Guide

SUSE Linux Enterprise Server for SAP
Applications 12 SP5

Guide

SUSE Linux Enterprise Server for SAP Applications 12 SP5

Publication Date: August 25, 2020

SUSE LLC
1800 South Novell Place
Provo, UT 84606
USA
https://documentation.suse.com ↗

Contents

About This Guide

SUSE® Linux Enterprise Server for SAP Applications is the reference platform for the software development of SAP. It is optimized for SAP applications. This guide provides detailed information about installing and customizing SUSE Linux Enterprise Server for SAP Applications.

SUSE Linux Enterprise High Availability Extension is also part of SUSE Linux Enterprise Server for SAP Applications.

1 Overview

The SUSE Linux Enterprise Server for SAP Applications Guide is divided into the following chapters:

What Is SUSE Linux Enterprise Server for SAP Applications?

An overview of SUSE Linux Enterprise Server for SAP Applications.

Planning the Installation

Information on hardware requirements, the installation workflow, partitioning, and other installation planning aspects.

Installing the Operating System

Installing the SUSE Linux Enterprise Server operating system that forms the basis of SUSE Linux Enterprise Server for SAP Applications.

Installing SAP Applications

Installing SAP applications on SUSE Linux Enterprise Server for SAP Applications, either directly after the installation of the operating system or in a running system.

Setting Up an Installation Server for SAP Media Sets

Setting up a server for all Installation Media used within your organization.

Setting Up an SAP HANA Cluster

Setting up an SAP HANA cluster with system replication using the YaST wizard.

Tuning

Tuning the system to enable the best performance for SAP applications.

Firewalling

Keeping the system and applications safe using firewalls.

Encrypting Directories Using `cryptctl`

Keeping server data encrypted at rest.

Protecting Against Malware With ClamSAP

Keeping your users safe from malware.

Connecting via RDP

Using remote access to administrate machines.

Creating Operating System Images

Creating KIWI images and scrubbing private data off fully configured images.

Important Log Files

Gives a short overview of important log files.

2 Available Documentation and Resources

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

- For the latest updates to this guide, see https://documentation.suse.com/sles-sap .

- This is an abridged version of the *SUSE Linux Enterprise Server for SAP Applications Guide*. For the latest version of the full *SUSE Linux Enterprise Server for SAP Applications Guide*, see https://documentation.suse.com/sles-sap .

- You can find numerous whitepapers, a best-practices guide, and other resources at the SUSE Linux Enterprise Server for SAP Applications resource library: https://www.suse.com/products/sles-for-sap/resource-library/ .

- For documentation on SUSE Linux Enterprise Server, see https://documentation.suse.com/sles-12 .

- For documentation on SUSE Linux Enterprise High Availability Extension, see https://documentation.suse.com/sle-ha-12 .

3 Giving Feedback

We welcome feedback for and contributions to this documentation. There are several channels for this:

Service Requests and Support

For services and support options available for your product, see http://www.suse.com/support/ ↗.

To open a service request, you need a SUSE subscription registered at SUSE Customer Center. Go to https://scc.suse.com/support/requests ↗, log in, and click *Create New*.

Bug Reports

Report issues with the documentation at https://bugzilla.suse.com/ ↗. Reporting issues requires a Bugzilla account.

To simplify this process, you can use the *Report Documentation Bug* links next to headlines in the HTML version of this document. These preselect the right product and category in Bugzilla and add a link to the current section. You can start typing your bug report right away.

Contributions

To contribute to this documentation, use the *Edit Source* links next to headlines in the HTML version of this document. They take you to the source code on GitHub, where you can open a pull request. Contributing requires a GitHub account.

Mail

You can also report errors and send feedback concerning the documentation to `doc-team@suse.com`. Include the document title, the product version, and the publication date of the document. Additionally, include the relevant section number and title (or provide the URL) and provide a concise description of the problem.

