop 12 or newer to download the media. This section only describes the download using a SUSE Linux Enterprise system.

8.1 Download Procedure

1. Access the SAP download site <https://support.sap.com/swdc>.
2. You are asked to login with your S-User and password (unless you have setup single sign on).
3. From the menu on the left, select Installations and Upgrades → A - Z Index.
4. In the middle of the window, select "N" for SAP NetWeaver, then select SAP NetWeaver.
5. Now you have to select your NetWeaver version. In our example, we select SAP NetWeaver 7.5. → Installation and Upgrade → Linux → SAP ASE . The version for your use case may differ. This example will follow the names and paths of SAP NetWeaver 7.5.

6. From the media list, select the media listed above and click “Add to Download Basket”.
7. From the menu on the left, select Support Packages and Patches → A - Z Index.
8. In the middle of the window, select "S" for SWPM, then select SOFTWARE PROVISIONING MANAGER.
9. Now you have to select your SWPM version. In our example, we select SOFTWARE PROVISIONING MGR 1.0 → Support Package Patches → Linux on x86_64 64bit. The version for your use case may differ. This example will follow the names and paths of SWPM 1.0.
10. From the media list, select the media listed above and click “Add to Download Basket”.
11. From the menu on the left, select Support Packages and Patches → A - Z Index.
12. In the middle of the window, select "S" for sapcar, then select SAPCAR.
13. Now you have to select your SAPCAR version. In our example, we select SAPCAR 7.21 → Linux on x86_64 64bit. The version for your use case may differ. This example will follow the names and paths of SWPM 1.0.
14. From the media list, select the media listed above and click “Add to Download Basket”.
The total size of installation sources is about 10GB for the chosen SAP NetWeaver 7.5 with SAP ASE. To unpack the archives, roughly twice the disk space is needed. Other products might need more space.
15. After some time, a pop-up with two buttons appears. Press "Download Basket".
16. Your selected media is shown in your download basket. If you haven't installed the SAP download manager yet, you will have to download and install it now.
 - Click “Get Download Manager” in this case. The SAP Download Manager Installation Guide is shown.
 - You need a Java version that fits SAP needs. Download the Linux version.
 - Follow the SAP Download Manger Installation Guide
17. Download your Download Basket

- Start the installed SAP Download Manager using the command `java -classpath DLManager.jar dlmanager.Application`.
- If you start the SAP Download Manager for the first time, you will need to provide some credentials such as the SAP Marketplace address (<http://support.sap.com>), your S-User, your S-User-Password, and the Data Store (directory to place the downloaded files).
- Press the "Start scheduled download" button (the button with two right arrows).
- Now it's time to be very patient.

8.2 Extract downloaded Archives

1. After the download you will have to unpack the downloaded files using unzip (for ZIP type), unrar (for EXE and RAR type) and sapcar (for SAR type). Unrar is able to skip the self extracting code in the EXE files and will include all files of a media set such as 1/2, 2/2.
2. Copy (rsync) the extracted files to your system to be installed or create NFS exports on the installation source and NFS mounts on the target systems. In our setup we use:
 - /sapcd/SWPM for the Software Provisioning Manager,
 - /sapcd/Kernel for the NW 7.5 kernel,
 - /sapcd/Sybase for the SAP ASE engine,
 - /sapcd/InstExp SAP NW 7.5 installation export.

9 Installation of SUSE Linux Enterprise Server for SAP Applications

9.1 Installing SLES for SAP



Note

Please read our [SAP note 1984787 - SUSE LINUX Enterprise Server 12: Installation notes \(http://service.sap.com/sap/support/notes/1984787\)](http://service.sap.com/sap/support/notes/1984787) for instructions on how to set-up SUSE Linux Enterprise Server for SAP Applications for SAP workloads.

- Boot \$SUSE Linux Enterprise Server for SAP Applications Installation (DVD, PXE).
- Welcome: leave language unchanged EN_US.utf8; keyboard can be adjusted; accept License Terms.
- Installation Mode: Select "New Installation"
- Clock and Time Zone: adjust to your time zone. In this case, we have configured Europe/Germany; set Hardware Clock to UTC.
- Server Base Scenario: select "Physical Machine".
- Installation Setting:
 - Software: deselect AppArmor; You could chose GNOME.
 - Software: select SAP Application Server Base, C/C++ Compiler and Tools, and High Availability.
 - Expert-> default multi-user.target
 - Partitioning: (sda: 42GB root; 8GB swap;) (/dev/disk/by-id/lun-sbdA, /dev/disk/by-id/lun-sbdB: unpartitioned 10MB reserved-for-sbd).
- Installation runs (phase 1).

- Reboot.
- Enter the password for the system administrator. Your password must match the SAP password requirements.
- Hostname and Domain Name: <nodename>, lab.ha-workshop.org; deselect "Change Hostname via DHCP"; do not select "Assign Hostname to Loopback IP".
- Network Configuration:
 - Disable IPv6,
 - Disable Firewall,
 - Optional: Other IP Addresses.
- Test Internet Connection: select "No, Skip This Test".
- Network Services Configuration: no changes; use all defaults.
- User Authentication Method: Local.
- Enter a new local user
- Release Notes: Read and continue :).
- Hardware Configuration: no changes; use all defaults.
- Installation Completes: no changes; use all defaults.

9.2 Additional Configuration

- NTP must be set up - in our setup pointing to the local NTP server.
- All nodes and service addresses must be resolved in /etc/hosts.

9.3 File System Layout

Important

The sizes in this section are only values measured in our cluster to give you an idea about the space needed on a demo system. The disk usage of your SAP NetWeaver installation may differ. This is also the reason why we do not provide a size for the transport directory. This size depends extremely on the usage.

On both cluster nodes, we only need about 2-10 GB disk space for files stored locally under /usr/sap. You may want to create its own file system for /usr/sap/<SID>/ASCS00 and /usr/sap/<SID>/ERS10.

On the NFS filer, the cluster will also need approx. 10 GB disk space plus additional space for transports.

TABLE 9.1: FILE SYSTEMS AND PLACES

Mounted under	Installed on	Controlled by	Size in cluster
/usr/sap/trans	NFS	autofs	n/a
/sapmnt/HA0	NFS	autofs	3.5 GB
/usr/sap/HA0/SYS	NFS	autofs	1 MB
/usr/sap/HA0/ASCS00	local on both nodes	OS	300 MB
/usr/sap/HA0/ERS10	local on both nodes	OS	250 MB
/usr/sap/HA0/PAS01	NFS	cluster	1.1 GB
/usr/sap/HA0/D02	NFS	cluster	1 GB

9.4 Update Channels and Software Selection

The system should be connected to Update Channels. The installed software packages should be updated. Either use the GUI of SUSE's systems administration tool YaST or the command line tool zypper. In the following, we describe how to do this with zypper.

```
# SUSEConnect
```

You should then have the following Channels in your system.

```
cl2n01:~ # zypper lr
# | Alias          | Name                    | Enabled | GPG Check | Refresh
---+-----+-----+-----+-----+-----
 1 | ...            | SLE-12-SAP-12-0       | Yes    | ( p) Yes  | Yes
 2 | ...            | SLE-12-SAP-Updates    | Yes    | ( p) Yes  | Yes
 3 | ...            | SLE-HA12-Pool         | Yes    | ( p) Yes  | No
 4 | ...            | SLE-HA12-Updates      | Yes    | ( p) Yes  | Yes
 5 | ...            | SLE12-SAP-Debuginfo-Pool | No     | ----     | No
 6 | ...            | SLE12-SAP-Pool        | Yes    | ( p) Yes  | No
 7 | ...            | SLES12-Debuginfo-Pool | No     | ----     | No
 8 | ...            | SLES12-Debuginfo-Updates | No     | ----     | Yes
 9 | ...            | SLES12-Pool           | Yes    | ( p) Yes  | No
10 | ...            | SLES12-Updates        | Yes    | ( r ) Yes | Yes
```

Then you can update from the channels:

```
# zypper refresh
# zypper list-updates
# zypper update -l
```

Install the following additional packages, using either YaST or zypper:

- sbd,
- sap_suse_cluster_connector,
- java-1_7_1-ibm,
- perl-TermReadLine-Gnu,
- ClusterTools2.

```
# zypper install sbd sap_suse_cluster_connector java-1_7_1-ibm perl-TermReadLine-Gnu
ClusterTools2
```

9.5 Runlevel and Services

Enable automatic start of recommended services:

```
# systemctl enable ntpd
# systemctl enable multipathd
# systemctl enable autofs
# systemctl enable pacemaker
# systemctl enable uidd.socket
```

Optional services like HAWK can be activated:

```
# systemctl enable lighttpd
```

Setup of autofs. Adjust the following lines to `/etc/auto.master`:

```
+auto.master
/- /etc/auto.direct
```

File `/etc/auto.direct`:

```
/sapmnt/HA0 -nfsvers=3,nosymlink,sync sapha0nfs:/export/sapmnt/HA0
/usr/sap/trans -nfsvers=3,nosymlink,sync sapha0nfs:/export/usr/sap/trans
/usr/sap/HA0/SYS -nfsvers=3,nosymlink,sync sapha0nfs:/export/usr/sap/HA0/SYS
```

9.6 Miscellaneous

Under the following circumstances, the IO scheduler should be changed from default CFQ to noop:

- System is running in virtual machine.
- System uses intelligent RAID controller with cache.
- System uses LUNs from SAN.
- System uses multipath IO.