4 Documentation Conventions

The following notices and typographic conventions are used in this document:

- `/etc/passwd`: Directory names and file names

- *`PLACEHOLDER`*: Replace *`PLACEHOLDER`* with the actual value

- `PATH`: An environment variable

- `ls`, `--help`: Commands, options, and parameters

- `user`: The name of user or group

- `package_name`: The name of a software package

- `Alt`, `Alt`–`F1`: A key to press or a key combination. Keys are shown in uppercase as on a keyboard.

- *File, File* › *Save As*: menu items, buttons

- `IBM Z, POWER` ⟩ This paragraph is only relevant for the architectures `IBM Z` and `POWER`. The arrows mark the beginning and the end of the text block. ◁

- *Chapter 1, "Example Chapter"*: A cross-reference to another chapter in this guide.

- Commands that must be run with `root` privileges. Often you can also prefix these commands with the **`sudo`** command to run them as non-privileged user.

  ```
  root # command
  tux > sudo command
  ```

- Commands that can be run by non-privileged users.

  ```
  tux > command
  ```

- Notices

  ### Warning: Warning Notice
  Vital information you must be aware of before proceeding. Warns you about security issues, potential loss of data, damage to hardware, or physical hazards.

  ### Important: Important Notice
  Important information you should be aware of before proceeding.

  ### Note: Note Notice
  Additional information, for example about differences in software versions.

💡 Tip: Tip Notice

Helpful information, like a guideline or a piece of practical advice.

- Compact Notices

  🌀 Additional information, for example about differences in software versions.

  💡 Helpful information, like a guideline or a piece of practical advice.

1 What Is SUSE Linux Enterprise Server for SAP Applications?

SUSE® Linux Enterprise Server for SAP Applications is a bundle of software and services that addresses the specific needs of SAP users. It is the only operating system that is optimized for all SAP software solutions.

Target use cases include:

- Unix to Linux migrations and replatforming

- SAP appliances

- SAP cloud deployments

SUSE Linux Enterprise Server for SAP Applications consists of software components and service offerings which are described in the following sections. The figure *Offerings of SUSE Linux Enterprise Server for SAP Applications* shows which software components and services are also available with other products from SUSE (green) and which ones are exclusively available with SUSE Linux Enterprise Server for SAP Applications (blue).
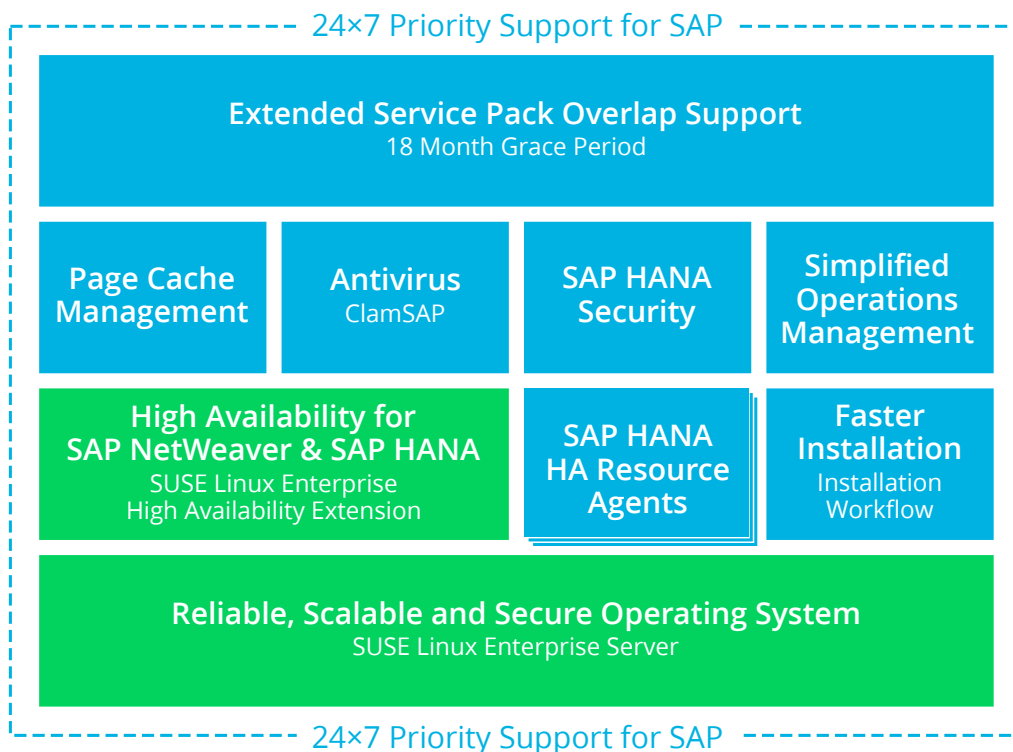
FIGURE 1.1: OFFERINGS OF SUSE LINUX ENTERPRISE SERVER FOR SAP APPLICATIONS

1.1   Software Components

As depicted in *Figure 1.1, "Offerings of SUSE Linux Enterprise Server for SAP Applications"*, SUSE Linux Enterprise Server for SAP Applications is based on SUSE Linux Enterprise Server. Additionally, it contains software components such as SUSE Linux Enterprise High Availability Extension, the kernel page-cache limit feature, and the installation workflow. These software components are briefly explained in the following sections.

1.1.1   SUSE Linux Enterprise Server

The current release is based on SUSE Linux Enterprise Server 12 SP5. SUSE Linux Enterprise Server is the most interoperable platform for mission-critical computing, both physical and virtual.

1.1.2   SUSE Linux Enterprise High Availability Extension

This component consists of:

- Flexible, policy-driven clustering

- Cluster-aware file system and volume management

- Continuous data replication

- Setup and installation

- Management and administration tools

- Resource agents, also for SAP

- Virtualization-aware

SUSE Linux Enterprise High Availability Extension provides two resource agents specifically for working with SAP applications:

- `SAPInstance` which allows starting and stopping instances of SAP products.

- `SAPDatabase` which allows starting and stopping all databases supported by SAP applications (SAP HANA, SAP MaxDB, SAP ASE, Oracle, Sybase, IBM DB2).

For more information about SUSE Linux Enterprise High Availability Extension, see the *Administration Guide,* available from https://documentation.suse.com/sle-ha-12 ↗. Additional White Papers and Best Practice Guides are available in the SUSE Linux Enterprise Server for SAP Applications Resource Library (https://www.suse.com/products/sles-for-sap/resource-library/ ↗).

1.1.3   Simplified SAP HANA System Replication Setup

SUSE Linux Enterprise Server for SAP Applications supports SAP HANA System Replication by using components of SUSE Linux Enterprise High Availability Extension and two additional resource agents (RA). Additionally, SUSE Linux Enterprise Server for SAP Applications ships with a YaST wizard that simplifies the cluster setup.

1.1.3.1   SAPHana Resource Agent

This resource agent from SUSE supports scale-up scenarios by checking the SAP HANA database instances for whether a takeover needs to happen. Unlike with the pure SAP solution, takeovers can be automated.

It is configured as a master/slave resource: The master assumes responsibility for the SAP HANA databases running in primary mode. The slave is responsible for instances that are operated in synchronous (secondary) status. In case of a takeover, the secondary (slave resource instance) can automatically be promoted to become the new primary (master resource instance).

This resource agent supports system replication for the following in scale-up scenarios:

- **Performance-Optimized Scenario.**   Two servers (A and B) in the same SUSE Linux Enterprise High Availability Extension cluster, one primary (A) and one secondary (B). The SAP HANA instance from the primary server (A) is replicated synchronously to the secondary server (B).