If in your setup, the systems are running as virtual machines and accesses LUNs on a intelligent NAS storage, the IO scheduler should be set to noop for all LUNs. This time we use YaST. Alternatively edit the files `/etc/default/grub`.

```
# yast
System
Boot Loader
Kernel Parameters
Optional Kernel Command Line Parameter
... elevator=noop
OK
OK
Quit
```

- Start the YaST Control Center: **yast**
- Select the system configuration **System**
- Select the **Boot loader** configuration
- Select the tab **Kernel Parameters**
- Select **Optional Kernel Command Line Parameter**
- Add the following parameter to the optional kernel command line parameters: **elevator=noop**
- Press **OK** (two times) and leave with **Quit**.

9.7 Check the Installation

Check that the correct base product with recommended software patterns is installed:

```
# cat /etc/products.d/baseproduct
# for f in base ha_sles x11 print_server sap_server Basis-Devel; do
    echo $f ; zypper search -t pattern | grep ^i...$f ;
done
```

Check that software update channels are set up:

```
# zypper repos
```

```
# zypper list-patches
```

Check that software packages have the correct version:

```
# for f in kernel-default multipath mdadm device-mapper glibc2 nfs-client libext2fs2
libuuid; do
    rpm -qa | grep $f;
done
# for f in sapconf sap-locale java gcc48 sbd corosync ; do
    rpm -qa | grep $f;
done
```

Check that recommended runlevel services are enabled:

```
# systemctl get-default | grep "multi-user.target"
# for f in ntpd multipathd autofs uidd.socket pacemaker; do
    systemctl is-enabled $f;
done
```

Check that unwanted runlevel services are disabled:

```
# for f in SuSEfirewall2; do
    systemctl is-enabled $f;
done
```

Check that the IO scheduler is noop:

```
# grep elevator /etc/default/grub
```

Check IP address and local name resolution:

```
# ip a s | grep global
# hostname -f
# grep -v "^#" /etc/hosts
```

Check NTP synchronization:

```
# grep server /etc/ntp.conf
# ntpq -p
```

Collect data for later re-use:

```
# cd
```

```
# supportconfig
# mv /var/log/nts_* .
# tar czf etc-$HOST.tgz /etc
# tar czf var-log-$HOST.tgz /var/log
# cs_sum_base_config >sum_base-$HOST.txt
# cp autoinst.xml autoinst.$HOST.xml
```

10 Prepare Storage for data and fencing

You need one or up to three tiny LUNs with capacity of 10 MB each. These are used as SBD devices later. The LUNs should not be partitioned.

11 Install SAP NetWeaver 7.50

Important

The described SAP NetWeaver installation procedure described in this document is only an example and does not replace the SAP installation guides available at the SAP Marketplace. The procedure is only described here to demonstrate how the reference system was set up.

11.1 Check operating system prerequisites

Check, if:

- a NFS server is installed and configured on another system,
- NFS File systems are mounted from NFS server (for the list of needed NFS file systems see also [Section 9.3, "File System Layout"](#)),
- all virtual host names (e.g. sapha0db, sapha0as, sapha0er, sapha0ci and sapha0d2) can be *locally* resolved on all nodes,

```
# from /etc/hosts
#...
172.16.12.190    sapha0db.lab.ha-workshop.org sapha0db
172.16.12.191    sapha0as.lab.ha-workshop.org sapha0as
172.16.12.194    sapha0er.lab.ha-workshop.org sapha0er
172.16.12.192    sapha0ci.lab.ha-workshop.org sapha0ci
172.16.12.193    sapha0d2.lab.ha-workshop.org sapha0d2
#...
```

- all virtual IP addresses (e.g. 172.16.12.190, 172.16.12.191, 172.16.12.194, 172.16.12.192 and 172.16.12.193) are configured on the installation node,

```
c12n01:~ # ip a s dev eth0
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
link/ether XX:XX:00:01:04:10 brd ff:ff:ff:ff:ff:ff
inet 172.16.12.101/24 brd 172.16.12.255 scope global eth0
inet 172.16.12.191/24 brd 172.16.12.255 scope global secondary eth0
inet 172.16.12.194/24 brd 172.16.12.255 scope global secondary eth0
inet 172.16.12.192/24 brd 172.16.12.255 scope global secondary eth0
inet 172.16.12.193/24 brd 172.16.12.255 scope global secondary eth0
inet6 fe80::5054:ff:fe01:410/64 scope link
```

```
valid_lft forever preferred_lft forever
```

- Database Host is installed and configured on another system (for example cl2n00),
- Database Host can be resolved and reached via IP address assigned to virtual hostname sapha0db.

11.2 Installing the ASCS Instance on Node cl2n01

- Create tmp directory for installation:

```
mkdir -p /tmp/NW750SYB/AS  
cd /tmp/NW750SYB/AS
```

- Start sapinst from the Software Provisioning Manager. Do not forget to use the parameter `SAPINST_USE_HOSTNAME = <Your-Virtual-HostName-For-ASCS>`. In our case, it is `sapha0as`. Do not use the nodes host name or the cluster would not be able to takeover that resource.

```
/sapcd/SWPM/sapinst SAPINST_USE_HOSTNAME=sapha0as
```

- Installation flow in sapinst.

- In the navigation tree select: SAP Netweaver 7.5 → SAP ASE → SAP Systems → Application Server ABAP → High-Availability System → ASCS Instance.
- Maybe some pop-ups about access rights for the tmp directory
- General SAP System Parameters: SAP System ID (SAPSID) HA0; SAP Mount Directory `/sapmnt`.
- DNS Domain Name: `lab.ha-workshop.org`.
- Media Browser - UC Kernel NW750: Browse to Kernel Media (`/sapcd/Kernel`).
- Enter the master password.
- SAP System Administrator: keep the settings
- ASCS Instance: keep the settings

- ABAP Message Server Ports: keep the settings
- Unpack Archives: keep the settings
- Parameter Summary: review all parameters.
- Task Progress: be patient.
- Pop-up: Completed.

The ASCS00 instance should be running now. Check by using the program **pgrep**:

```
cl2n01:~ # pgrep -fl HA0_ASCS
3496 sapstartsrv
3903 sapstart
3916 ms.sapHA0_ASCS00
3917 en.sapHA0_ASCS00
```

You can see that the message server and enqueue server of the SAP system HA0 and instance ASCS00 are running.

11.3 Installing SAP ASE Database

Important

The SAP ASE installation is carried out on the system planned for the database (in our case cl2n00). This node is not a member of the cluster discussed in this document, but the database could also be running in a highly available setup using SLE HAE.

- Create tmp directory for installation:

```
mkdir -p /tmp/NW750SYB/DB
cd /tmp/NW750SYB/DB
```

- Start sapinst from the Software Provisioning Manager. Do not forget to use the parameter `SAPINST_USE_HOSTNAME = <Your-Virtual-HostName-For-DB>`. In our case it is `sapha0db`. Do not use the node's host name.

```
/sapcd/SWPM/sapinst SAPINST_USE_HOSTNAME=sapha0db
```

- Installation flow in sapinst:
 - In the navigation tree select: SAP Netweaver 7.5 → SAP ASE → SAP Systems → Application Server ABAP → High-Availability System → Database Instance.
 - Maybe some pop-ups about access rights for the tmp directory.
 - SAP System Identification: Profiles Available: selected; Profile Directory: /usr/sap/HA0/SYS/profile.
 - Enter the master password.
 - SAP System Administrator: keep the settings
 - Media Browser - UC Kernel NW750: Browse to Kernel Media (/sapcd/Kernel).
 - SAP Database Administrator: keep the settings
 - Media Browser - Installation Export NW750 (folder EXP1): Browse to Export Folder EXP1 (/sapcd/InstExp).
 - Media Browser - SAP ASE: Browse to ASE Media, Folder SYBASE_LINUX_X86_64 (/sapcd/Sybase).
 - Automatic Database Expansion: keep the settings
 - SAP ASE Database System Parameters: keep the settings
 - Database User Passwords: keep the settings
 - Database Server: keep the settings
 - Declustering/Depooling Option: keep the settings
 - SAP ASE Database Statistics: keep the settings
 - Unpack Archives: select all
 - Parameter Summary: review all parameters.
 - Task Progress: be *very* patient.
 - Pop-up: Completed.

11.4 Installing Enqueue-Replication Instance on Node cl2n01

- Create tmp directory for installation:

```
mkdir -p /tmp/NW750SYB/ER
cd /tmp/NW750SYB/ER
```

- Start sapinst from the Software Provisioning Manager. Do not forget to use the parameter `SAPINST_USE_HOSTNAME = <Your-Virtual-HostName-For-ER>`. In our case it is `sapHA0er`. Do not use the node's host name or the cluster will not be able to takeover this resource.

```
/sapcd/SWPM/sapinst SAPINST_USE_HOSTNAME=sapha0er
```

- Installation flow in sapinst.

- SAP Netweaver 7.5 → SAP ASE → SAP Systems → Application Server ABAP → High-Availability System → Enqueue Replication Server Instance
- Maybe some pop-ups about access rights for the tmp directory.
- General SAP System Parameters: Profile Directory `/usr/sap/HA0/SYS/profile`.
- Existing SAP System Instances: HA0/ASCS00 selected.
- Media Browser - UC Kernel NW750: Browse to Kernel Media (`/sapcd/Kernel`).
- ERS Instance: number of the ERS Instance: 10.
- Activate Changes / Attention: "Automatic Instance and Service Restart" is selected.
- Parameter Summary: review all parameters.
- Task Progress: be patient.
- Pop-up: Completed.

The ERS10 instance should be running now. Use the program **pgrep** to check:

```
cl2n01:~ # pgrep -fl HA0_ERS
20659 sapstartsrv
21372 sapstart
```

You can see that the enqueue replication process for SAP system HA0 and instance ERS10 is running.

11.5 Installing Central Instance on Node cl2n01

- Create tmp directory for installation:

```
mkdir -p /tmp/NW750SYB/PAS
cd /tmp/NW750SYB/PAS
```

- Start sapinst from the Software Provisioning Manager. Do not forget to use the parameter `SAPINST_USE_HOSTNAME = <Your-Virtual-HostName-For-CI>`. In our case it is `sapHA0pas`. Do not use the node's host name or the cluster will not be able to takeover this resource.

```
/sapcd/SWPM/sapinst SAPINST_USE_HOSTNAME=sapHA0pas
```

- Installation Flow in sapinst

- SAP Netweaver 7.5 → SAP Systems → Application Server ABAP → High-Availability System → Primary Application Server Instance.
- Maybe some pop-ups about access rights for the tmp directory.
- General SAP System Parameters: Profile Directory `/usr/sap/HA0/SYS/profile`.
- Enter the master password.
- Media Browser - UC Kernel NW750: Browse to Kernel Media.
- Primary Application Server Instance: keep the settings
- ABAP Message Server Ports: keep the settings
- ICM User Management: keep the settings
- SLD Destination for Diagnostic Agent: Select "No SLD destination".
- Message Server Access Control List: keep the settings

- Media Browser: Installation Export
- SAP System DDIC Users: keep the settings
- Database User Password: use your database password
- Secure Storage Key Generation: Select "Default Key"
- Unpack Archives: Select all
- Install Diagnostic Agent: no
- Parameter Summary: review all parameters.
- Task Progress: be patient.
- Pop-up: Completed.

The PAS01 instance should be running now. Use the program pgrep to check:

```

cl2n01:~ # pgrep -fl HA0_D00
4473 sapstart
4493 HA0_00_DP
4494 ig.sapHA0_D00
4497 igsmux_mt
4498 igspw_mt
4499 igspw_mt
4504 gwr
4519 icman
4521 HA0_00_DIA_W0
4523 HA0_00_DIA_W1
4524 HA0_00_DIA_W2
4525 HA0_00_DIA_W3
4526 HA0_00_DIA_W4
4527 HA0_00_DIA_W5
4528 HA0_00_DIA_W6
4529 HA0_00_DIA_W7
4530 HA0_00_DIA_W8
4531 HA0_00_DIA_W9
4532 HA0_00_UPD_W10
4533 HA0_00_BTC_W11
4534 HA0_00_BTC_W12
4535 HA0_00_BTC_W13
4536 HA0_00_SP0_W14
4537 HA0_00_UP2_W15
27333 sapstartsrv

```

You can see that the dispatcher and work processes as well as the gateway, icman and igs-processes for SAP system HA0 and instance PAS01 are running.

11.6 Installing Additional Dialog Instances on Node cl2n01

- Create tmp directory for installation:

```
mkdir -p /tmp/NW750SYB/D02  
cd /tmp/NW750SYB/D02
```

- Start sapinst from the Software Provisioning Manager. Do not forget to use the parameter `SAPINST_USE_HOSTNAME = <Your-Virtual-HostName-For-D2>`. In our case it is `sapHA0d2`. Do not use the node's host name or the cluster will not be able to takeover this resource.

```
/sapcd/SWPM/sapinst SAPINST_USE_HOSTNAME=sapha0d2
```

- Installation flow in sapinst.
 - SAP Netweaver 7.5 → SAP Systems → Application Server ABAP → High-Availability System → Additional Application Server Instance.
 - Maybe some pop-ups about access rights for the tmp directory.
 - General SAP System Parameters: Profile Directory `/usr/sap/HA0/SYS/profile`.
 - Media Browser - UC Kernel NW750: Browse to Kernel Media.
 - Enter the master password.
 - Additional Application Server Instance: keep the settings
 - Message Server Access Control List: keep the settings
 - Unpack Archives: select all
 - Install Diagnostics Agent: keep the settings
 - Parameter Summary: review all parameters.
 - Task Progress: be patient.
 - Pop-up: Completed.

11.7 Update SAP Kernel and SAPHOSTAGENT Package

After the SAP NetWeaver installation, you should at least update SAP kernel and SAPHOSTAGENT. Refer to the SAP instructions available at the SAP Marketplace for download and update procedure.

11.8 SAP Profile Changes to Match the Master-Slave Scenario

In the master-slave scenario, the enqueue server must not be restarted locally. This is why we need to change the start directive for the enqueue server in the ASCS instance profile (/sapmnt/HA0/profile/HA0_ASCS00_sapha0as):

```
#-----  
# Start SAP message server  
#-----  
_MS = ms.sap$(SAPSYSTEMNAME)_$(INSTANCE_NAME)  
Execute_02 = local rm -f $_MS  
Execute_03 = local ln -s -f $(DIR_EXECUTABLE)/msg_server$(FT_EXE) $_MS  
Restart_Program_00 = local $_MS pf=$_PF  
#-----  
# Start SAP enqueue server  
#-----  
_EN = en.sap$(SAPSYSTEMNAME)_$(INSTANCE_NAME)  
Execute_04 = local rm -f $_EN  
Execute_05 = local ln -s -f $(DIR_EXECUTABLE)/enserver$(FT_EXE) $_EN  
Start_Program_01 = local $_EN pf=$_PF  
Max_Program_Restart = 03
```

The changes made here are:

- Use of **Start_Program_nn** for the enqueue server instead of **Restart_Program_nn** to avoid the local restart of the enqueue server after the process has failed.
- The line **Max_Program_Restart = 03** defines that **sapstart** only restarts processes three times. This is useful for takeover of the ASCS instance if the message server is failing multiple times.

11.9 Integrating the `sap_suse_cluster_connector`

To integrate the `sap_suse_cluster_connector`, you have to allow the `<sid>adm` user to communicate with the cluster. Additionally, the SAP command `sapstartsrv` must be instructed to load and use the `saphascriptco` library.

1. Add all `<sid>adm` users to group `haclient`. Example for `ha0adm`:

```
usermod -aG haclient ha0adm
```

2. The specific values of the profile parameters differ, depending on the version of the SAP kernel:

SAP Instance Profile Changes for dialog instances:

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

You need at least `sap_suse_cluster_connector` version 1.1.x.

12 Additional Work on the Second Node

12.1 Duplicate SAP Related Files to the Second Node

Create users and groups on the other node as they were created by the SAP installation on the first node. Use the same user ID and group ID. Be careful to use the same user-IDs and group-IDs in `/etc/passwd` and `/etc/group`.

Check login user profiles of `<sid>adm` and `sqd<sid>` on `<node1>`; they should be similar. Depending on the SWPM used for the SAP installation, the login profiles for the SAP Administrator user (`<sid>adm`) and the database administrator user might be different. In older and non-high availability installations, the user login profiles look similar to `.sapenv_hostname.csh`. Optionally repair profiles. Ensure that the same user login profiles for the `<sid>adm` and `sqd<sid>` user are on all nodes. You can do so by copying the `/home/<sid>adm` to `<node2>`.

Copy the `/etc/services` file or its values, which were adjusted by the `sapinst` (see SAP related entries at the end of the file), to all nodes.

There are other directories within the SAP file system, which have to be configured. These directories belong to specific SAP services, so their configuration depends on the particular SAP landscape. To set up systems quickly in no greater context or SAP landscape, it is sufficient to just copy them to the other node.