- **Cost-Optimized Scenario.**   The basic setup of A and B is the same as in the *Performance-Optimized Scenario*. However, the secondary server (B) is also used for non-productive purposes, such as for an additional SAP HANA database for development or QA. The production database is only kept on permanent memory, such as a hard disk. If a takeover needs

to occur, the non-productive server will be stopped before the takeover is processed. The system resources for the productive database are then increased as quickly as possible via an SAP hook call-out script.

- **Chain/Multi-Tier Scenario.** Three servers (A, B, and C), of which two are located in the same SUSE Linux Enterprise High Availability Extension cluster (A and B). The third server (C) is located externally. The SAP HANA system on the primary server (A) is replicated synchronously to the secondary server (B). The secondary server (B) is replicated asynchronously to the external server (C).

  If a takeover from A to B occurs, the connection between B and C remains untouched. However, B is not allowed to be the source for two servers (A and C). This would be a "star" topology which is not supported with current SAP HANA versions (such as SPS11). Using SAP HANA commands, you can then manually decide what to do:

  - The connection between B and C can be broken, so that B can connect to A.

  - If replication to the external site (C) is more important than local system replication, the connection between B and C can be kept.

For all of the scenarios, SUSE Linux Enterprise Server for SAP Applications supports both single-tenant and multi-tenant (MDC) SAP HANA databases. That is, you can use SAP HANA databases that serve multiple SAP applications.

1.1.3.2 `SAPHanaTopology` Resource Agent

To make configuring the cluster as simple as possible, SUSE has developed the `SAPHanaTopology` resource agent. This agent runs on all nodes of a SUSE Linux Enterprise High Availability Extension cluster. It gathers information about the status and configurations of SAP HANA system replications. It is designed as a normal (stateless) clone.

1.1.3.3 YaST Wizard to Set Up SAP HANA Clusters

SUSE Linux Enterprise Server for SAP Applications now additionally ships a YaST wizard that manages the initial setup of such clusters according to best practices. The wizard is part of the package `yast2-sap-ha` and can be started using YaST, via *HA Setup for SAP Products*.

For more information, see *Chapter 6, Setting Up an SAP HANA Cluster*.

1.1.3.4 For More Information

For more information, see:

- The *Administration Guide* at https://documentation.suse.com/sles-12 ↗.

- The *Best Practices* in the Resource Library at https://www.suse.com/products/sles-for-sap/resource-library/ ↗. In particular, see *Setting up a SAP HANA SR performance optimized infrastructure* and *Setting up a SAP HANA SR cost optimized infrastructure*.

1.1.4 Installation Workflow

The installation workflow offers a guided installation path for both the SUSE Linux Enterprise Server operating system and the SAP application. For more information, see *Section 2.5, "Overview of the Installation Workflow"*.

Additionally, the installation workflow can be extended by third-party vendors or customers using Supplementary Media. For more information about creating Supplementary Media, see *Appendix C, Supplementary Media*.

1.1.5 Page-Cache Limit

You can limit the kernel file system cache size to influence swapping behavior. With this feature, you can often gain better performance by allocating memory to an application.

For more information, see *Section 7.1, "Kernel: Page-Cache Limit"*.

1.1.6 Malware Protection with ClamSAP

ClamSAP integrates the ClamAV anti-malware toolkit into SAP NetWeaver and SAP Mobile Platform applications to enable cross-platform threat detection. For example, you can use ClamSAP to allow an SAP application to scan for malicious uploads in HTTP uploads.

For more information, see *Chapter 10, Protecting Against Malware With ClamSAP*.

1.1.7 SAP HANA Security

SUSE Linux Enterprise Server for SAP Applications contains additional features to help set up well-secured SAP HANA installations.

1.1.7.1 Firewall for SAP HANA

Securing SAP HANA can require many additional firewall rules. To simplify firewall setups for SAP HANA, SUSE Linux Enterprise Server for SAP Applications contains the package `HANA-Firewall`. It provides preconfigured rules and integrates with SuSEfirewall2.

For more information, see *Section 8.2, "Configuring HANA-Firewall"*.

1.1.7.2 Hardening Guide for SAP HANA

For information on hardening the underlying operating system, see the SUSE Linux Enterprise Server for SAP Applications resource library: https://www.suse.com/products/sles-for-sap/resource-library/ 🡵 . There, find the document *OS Security Hardening for SAP HANA*.

1.1.8 Simplified Operations Management

SUSE Linux Enterprise Server for SAP Applications combines several features that enable simplified operations management.

1.1.8.1 System Tuning with `sapconf`

The system tuning application `sapconf` allows you to automatically and comprehensively tune your system as recommended by SAP (for use with SAP S/4HANA, SAP NetWeaver, or SAP HANA/SAP BusinessOne). To do this, `sapconf` activates `tuned` profiles. These allow tuning several kernel parameters, depending on the hardware components you are using, such as the amount of available RAM.

For more information, see *Section 7.2, "Tuning Systems with `sapconf4`"*.

1.1.8.2 Storage Encryption for Hosted Applications with `cryptctl`

Today, databases and similar applications are often hosted on external servers that are serviced by third-party staff. Certain data center maintenance tasks require third-party staff to directly access affected systems. In such cases, privacy requirements necessitate disk encryption.

`cryptctl` allows encrypting sensitive directories using LUKS and offers the following additional features:

- Encryption keys are located on a central server which can be located on customer premises.

- Encrypted partitions are automatically remounted after an unplanned reboot.

For more information, see *Chapter 9, Encrypting Directories Using* `cryptctl`.

1.1.8.3   Patterns Providing Dependencies of SAP Applications

To simplify working with software dependencies of SAP applications, SUSE has created patterns that combine relevant dependency RPM packages for specific applications:

- *SAP BusinessOne Server Base*

- *SAP HANA Server Base*

- *SAP NetWeaver Server Base*

❗ Important: Packages May Be Missing from Patterns

The selection of packages of the software patterns are defined while a specific release (Service Pack or major version) of SUSE Linux Enterprise Server for SAP Applications is developed. This package selection is stable over the lifetime of this particular release. When working with SAP applications that have been released more recently than your SUSE Linux Enterprise Server for SAP Applications version, dependencies can be missing from the patterns.

For definitive information about the dependencies of your SAP application, see the documentation provided to you by SAP.

1.1.8.4   `ClusterTools2`

`ClusterTools2` provides tools that help set up and manage a Corosync/`pacemaker` cluster. Among them are the command `wow` which helps create highly available system resources, and `ClusterService` which allows managing a cluster.

Additionally, `ClusterTools2` provides scripts that automate common cluster tasks:

- Scripts that perform checks. For example, to find out whether a system is set up correctly for creating a `pacemaker` cluster.

- Scripts that simplify configuration. For example, to create a Corosync configuration.

- Scripts that monitor the system and scripts that show or collect system information. For example, to find known error patterns in log files.

For more information, see the man page of the respective tool, included with the package `ClusterTools2`. Also see the project home page at https://github.com/fmherschel/ClusterTools2 ↗.

1.2   Software Repository Setup

Software included with operating systems based on SUSE Linux Enterprise is delivered as RPM packages, a form of installation package that can have dependencies on other packages. On a server or an installation medium, these packages are stored in software repositories (sometimes also called "channels").

By default, computers running SUSE Linux Enterprise Server for SAP Applications are set up to receive packages from multiple repositories. Of each of the standard repositories, there is a "Pool" variant that represents the state of the software when it was first shipped. There is also an "Update" variant that includes the newest maintenance updates for the software in the "Pool" variant.