Create all needed mount points on the second node (`/sapmnt`, `/sapdb`, `/usr/sap`).

An empty work directory (`/usr/sap/<SID>/<Instance><Number>/work`) of an SAP instance leads to a monitoring error of the SAPIInstance resource agent. Every instance has to be started manually once so the correct entries will be written to the work directory. After that you can do a manual shutdown of the instances and then the cluster is ready to control them.

`sapinst` will create a script `/etc/rc.d/sapinit` and a configuration file `/usr/sap/sapservices`. Both are not used by the SAPIInstance Resource Agent. Make sure `/etc/rc.d/sapinit` is consistent on both nodes but not activated in a runlevel for startup of the operating system.

12.1.1 Additional files and directories for MaxDB

In the described installation we use ASE and could skip this section. However, if you like to integrate MaxDB into the cluster, you also need to prepare the following steps:

In case of choosing MaxDB as database, files and directories have to be synchronized, too. Copy the file /etc/opt/sdb and the directory structure /usr/spool/sql to the other node.

Make the directories (with their content) /sapdb/programs/lib and /sapdb/programs/runtime available even if the file system /sapdb is not mounted. To do so, mount /sapdb, copy the directories to a temporary directory, **unmount /sapdb** and copy them locally to /sapdb/programs/lib and /sapdb/programs/runtime. Do so on every node.

13 Configuration of SUSE Linux Enterprise High Availability Extension and SAP NetWeaver integration

13.1 Basic Cluster Configuration

The first step is to setup the base cluster framework. For convenience, use YaST2 or the `ha-cluster-init` script. Depending on your environment, you might need to adjust the SBD configuration later and select an appropriate hardware watchdog. You should later add a second ring to the cluster communication and also change to UCAST communication as well.

Create an initial setup, using `ha-cluster-init`:

```
# ha-cluster-init
Enabling sshd.service
Generating ssh key
Configuring csync2
Generating csync2 shared key (this may take a while)...done
Enabling csync2.socket
csync2 checking files

Configure Corosync:
This will configure the cluster messaging layer. You will need
to specify a network address over which to communicate (default
is eth0's network, but you can use the network address of any
active interface), a multicast address and multicast port.

Network address to bind to (e.g.: 192.168.1.0) [10.20.88.0] 192.168.124.0
Multicast address (e.g.: 239.x.x.x) [239.249.189.242]
Multicast port [5405]

Configure SBD:
If you have shared storage, for example a SAN or iSCSI target,
you can use it avoid split-brain scenarios by configuring SBD.
This requires a 1 MB partition, accessible to all nodes in the
cluster. The device path must be persistent and consistent
across all nodes in the cluster, so /dev/disk/by-id/* devices
are a good choice. Note that all data on the partition you
specify here will be destroyed.

Do you wish to use SBD? [y/N] y
Path to storage device (e.g. /dev/disk/by-id/...) [] /dev/disk/by-id/lun-sbdA
All data on /dev/disk/by-id/lun-sbdA will be destroyed
Are you sure you wish to use this device [y/N] y
Initializing SBD.....done
```

```
Enabling hawk.service
HA Web Konsole is now running, to see cluster status go to:
https://10.20.88.51:7630/
Log in with username 'hacluster', password 'linux'
WARNING: You should change the hacluster password to something more secure!
Enabling pacemaker.service
Waiting for cluster.....done
Loading initial configuration
Done (log saved to /var/log/ha-cluster-bootstrap.log)

Change the hacluster password
```

So far we have configured the basic cluster framework including:

- ssh keys,
- csync2 to transfer configuration files,
- SBD (at least one device),
- hardware watchdog (or softdog in some kinds of VM),
- corosync (at least one ring),
- HAWK web interface.

As requested by ha-cluster-init, we change the passwords of the user hacluster.

13.2 Adapting the Configuration Manually

13.2.1 Change the corosync configuration to UCAST and two rings

For the second ring we add a second interface definition inside of "totem".

For the unicast communication we add the transport definition inside "totem" (transport: udpu).

```
# /etc/corosync/corosync.conf
totem {
    ...
```

```

interface {
    #Network Address to be bind for this interface setting
    bindnetaddr:    192.168.124.0

    #The multicast port to be used
    mcastport:     5405

    #The ringnumber assigned to this interface setting
    ringnumber:    0

    #Time-to-live for cluster communication packets
    ttl:           1

}
interface {
    #Network Address to be bind for this interface setting
    bindnetaddr:    192.168.125.0

    #The multicast port to be used
    mcastport:     5405

    #The ringnumber assigned to this interface setting
    ringnumber:    1

    #Time-to-live for cluster communication packets
    ttl:           1

}
...

#Transport protocol
transport:        udpu
}

```

For the unicast communication we also need a nodelist with all communication addresses (here 2 per node). The section is called nodelist. The nodelist is at "top-level" of the configuration (same as totem and others).

```

totem {
    ...
}
nodelist {
    node {

```

```

#ring0 address
ring0_addr:    192.168.124.1

#ring1 address
ring1_addr:    192.168.125.1

}
node {
#ring0 address
ring0_addr:    192.168.124.2

#ring1 address
ring1_addr:    192.168.125.2

}
}

```

13.2.2 Setup the use of more than one sbd

To set up the use of more than one sbd, we add the other device to the variable **SBD_DEVICE** in `/etc/sysconfig/sbd`. The device paths are separated by semi-colon (;).

```

# /etc/sysconfig/sbd
SBD_DEVICE="/dev/disk/by-id/lun-sbdA;/dev/disk/by-id/lun-sbdB"

```

It is required to use a watchdog together with SBD. Therefore the **-W** option is needed. With current pacemaker packages, you can also use the **-P** option (*Check Pacemaker quorum and node health*), which enables the cluster nodes not to self-fence if SBD devices are lost, but pacemaker communication is still available. If **-S** is set to **1** (one), the sbd daemon will only start if the node was previously shut down cleanly (as indicated by an exit request message in the slot), or if the slot is empty.

TABLE 13.1: SBD OPTIONS

Parameter	Description
-W	Use watchdog. It is mandatory to use a watchdog. SBD does not work reliable without watchdog. Please refer to the SLES manual and SUSE TIDs 7016880 for setting up a watchdog. This is equivalent to <code>SBD_WATCHDOG="yes"</code>
-S 1	Start mode. If set to one, sbd will only start if the node was previously shutdown cleanly or if the slot is empty. This is equivalent to <code>SBD_STARTMODE="clean"</code>
-P	Check Pacemaker quorum and node health. This is equivalent to <code>SBD_PACEMAKER="yes"</code>

```
# vi /etc/sysconfig/sbd
# /etc/sysconfig/sbd
# Pls. use /dev/disk/by-id/... here. No trailing ";".
SBD_DEVICE="/dev/disk/by-id/lun-sbdA;/dev/disk/by-id/lun-sbdB"
# Pls. mind the gap between -S and 1.
SBD_WATCHDOG="yes"
SBD_PACEMAKER="yes"
SBD_STARTMODE="clean"
SBD_OPTS=""
```

 Note

This equates to the SUSE Linux Enterprise 11 settings `SBD_OPTS="-W -P -S 1"`.

Further information on how to configure SBD and activate the correct watchdog is given in the product documentation at https://www.suse.com/documentation/sle-ha-12/single-html/book_sleha/book_sleha.html#sec.ha.storage.protect.fencing and <https://www.suse.com/support/kb/doc.php?id=7016305>.

You might also check manpages `stonith_sbd(7)` and `sbd(8)` for complete documentation of the available options.

13.3 Join the 2node to the Cluster

The SBD daemon always needs a watchdog. The appropriate watchdog kernel module could be defined manually. For this document, the setup was done in a virtual machine. Therefore we use the Linux kernel's software watchdog module `softdog.ko`. On physical machines another watchdog might be used.

```
# vi /etc/modules-load.d/watchdog.conf
softdog
# modprobe softdog
```

If the watchdog is activated, the respective kernel module and a device node has to show up.

```
# lsmod | grep -e wdt -e dog
softdog
# ls -l /dev/watchdog
crw-rw---- 1 root root 10, 130 Mar 08 08:08 /dev/watchdog
# grep -e wdt -e dog /etc/modules-load.d/watchdog.conf
softdog
```

If the watchdog module is not automatically loaded during system boot, please add **`modprobe watchdog-module`** to `/etc/init.d/boot.local`. For example for `softdog`:

```
modprobe softdog
```

Further information on how to activate the correct watchdog on a certain machine is given in the product documentation at https://www.suse.com/documentation/sle-ha-12/single-html/book_sleha/book_sleha.html#pro.ha.storage.protect.watchdog.