If you registered your system during installation, your repository setup should include the following:

TABLE 1.1: STANDARD REPOSITORIES

Content	Base repository ("Pool")	Update repository
Base packages of SUSE Linux Enterprise Server	`SLES12-SPSP5-Pool`	`SLES12-SPSP5-Updates`
Packages specific to SUSE Linux Enterprise Server for SAP Applications	`SLE12-SPSP5-SAP-Pool`	`SLE-12-SPSP5-SAP-Up-dates`

Content	Base repository (“Pool”)	Update repository
Packages specific to SUSE Linux Enterprise High Availability Extension	`SLE-HA12-SPSP5-Pool`	`SLE-HA12-SPSP5-Updates`

The tables in this section do not show *Debuginfo* and *Source* repositories which are also set up but disabled by default. The *Debuginfo* repositories contain packages that can be used for debugging regular packages. The *Source* repositories contain source code for packages.

Depending on your installation method, you may also see `SLE-12-SPSP5-SAP-12.SP5-0` which is the installation medium. It contains packages from all of the base software repositories listed above.

Because there are own repositories for SUSE Linux Enterprise Server for SAP Applications, SUSE can ship packages and patches that are specific to SUSE Linux Enterprise Server for SAP Applications.

> ### Note: ESPOS Updates Shipped Directly in Update Repositories
> Unlike in SUSE Linux Enterprise Server for SAP Applications 11, updates related to Extended Service Pack Overlay Support (ESPOS) are shipped directly from the `Update` repositories. This means there is no separate ESPOS repository to set up.

In addition to the standard repositories, you can enable the following SLE Modules and SLE Extensions (either during the installation or from the running system by using YaST or the command `SUSEConnect`).

TABLE 1.2: **MODULE REPOSITORIES**

Content	Base repository (“Pool”)	Update repository
Advanced Systems Management Module: CFEngine, Puppet and the Machinery tool	`SLE-Module-Adv-Systems-Management12-Pool`	`SLE-Module-Adv-Systems-Management12-Updates`
Containers Module: Docker, tools, prepackaged images	`SLE-Module-Containers12-Pool`	`SLE-Module-Containers12-Updates`

Content	Base repository ("Pool")	Update repository
HPC Module: tools and libraries related to High Performance Computing	`SLE-Module-HPC12-Pool`	`SLE-Module-HPC12-Updates`
Legacy Module: Sendmail, old IMAP stack, old Java, …	`SLE-Module-Legacy12-Pool`	`SLE-Module-Legacy12-Updates`
Public Cloud Module: public cloud initialization code and tools	`SLE-Module-Public-Cloud12-Pool`	`SLE-Module-Public-Cloud12-Updates`
Toolchain Module: GNU Compiler Collection (GCC)	`SLE-Module-Toolchain12-Pool`	`SLE-Module-Toolchain12-Updates`
Web and Scripting Module: PHP, Python, Ruby on Rails	`SLE-Module-Web-Scripting12-Pool`	`SLE-Module-Web-Scripting12-Updates`

TABLE 1.3: **EXTENSION REPOSITORIES**

Content	Base repository ("Pool")	Update repository
SLE SDK	`SLE-SDK12-SPSP5-Pool`	`SLE-SDK12-SPSP5-Updates`
SUSE Package Hub	`SUSE-Package-Hub-12-SPSP5-Pool`	`SUSE-Package-Hub-12-SPSP5-Updates`

For more information about PackageHub, see *Section A.3, "PackageHub"*. For information about life cycle and support of modules and extensions, see *Section 1.3, "Included Services"*.

1.3 Included Services

Extended Service Pack Overlap Support (ESPOS)

> Subscriptions for SUSE Linux Enterprise Server for SAP Applications include Extended Service Pack Overlap Support. It extends the overlap between the support periods of two consecutive service packs by three years. During this period, you receive support and all relevant maintenance updates under the terms of Long Term Service Pack Support (LTSS). This feature is only available with SUSE Linux Enterprise Server for SAP Applications.

Extended Service Pack Overlap Support allows you to perform service pack migrations within three and a half (3½) years instead of only six months. This enables you to schedule migrations more easily and perform testing before a migration under less restrictive time constraints.