13.4 Start the Cluster for the first Time

Check on both nodes, if the cluster is running:

```
systemctl status pacemaker
```

If the cluster is not running, it's time to start the cluster for the first time on both nodes:

```
systemctl start pacemaker
```

Check the cluster status with `crm_mon`. We use the option `-r` to also see resources, which are configured but stopped:

```
crm_mon -r
```

The command will show the "empty" cluster and will print something like in the following screen output. The most interesting information for now is that there are two nodes in status "online" and the message "partition with quorum".

```
=====
Last updated: Fri Sep  7 18:34:49 2012
Last change: Thu Sep  6 14:48:24 2012 by ha0adm via cibadmin on cl2n01
Stack: openais
Current DC: cl2n01 - partition with quorum
Version: 1.1.6-b988976485d15cb702c9307df55512d323831a5e
2 Nodes configured, 2 expected votes
10 Resources configured.
=====

Online: [ cl2n02 cl2n01 ]

stonith-sbd      (stonith:external/sbd): Started cl2n01
```

Some more checks like `cs_make_sbd_devices --dump`.

```
# cs_make_sbd_devices --dump
==Dumping header on disk /dev/disk/by-id/lun-sbdA
Header version      : 2
Number of slots     : 255
Sector size         : 512
Timeout (watchdog)  : 30
Timeout (allocate)  : 2
Timeout (loop)      : 2
Timeout (msgwait)   : 130
==Header on disk /dev/disk/by-id/lun-sbdA is dumped
0      cl2n02  clear
1      cl2n01  clear
==Dumping header on disk /dev/disk/by-id/lun-sbdB
Header version      : 2
Number of slots     : 255
Sector size         : 512
```

```
Timeout (watchdog) : 30
Timeout (allocate) : 2
Timeout (loop)     : 2
Timeout (msgwait)  : 130
==Header on disk /dev/disk/by-id/lun-sbdb is dumped
0      cl2n01  clear
1      cl2n02  clear
```

13.5 Configure Bootstrap and STONITH Using crm Shell

There are multiple ways to configure resources in a pacemaker cluster. In the following we describe the method using the crm shell. For all of the following sections to configure parts of the cluster proceed the following steps:

- Start the crm shell and enter the configure action: **crm configure**
- Copy and paste and all prepared crm configuration sequences into the crm shell
- Quit the crm shell (**quit**). You might be asked, if you want to commit your changes.

An other method is to save the crm configuration sequences into a text file adding **configure** at the beginning of each object (like **configure property ...**). The created text file could then be loaded to the cluster using **crm -f crm-command-file**

We define some cluster wide properties. The most important definitions are to activate STONITH, to define the no-quorum-policy, and to adjust the STONITH timeout values. The values here are examples only, which work in our demo cluster. These values must be tuned to your hardware and SAN environment.

```
property $id="cib-bootstrap-options" \  
    stonith-enabled="true" \  
    no-quorum-policy="ignore" \  
    stonith-action="reboot" \  
    stonith-timeout="150s"
```

For resources, we define the default stickyness of 1000. Depending on other scoring values, this should tell the cluster not to move resources without having either a loss of a node or an explicit migration request.

```
rsc_defaults $id="rsc-options" \  
    stickiness="1000"
```

```
resource-stickiness="1000" \  
migration-threshold="5"
```

Now we define the operation defaults. We define 10 minutes (600 seconds) as default timeout and we set the cluster to show pending actions. This is needed by the `sap_suse_cluster_connector` for proper cluster communication.

```
op_defaults $id="op-options" \  
  timeout="600" \  
  record-pending="true"
```

13.6 Configure Resources Using crm Shell

First we define the group of resources needed, before the ASCS instance can be started. The most important parameter here are `ip="172.16.12.191"`. This of course needs to be adapted to your environment.

```
primitive rsc_ip_HA0_sapha0as ocf:heartbeat:IPaddr2 \  
  params ip="172.16.12.191" \  
  op monitor interval="10s" timeout="20s" on_fail="restart" \  
group grp_sap_as_HA0 rsc_ip_HA0_sapha0as \  
  meta target-role="Started" is-managed="true" \  
  resource-stickiness="1000"
```

Now we define the group of resources needed, before the ERS instance can be started. The most important parameter here are `ip="172.16.12.194"`. This of course needs to be adapted to your environment.

```
primitive rsc_ip_HA0_sapha0er ocf:heartbeat:IPaddr2 \  
  params ip="172.16.12.194" \  
  op monitor interval="10s" timeout="20s" on_fail="restart" \  
group grp_sap_er_HA0 rsc_ip_HA0_sapha0er \  
  meta target-role="Started" is-managed="true" \  
  resource-stickiness="1000"
```

The next step is to define the master/slave resource to have an enqueue/enqueue-replication construct. The most important parameters are the setting `InstanceName`, `START_PROFILE` (with full path) and the equivalent parameters for the ERS instance.

```
primitive rsc_sap_HA0_ASCS00 ocf:heartbeat:SAPInstance \  
  operations $id="rsc_sap_HA0_ASCS00-operations" \  
  op monitor interval="11" role="Slave" timeout="60" \  
  op monitor interval="13" role="Master" timeout="60" \  
  resource-stickiness="1000"
```

```

params \
  InstanceName="HA0_ASCS00_sapha0as" \
  START_PROFILE="/usr/sap/HA0/SYS/profile/HA0_ASCS00_sapha0as" \
  ERS_InstanceName="HA0_ERS10_sapha0er"
  ERS_START_PROFILE="/usr/sap/HA0/SYS/profile/HA0_ERS10_sapha0er"

ms msl_sap_enqrepl_HA0 rsc_sap_HA0_ASCS00 \
  meta clone-max="2" target-role="Started" master-max="1" \
  is-managed="true"

```

To add the primary application server, we need to define all depending resources in a group. Parameters to be adapted here are: the ip-address, the mount-point and device as well as the instance name and start profile.

```

primitive rsc_ip_HA0_sapha0ci ocf:heartbeat:IPaddr2 \
  params ip="172.16.12.192" \
  op monitor interval="10s" timeout="20s" on_fail="restart"

primitive rsc_fs_HA0_pas01 ocf:heartbeat:Filesystem \
  operations $id="rsc_fs_HA0_pas01-operations" \
  op monitor interval="20" timeout="40" \
  params device="sapha0nfs:/export/usr/sap/HA0/PAS01" fstype="nfs" \
  directory="/usr/sap/HA0/PAS01"

primitive rsc_sap_HA0_PAS01 ocf:heartbeat:SAPInstance \
  operations $id="rsc_sap_HA0_PAS01-operations" \
  op monitor interval="120" timeout="60" \
  params InstanceName="HA0_PAS01_sapha0ci" \
  START_PROFILE="/usr/sap/HA0/SYS/profile/HA0_PAS01_sapha0ci"

group grp_sap_ci_HA0 rsc_ip_HA0_sapha0ci rsc_fs_HA0_pas01 \
  rsc_sap_HA0_PAS01

```

For each additional SAP application server, we need to define an additional group of resources. Same parameters to be adapted, now for the second dialog instance.

```

primitive rsc_ip_HA0_sapha0d2 ocf:heartbeat:IPaddr2 \
  params ip="172.16.12.193" \
  op monitor interval="10s" timeout="20s" on_fail="restart"

primitive rsc_fs_HA0_d02 ocf:heartbeat:Filesystem \
  operations $id="rsc_fs_HA0_d02-operations" \
  op monitor interval="20" timeout="40" \
  params device="sapha0nfs:/export/usr/sap/HA0/D02" \
  directory="/usr/sap/HA0/D02" fstype="nfs"

primitive rsc_sap_HA0_D02 ocf:heartbeat:SAPInstance \
  operations $id="rsc_sap_HA0_D02-operations" \
  op start interval="0" timeout="240" \
  op monitor interval="120" timeout="60" \
  params InstanceName="HA0_D02_sapha0d2"
  START_PROFILE="/usr/sap/HA0/SYS/profile/HA0_D02_sapha0d2" \

```

```
meta target-role="Started"

group grp_sap_d2_HA0 rsc_ip_HA0_sapha0d2 rsc_fs_HA0_d02 rsc_sap_HA0_D02
```

Check if you have the stonith-sbd resource defined like this. It is important to set the timeout and the start-delay parameter for the start operation.

```
primitive stonith-sbd stonith:external/sbd \
    op start interval="0" timeout="15" start-delay="5"
```

Now we have defined all needed resources.

13.7 Configure Constraints Using crm Shell

Now we define the rules to place the resources for an initial cluster start. We have to score all SAP resource groups:

```
location loc_grp_sap_as_HA0_cl2n01 grp_sap_as_HA0 10: cl2n01
location loc_grp_sap_di_HA0_cl2n01 grp_sap_as_HA0 10: cl2n01
location loc_grp_sap_ci_HA0_cl2n02 grp_sap_ci_HA0 10: cl2n02
```

To "bind" the IP address for the ASCS instance to the Master:

```
colocation col_grp_sap_as_HA0_msl_sap_enqrepl_HA0_MASTER 2000: \
    grp_sap_as_HA0 msl_sap_enqrepl_HA0:Master
```

To "bind" the IP address for the ERS instance to the Slave:

```
colocation col_grp_sap_er_HA0_msl_sap_enqrepl_HA0_SLAVE 2000: \
    grp_sap_er_HA0 msl_sap_enqrepl_HA0:Slave
```

And finally we define the start/stop order of resources:

```
order ord_grp_sap_as_HA0_msl_sap_enq_repl Optional: grp_sap_as_HA0:start \
    msl_sap_enqrepl_HA0:promote symmetrical=true
order ord_msl_sap_enqrepl_HA0_promote_rsc_sap_HA0_D02_start Optional: \
    msl_sap_enqrepl_HA0:promote rsc_sap_HA0_D02
order ord_msl_sap_enqrepl_HA0_promote_rsc_sap_HA0_PAS01_start Optional: \
    msl_sap_enqrepl_HA0:promote rsc_sap_HA0_PAS01
```

14 Administration and Troubleshooting

TABLE 14.1: ADMINISTRATION REMARKS AND TROUBLE SHOOTING

Case	Action/Remark
How to handle sap-startsrv if it has been stopped before? sapcontrol output like "FAIL: NIECONN_REFUSED"	Cluster procedure: login as root and call <u>ClusterService RUP <res></u> If the side had sapstartsrv active before, login as sidam and call <u>sapcontrol -nr <nr> -function StartService <SID></u>
How to clean resources failures?	If crm_mon -1 outputs also failed actions, you can clean this status by running <u>ClusterService RCL <res></u> For clones or master/slave resources, use the primitive instance name like rsc_sapHA0_ASCS00:0. This only cleans the status of the failed resource and does not clean-up the status of the complete master/slave construct. The clean-up procedure is more robust, because instances without a failure are not touched.
Procedure to set complete cluster into maintenance mode for SAP.	If you completely need to deactivate the cluster software, perform the following procedure: <ul style="list-style-type: none"> • For all resource groups and master/slave resources: <u>ClusterService RSU <grp-or-ms></u> • Wait until cluster status is <u>S_IDLE</u>. • Stop cluster software on <i>second</i> node: <u>rcopenais stop</u> • Stop cluster software on <i>first</i> node: <u>rcopenais stop</u>

Case	Action/Remark
<p>Get cluster back to production after the cluster has been set to maintenance mode.</p>	<p>Get cluster back to production after the cluster has been set to maintenance mode.</p> <ul style="list-style-type: none"> • Start cluster software on <i>first</i> node: <u>rcopenais start</u> • After some seconds, start cluster software on <i>second</i> node: <u>rcopenais start</u> • For all resource groups and master/slave resources: <u>ClusterService RSM <grp-or-ms></u> • Wait until cluster status is <u>S_IDLE</u>.
<p>How to block resources to be controlled manually?</p>	<p>To block a SAP resource, add the cluster resource name to the config variable <u>BLOCK_RESOURCES</u> in file <u>/etc/sysconfig/sap_suse_cluster_connector</u></p>
<p>Cluster resource naming requirements</p>	<p>Cluster resource name <i>must not</i> be prefixes of other cluster resource names.</p> <p>If you are using a name scheme like <code>rsc_sap<SID>_<INSTANCENAME></code>, you are on the safe side, because instance names including the instance number can also not be prefixes of other instance names.</p>
<p>Restrictions about migration rules or location rules.</p>	<p>While the master-slave resource must always be able to switch to another cluster node in case of a fail of the SAP instance, it is <i>strongly recommended not to</i> use migration rules for a longer time period and not to use location constraints with high scoring. Location constraints with low scoring such as "10" are valid. Migration constraints <i>must</i> be removed after the cluster is back in status <u>S_IDLE</u>.</p>

15 Cluster Tests

15.1 Mandatory Test Cases for SAP Enqueue Replication Server Integration

SUSE recommends to check an installed cluster for correct integration of SAP Enqueue Replication Server. The following test cases can be used to make sure that the most important mandatory parameters are set correctly. Before a case is tested, the cluster and its resources should be in a proper state. If no different situation is defined, all failures have to be repaired, all failcounts have to be cleared, and all migration constraints have to be removed.

TABLE 15.1: TEST CASES FOR SAP ENQUEUE REPLICATION SERVER INTEGRATION

#	Description	Action / Prerequisites / Result
01	Kill Message Server (Category Recoverable, no failover).	A: Identify the Message Servers PID: <code>pgrep -fl ms.sapHA0</code> then terminate the process: <code>kill -9 <pid-ms></code> P: Cluster in origin state (both nodes up and active, all instances up, Cluster in Status S_IDLE) and no migration rules left. R: Message server is restarted by sapstart. Killing multiple times will force a failover of the ASCS instance.
02	Kill Enqueue Server (Category Irrecoverable, failover required): shutdown Server, on which Enqueue Service is running.	A: Identify the Message Servers PID: <code>pgrep -fl en.sapHA0</code> then terminate the process: <code>kill -9 <pid-en></code> P: Cluster in origin state. R: Cluster will move instance ASCS to the other node.

#	Description	Action / Prerequisites / Result
03	Kill Enqueue Replication Server (Category Recoverable, no failover).	A: Identify the Message Servers PID: <u>pgrep -fl er.sapHA0</u> then terminate the process: <u>kill -9 <pid-er></u> . P: Cluster in origin state. R: enqueue replication server is restarted by sapstart.
04	Kill Enqueue sapstartsrv service (Category Recoverable, no failover) to test the integration of sapstartsrv into 3rd party failover solution and as provider of status information.	A: Identify the Message Servers PID: <u>pgrep -fl "ASCS00.*sapstartsrv"</u> then terminate the process: <u>kill -9 <pid-ssrv></u> . P: Cluster in origin state. R: Next monitor will restart sapstartsrv. Note: A SIGINT (2) Signal to the sapstartsrv service seems to be identified as a "Shutdown" by the Sapstartsrv.
05	Startup test: show proper loading of Shared Library.	A: <u>ClusterService RUP <resource></u> P: Cluster in origin state, but one instance stopped. R: Library is properly loaded and sap_suse_cluster_connector is called with command "init"; instance is started.
06	Runtime Test Shared Library	A: <u>sapcontrol -<nr> -function "Start Stop"</u> P: Cluster is in origin state. R: Instance addressed by the sapcontrol-call is started/stopped.
07	Apply kernel patch in DIR_CT_RUN and restart system (no failover required.)	A: SAP method to install a kernel update. P: Cluster is in origin state. R: Instances started after the SAP kernel update are running the new kernel code (<u>sapcpe</u> copies from <u>/sapmnt/<SID>/...</u>)
08	Simulate cluster failure with HA vendor software.	A: n/a P: n/a

#	Description	Action / Prerequisites / Result
		R: n/a
09	Call HA vendor function to move clustered SCS/ASCS instance to a different node.	A: <u>Clusterervice RMI <group-for-as></u> ; wait until cluster has moved the instance and is S_IDLE again; use <u>ClusterService RUM <group-for-as></u> to delete the client-preferred location rule. P: Cluster is in origin state. R: Cluster will move instance ASCS to the other node.
10	Test SCS/ASCS hardware or OS failure.	A: KVM-HOST: <u>virsh destroy <node></u> P: Cluster is in origin state. R: Remaining node does run the resources which where running on the failing node.
11	Temporarily disable HA vendor failover software.	A: For ALL groups and master-slaves: <u>ClusterService RSU <resource></u> ; then on both nodes: <u>rcopenais stop</u> (delay between the two nodes ~1 minute). P: Cluster is in origin state. R: Cluster Framework stopped, all SAP instances started before are still running.
12	New directory structure of Kernel 7.20 is used.	A: n/a P: n/a R: n/a
13	Message Server could not be restarted locally.	A: change mode of the message server binary or other action so message server binary file is not longer available for restart by sapstart. Kill the message server like described above. P: Cluster in origin state. R: Cluster needs to take over the ASCS00 instance, because it could not be recovered locally.

#	Description	Action / Prerequisites / Result
14	Enqueue Replication Server could not be restarted locally.	<p>A: Change mode of the enqueue replication server binary or other action so enqueue replication server binary file is no longer available for restart by sapstart.</p> <p>P: Cluster in origin state.</p> <p>R: Cluster needs to take over the ERS10 instance, because it could not be recovered locally.</p>

15.2 Recommended Tests for SLE High Availability Extension Basic Functionality

SUSE recommends to check an installed cluster for correct function of the SLE High Availability Extension basic infrastructure. This should be done before the integration test is conducted. The following test cases can be used in an environment similar to the certification. Since customers have different basic setups, the test cases may be adjusted. Before a case is tested, the cluster and its resources should be in a proper state. If no different situation is defined, all failures have to be repaired, all failcounts have to be cleared, and all migration constraints have to be removed.