At an additional cost, SUSE also offers Long Term Service Pack Support (LTSS). With LTSS, you receive support for a particular service pack after the ESPOS period ends. SUSE Linux Enterprise Server for SAP Applications includes one and a half (1½) years of general support and three years of ESPOS for each service pack. Therefore, LTSS may only be offered for the last service pack. For more information about the life cycle of SUSE products, see:

- General explanation: *Deployment Guide, Part "Updating and Upgrading SUSE Linux Enterprise", Chapter "Life Cycle and Support"* (https://documentation.suse.com/sles-12 ↗)

- Exact dates: https://www.suse.com/lifecycle/ ↗

SUSE Linux Enterprise Server Priority Support for SAP Applications

Subscriptions for SUSE Linux Enterprise Server for SAP Applications include SUSE Linux Enterprise Server Priority Support for SAP Applications. It offers technical support for SUSE Linux Enterprise Server for SAP Applications directly from SAP. The joint support infrastructure is provided by support engineers from SUSE Technical Support and SAP. It is based upon SAP Resolve and offers seamless communication with both SAP and SUSE. This "One Face to the Customer" support model reduces complexity and lowers the total cost of ownership.

For more information, see *SAP Note 1056161: SUSE Priority Support for SAP Applications* (https://launchpad.support.sap.com/#/notes/1056161 ↗).

❗ Important: Life Cycle and Support for Modules and Extensions

Modules and extensions have a different life cycle than SLES-SAP and SUSE provides different support offerings for them:

- Modules:

    - **Life Cycle.** Varies depending on the module.

    - **Support.** Only up-to-date packages are supported. Support is included with your subscription for SUSE Linux Enterprise Server for SAP Applications. You do not need an additional registration key.

- Extensions

- **Life Cycle.** Releases are usually coordinated with SUSE Linux Enterprise Server for SAP Applications.

- **Support.** Support is available but not included with your subscription for SUSE Linux Enterprise Server for SAP Applications. You need an additional registration key.

- Unsupported Extensions (PackageHub and SUSE Linux Enterprise Software Development Kit)

  - **Life Cycle.** Releases are usually coordinated with SUSE Linux Enterprise Server for SAP Applications.

  - **Support.** There is no support beyond fixes for security and packaging issues. You do not need an additional registration key.

2 Planning the Installation

Read this chapter carefully, as it helps you plan the installation: It lists requirements and helps you collect data about your system.

2.1 Hardware Requirements

This section lists minimum hardware requirements for the installation of SUSE Linux Enterprise Server for SAP Applications. It gives basic guidance on the expected hardware requirements of certain SAP software. For the most up-to-date information about the hardware requirements of SAP software, see the official sizing guidelines at https://service.sap.com/sizing↗.

CPU

> Intel 64/AMD64
>
> IBM POWER servers with PowerKVM compatibility

Hard Disk

> SUSE Linux Enterprise Server for SAP Applications requires at least 45 GB of hard disk space for the system volume. In addition to that, reserve an appropriate amount of hard disk space for the swap partition.
>
> To install an SAP application such as SAP NetWeaver, you need at least 200 GB of free disk space (in addition to the required space for the operating system for the application's `/data` partition).
>
> To install SAP HANA, you need either:

- An SAP BusinessOne-certified machine

- A compatible machine that meets the requirements for SAP HANA TDI (Tailored Datacenter Integration). That is, you need the following amounts of free disk space in addition to the required space for the operating system:

  - 52 GB of free disk space for the partition `/usr/sap`

  - Space for three partitions for SAP HANA data: `/hana/data` (same size as RAM), `/hana/log` (same size as RAM up to a maximum of 512 GB), and `/hana/shared` (same size as RAM up to a maximum of 1 TB).