TABLE 15.2: TEST CASES SLE HIGH AVAILABILITY EXTENSION FOR BASIC FUNCTIONALITY

#	Description	Action / Prerequisites / Expected Result
01	fence other node	<p>A: crm node fence \$node</p> <p>P: cluster up and running</p> <p>R: node \$node gets rebooted. Cluster migrates RGs to other node.</p>
02	watchdog functionality check	<p>A: killall -9 sbd</p> <p>P: cluster up and running</p> <p>R: node gets rebooted by watchdog. Cluster migrates RGs to other node.</p>

#	Description	Action / Prerequisites / Expected Result
11	Admin. Start of Cluster manager on both nodes	A: rcopenais start , on both nodes P: no cluster running on any node R: Both nodes join cluster. RGs start on their respective nodes.
12	Admin. Stop of cluster manager on both nodes	A: rcopenais stop, on both nodes P: cluster up and running R: Both nodes leave cluster. RGs stopped.
17	Failing of cluster manager on one node	A: killall -9 corosync P: cluster up and running R: Cluster manager stonithes the node via sbd. Cluster migrates RGs to other node.
19	Admin reboot of one cluster node	A: init 6 P: cluster up and running R: Node gets rebooted. Cluster migrates RGs to other node. After boot cluster manager has to be started manually.
41	Failing of one of two cluster interconnect links	A: Deactivate one port on LAN switch, unplug cable, iptables rule, NIC link. NOT: ifup/ifdown P: cluster up and running R: Resources are not affected.
42	Repair of cluster interconnect link.	A: Reactivate one port on LAN switch, plug cable, iptables rule, NIC link. P: Cluster up and running. One corosync ring faulty (as in 41). R: Corosync ring is recovered automatically. Resources are not affected.
46	Failure of one SBD link (out of two)	A: Deactivate one SBD LUN or set iptables rule for iSCSI LUN.

#	Description	Action / Prerequisites / Expected Result
		<p>P: cluster up and running.</p> <p>R: Resources are not affected. Cluster manager is not affected. SBD is not affected in general.</p>
57	Failing of all SBD links	<p>A: Deactivate all SBD LUNs or set iptables rule for iSCSI LUN.</p> <p>P: cluster up and running.</p> <p>R: Node loses SBD devices. Resources are not affected.</p>
59	Recovery of SBD links	<p>A: Reactivate all SBD LUNs or remove iptables rule for iSCSI LUN.</p> <p>P: cluster up and running.</p> <p>R: Node recovers all SBD devices automatically. Resources are not affected.</p>

III Appendix

- A Software Downloads 72
- B SUSE Products Online Documentation 73
- C SAP Notes 74
- D Links to SLES for SAP, SAP Software, SAP ASE 76
- E Licenses 78

A Software Downloads

Product	URL
SLES for SAP	https://www.suse.com/products/sles-for-sap/download/ ↗
SLES	https://www.suse.com/products/server/download/ ↗
SLES High Availability Extension	https://www.suse.com/products/highavailability/download/ ↗
SAP NetWeaver 7.5	http://service.sap.com/swdc ↗
SAP ASE	http://service.sap.com/swdc ↗

B SUSE Products Online Documentation

Documentation for SUSE products is available at <http://www.suse.com/documentation> in HTML and PDF formats.

Product	URL
SLES for SAP 12	https://www.suse.com/documentation/sles-for-sap-12/ ↗
SLES 12	https://www.suse.com/documentation/sles-12/ ↗
SLE High Availability Extension 12	https://www.suse.com/documentation/sle-ha-12/ ↗

C SAP Notes

The general installation of SAP on Linux is described in the SAP Note 171356 - SAP software on Linux: Essential information. This SAP note also points to some SAP notes with more detailed information about hardware platforms and Linux enterprise distributions. A good entry point for installing SAP on SUSE Linux Enterprise Server 12 is SAP Note 1984787. SAP Notes are available at the SAP Service Marketplace (<https://support.sap.com> (<https://support.sap.com/>) ). You need an account to access this information.

SAP Note	Title
1984787	SUSE LINUX Enterprise Server 12 Installation notes
1763512	Support details for SUSE Linux Enterprise High Availability
171356	SAP Software on Linux: General information
516716	Linux: Problems with locales after glibc update
1014480	XXX SAP Management Console (SAP MC)
784391	SAP support terms and 3rd-party Linux kernel drivers
1067221	Composite SAP Note for heterogeneous installation
1275776	Linux: Preparing SLES for SAP environments
941595	XXX Download J2SE 1.4.2 for the x64 platform
1172419	XXX Linux: Supported Java versions on the x86_64 platform
1240081	Java Cryptography Extension (JCE) Jurisdiction Policy Files
2206460	Release restrictions for SAP NetWeaver 7.5
2212573	SAP NetWeaver 7.5 Documentation
940420	FAQ: Database structure check (CHECK DATA/VERIFY)
785925	XXX SAP Web AS 6.40 SR1 ABAP Installation on UNIX

SAP Note	Title
1008828	XXX ACC 7.1 PI / Adaptive Computing Controller Collective Note
877795	XXX Problems with sapstartsrv from Release 7.00 and 6.40 patch 169
995116	XXX Backward porting of sapstartsrv for earlier releases
1122387	Linux: SAP Support in virtualized environments.
1554717	SYB: Planning information for SAP on ASE
1672366	XXX SYB: SAP NetWeaver 7.3 EHP1 on Sybase ASE (obsolete)
1585981	SYB: Ensuring Recoverability for Sybase ASE
1588316	SYB: Configure automatic database and log backups
1618817	SYB: How to restore an SAP ASE database server (UNIX)
1633491	SYB: Timestamp, BigDecimal: Problem during serialization
1716201	XXX Obsolete - Update control.xml files for installation: Sybase ASE
1606654	XXX SYB: SAP Business Suite 7i2010 on Sybase ASE (obsolete)
1590719	SYB: Updates for SAP Adaptive Server Enterprise (SAP ASE)
1599814	SYB: Installing Service Packs for SAP ASE (UNIX + Linux)
1602547	SYB: Versions of the syb_update_db script
1598817	SYB: Removal of sample client library applications
1763512	Linux: Support details for SUSE Linux Enterprise High Availability

D Links to SLES for SAP, SAP Software, SAP ASE

SLES for SAP:

<https://www.suse.com/products/server/> ↗

<https://www.suse.com/susePSC/home> ↗

<https://www.suse.com/yesssearch/Search.jsp> ↗

http://support.novell.com/products/server/supported_packages ↗

<https://www.suse.com/partners/alliance-partners/sap/> ↗

http://www.novell.com/docrep/2010/04/x2_Novell_SAP_Brochure_SinglePage%20Online.pdf ↗

http://www.novell.com/docrep/2010/07/SLES_for_SAP_Applications_FAQ.pdf ↗

<http://www.novell.com/docrep/2009/09/>

RT_WP_Linux_Trends_in_SAP_DC_200908231_Final_English_en.doc ↗

http://www.novell.com/docrep/2009/05/SUSE-Linux-Enterprise-11_Technical-Presentation_en_en.odp ↗

<http://www.novell.com/docrep/2009/11/>

Enterprise_Linux_Servers_Solution_Presentation_f_110409_en.pdf ↗

<http://www.linux-ha.org/> ↗

<http://clusterlabs.org/doc/> ↗

<http://www.clusterlabs.org/> ↗

<https://raid.wiki.kernel.org/index.php> ↗

<https://www.suse.com/support/kb/doc.php?id=7016880> ↗

<https://www.suse.com/support/kb/doc.php?id=7008216> ↗

<https://www.suse.com/support/kb/doc.php?id=7011346> ↗

<https://www.suse.com/support/kb/doc.php?id=7009485> ↗

<https://www.suse.com/support/kb/doc.php?id=7004817> ↗

<https://www.suse.com/support/kb/doc.php?id=D7007614> ↗

<https://www.suse.com/support/kb/doc.php?id=7016305> ↗

<https://www.suse.com/support/kb/doc.php?id=3155529> ↗

<https://www.suse.com/support/kb/doc.php?id=7015898> ↗

<https://www.suse.com/support/kb/doc.php?id=7016333> ↗

SLES for SAP:

<https://www.suse.com/support/kb/doc.php?id=7014247> ↗

<https://www.suse.com/support/kb/doc.php?id=7016707> ↗

<https://www.suse.com/support/kb/doc.php?id=7016761> ↗

SAP:

<https://www.sdn.sap.com/irj/sdn/nw-products> ↗

<http://scn.sap.com/docs/DOC-67382> ↗

<https://scn.sap.com/community/netweaver> ↗

<https://scn.sap.com/community/linux> ↗

<https://www.sdn.sap.com/irj/scn/weblogs?blog=/pub/wlg/13603> ↗

<https://support.sap.com/> ↗

<https://websmp201.sap-ag.de/pam> ↗

<https://sdn.sap.com/> ↗

<http://service.sap.com/osdbmigration> ↗

<http://www.sap.com/solutions/benchmark/index.epx> ↗

<https://scn.sap.com/docs/DOC-27321> ↗

<https://scn.sap.com/docs/DOC-31701> ↗

SAP ASE:

<http://infocenter.sybase.com/help/index.jsp> ↗

http://en.wikipedia.org/wiki/Adaptive_Server_Enterprise ↗

http://www.isug.com/Sybase_FAQ/ASE/index.html ↗

E Licenses

GNU Free Documentation License

Version 1.2, November 2002

Copyright (C) 2000,2001,2002 Free Software Foundation, Inc. 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

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.

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.

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.

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.

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.

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”.

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.

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.

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.

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.

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

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

```
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.

GNU General Public License

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc. 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a). You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b). You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c). If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

a). Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

b). Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c). Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program 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.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public 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.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
one line to give the program's name and an idea of what it does.  
Copyright (C) yyyy name of author
```

```
This program is free software; you can redistribute it and/or  
modify it under the terms of the GNU General Public License  
as published by the Free Software Foundation; either version 2  
of the License, or (at your option) any later version.
```

```
This program is distributed in the hope that it will be useful,  
but WITHOUT ANY WARRANTY; without even the implied warranty of  
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the  
GNU General Public License for more details.
```

```
You should have received a copy of the GNU General Public License  
along with this program; if not, write to the Free Software  
Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.
```

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) year name of author  
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details  
type `show w'. This is free software, and you are welcome  
to redistribute it under certain conditions; type `show c'  
for details.
```

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w' and `show c'; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright  
interest in the program `Gnomovision'  
(which makes passes at compilers) written  
by James Hacker.
```

signature of Ty Coon, 1 April 1989
Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the [GNU Lesser General Public License \(http://www.fsf.org/licenses/lgpl.html\)](http://www.fsf.org/licenses/lgpl.html) instead of this License.

Terminology

active/active, active/passive

A concept of how services are running on nodes. An active-passive scenario means that one or more services are running on the active node and the passive node waits for the active node to fail. Active-active means that each node is active and passive at the same time.

cluster

A high-performance cluster is a group of computers (real or virtual) sharing the application load in order to achieve faster results. A high-availability cluster is designed primarily to secure the highest possible availability of services.

cluster information base (CIB)

A representation of the whole cluster configuration and status (node membership, resources, constraints, etc.) written in XML and residing in memory. A master CIB is kept and maintained on the *designated coordinator (DC)* and replicated to the other nodes.

cluster partition

Whenever communication fails between one or more nodes and the rest of the cluster, a cluster partition occurs. The nodes of a cluster are split in partitions but are still alive. They can only communicate with nodes in the same partition and are unaware of the separated nodes. If the loss of the nodes on the other partition cannot be confirmed, a split brain scenario develops (see also *split brain*).

cluster resource manager (CRM)

The main management entity responsible for coordinating all non-local interactions. Each node of the cluster has its own CRM, but the one running on the DC is the one elected to relay decisions to the other non-local CRMs and process their input. A CRM interacts with a number of components: local resource managers, both on its own node and on the other nodes, non-local CRMs, administrative commands, the fencing functionality, and the membership layer.

consensus cluster membership (CCM)

The CCM determines which nodes make up the cluster and shares this information across the cluster. Any new addition and any loss of nodes or quorum is delivered by the CCM. A CCM module runs on each node of the cluster.

designated coordinator (DC)

The “master” node. This node is where the master copy of the CIB is kept. All other nodes get their configuration and resource allocation information from the current DC. The DC is elected from all nodes in the cluster after a membership change.

distributed lock manager (DLM)

DLM coordinates disk access for clustered file systems and administers file locking to increase performance and availability.

distributed replicated block device (DRBD)

DRBD is a block device designed for building high availability clusters. The whole block device is mirrored via a dedicated network and is seen as a network RAID-1.

failover

Occurs when a resource or node fails on one machine and the affected resources are started on another node.

fencing

Describes the concept of preventing access to a shared resource by isolated or failing cluster members. Should a cluster node fail, it will be shut down or reset to prevent it from causing trouble. This way, resources are locked out of a node whose status is uncertain.

Heartbeat resource agent

Heartbeat resource agents were widely used with Heartbeat version 1. Their use is deprecated, but still supported in version 2. A Heartbeat resource agent can perform `start`, `stop`, and `status` operations and resides under `/etc/ha.d/resource.d` or `/etc/init.d`. For more information about Heartbeat resource agents, refer to <http://www.linux-ha.org/HeartbeatResourceAgent> (see also *OCF resource agent*).

high availability

High availability is a system design approach and associated service implementation that ensures a prearranged level of operational performance will be met during a contractual measurement period.

Availability is a key aspect of service quality. Availability is usually calculated based on a model involving the Availability Ratio and techniques such as Fault Tree Analysis.

See also: http://en.wikipedia.org/wiki/High_availability/ http://www.itlibrary.org/index.php?page=Availability_Management

local resource manager (LRM)

The local resource manager (LRM) is responsible for performing operations on resources. It uses the resource agent scripts to carry out these operations. The LRM is “dumb” in that it does not know of any policy. It needs the DC to tell it what to do.

Linux Standard Base (LSB)

The goal of the Linux Standard Base is to develop and promote a set of standards that will increase compatibility among Linux distributions.

LSB resource agent

LSB resource agents are standard LSB init scripts. LSB init scripts are not limited to use in a high availability context. Any LSB-compliant Linux system uses LSB init scripts to control services. Any LSB resource agent supports the options start, stop, restart, status and force-reload and may optionally provide try-restart and reload as well. LSB resource agents are located in /etc/init.d. Find more information about LSB resource agents and the actual specification at <http://www.linux-ha.org/LSBResourceAgent> and http://www.linux-foundation.org/spec/refspecs/LSB_3.0.0/LSB-Core-generic/LSB-Core-generic/inisrptact.html (see also *OCF resource agent* and *Heartbeat resource agent*).

node

Any computer (real or virtual) that is a member of a cluster and invisible to the user.

policy engine (PE)

The policy engine computes the actions that need to be taken to implement policy changes in the CIB. This information is then passed on to the transaction engine, which in turn implements the policy changes in the cluster setup. The PE always runs on the DC.

OCF resource agent

OCF resource agents are similar to LSB resource agents (init scripts). Any OCF resource agent must support start, stop, and status (sometimes called monitor) options. Additionally, they support a metadata option that returns the description of the resource agent type in XML. Additional options may be supported, but are not mandatory. OCF resource agents reside in /usr/lib/ocf/resource.d/<provider>. Find more information about OCF resource agents and a draft of the specification at <http://www.linux-ha.org/OCFResourceAgent> and <http://www.opencf.org/cgi-bin/viewcvs.cgi/specs/ra/resource-agent-api.txt?rev=HEAD> (see also *Heartbeat resource agent*).

quorum

In a cluster, a cluster partition is defined to have quorum (is “quorate”) if it has the majority of nodes (or votes). Quorum distinguishes exactly one partition. It is part of the algorithm to prevent several disconnected partitions or nodes from proceeding and causing data and service corruption (split brain). Quorum is a prerequisite for fencing, which then ensures that quorum is indeed unique.

resource

Any type of service or application that is known to Heartbeat. Examples include an IP address, a file system, or a database.

resource agent (RA)

A resource agent (RA) is a script acting as a proxy to manage a resource. There are three different kinds of resource agents: OCF (Open Cluster Framework) resource agents, LSB resource agents (Standard LSB init scripts), and Heartbeat resource agents (Heartbeat v1 resources).

Single Point of Failure (SPOF)

A single point of failure (SPOF) is any component of a cluster that, should it fail, triggers the failure of the entire cluster.

split brain

A scenario in which the cluster nodes are divided into two or more groups that do not know of each other (either through a software or hardware failure). STONITH prevents a split brain situation from badly affecting the entire cluster. Also known as a “partitioned cluster” scenario.

The term split brain is also used in DRBD but means that the two nodes contain different data.

STONITH

The acronym for “Shoot the other node in the head”, which refers to the fencing mechanism that shuts down a misbehaving node to prevent it from causing trouble in a cluster.

STONITH Block Device (SBD)

Implementation of STONITH, using a shared SAN LUN for issuing shutdown instruction to certain nodes. This implementation always needs an a watchdog running on each node to be save.

transition engine (TE)

The transition engine (TE) receives policy directives from the PE and carries them out. The TE always runs on the DC. From there, it instructs the local resource managers on the other nodes which actions to take.

watchdog

A watchdog timer, or simply a watchdog, is an electronic timer that is used to detect and recover from computer malfunctions. During normal operation, the cluster regularly restarts the watchdog timer to prevent it from elapsing, or "timing out". If, due to a hardware fault or program error, the cluster fails to restart the watchdog, the timer will shutdown the computer.