For more information about SAP HANA TDI storage requirements, see:
https://www.sap.com/docs/
download/2015/03/74cdb554-5a7c-0010-82c7-eda71af511fa.pdf ↗

RAM

The SUSE Linux Enterprise Server operating system itself requires a minimum of 1024 MB
of total RAM or a minimum of 512 MB of RAM per CPU core (choose whichever is higher).
Any SAP software you install will require additional RAM.

To install SAP HANA, your machine needs a minimum of 24 GB of RAM.

For more information about configuring hardware for SAP HANA, see *SAP Note 1944415: Hard-
ware Configuration Guide and Software Installation Guide for SUSE Linux Enterprise Server with SAP
HANA and SAP Business One* (https://launchpad.support.sap.com/#/notes/1944415 ↗).

For more information about partitioning, see *Section 2.7, "Partitioning"*.

2.2 Downloading the Installation Image

1. Download the ISO image of SUSE® Linux Enterprise Server for SAP Applications 12 SP5
   DVD 1 (electronic media kit) from https://www.suse.com/products/sles-for-sap/ ↗.

2. Burn the image onto a physical DVD and ensure that it is bootable. Alternatively, use a
   virtual DVD-ROM device for installation in a virtual machine.

2.3 Offline Migration

The migration paths for SUSE Linux Enterprise Server are identical to those for SUSE Linux
Enterprise Server for SAP Applications. Find detailed information in the *Upgrade Guide* at https://
documentation.suse.com/sles/html/SLES-all/cha-upgrade-paths.html ↗.

2.4 Installation Methods

There are multiple ways of installing SUSE Linux Enterprise Server for SAP Applications:

- *Using the Installation Workflow* (standard way of installation)

- *Using an External AutoYaST Profile*

2.5 Overview of the Installation Workflow

The installation workflow of SUSE Linux Enterprise Server for SAP Applications consists of the following steps:

1. Installation of the operating system (SUSE Linux Enterprise Server). See *Section 3.1, "Using the Installation Workflow"*.

2. SAP Installation Wizard, part 1: Copying all required SAP media to the local disk or selecting a shared storage medium to use. See *Section 4.3, "Using the SAP Installation Wizard"*, in particular *Step 1*.

3. SAP Installation Wizard, part 2: Collecting all parameters for the actual installation by querying the user interactively. See *Section 4.3, "Using the SAP Installation Wizard"*, in particular *Step 10*.

4. SAP Installation Wizard, part 3: Running the SAP Installer. See *Section 4.3, "Using the SAP Installation Wizard"*, in particular *Step 13*.

> ⓘ **Important: Installation of Only SUSE Linux Enterprise Server**
>
> You can choose to only install a base SUSE Linux Enterprise Server system. In that case, only the first step of the installation workflow is executed.
>
> This can be necessary when you want to install an Oracle database on a SUSE Linux Enterprise Server for SAP Applications machine. To do so, install the base product SUSE Linux Enterprise Server first, then install the Oracle database and later convert your installation to SLES-SAP. This is necessary because the installer of Oracle databases queries for the existence of certain files, not all of which are included in a SLES-SAP installation.
>
> For more information about converting, see *Section 3.4, "Converting a SLES Installation to a SLES-SAP Installation"*.

Most of these steps do not need to be run immediately after each other, which allows for flexibility in how you install systems. This means that you can prepare a single installation as a first step and then continue from there. For example:

- Install the operating system (SUSE Linux Enterprise Server) only.

or

- Install the operating system (SUSE Linux Enterprise Server), copy SAP media, and collect SAP installation parameters.

Then, create disk images, copy them to other systems, and adjust SAP installation parameters. Finally, finish the installation on each machine individually.

2.6 Required Data for Installing

Operating System

The SUSE Linux Enterprise Server installation requires the following data for every physical server:

- Network configuration parameters, such as host name, domain, IP address, subnet mask, domain searchlist (DNS), IP for name server, IP for gateway

- Administrator (`root`) password for the SUSE Linux Enterprise Server installation

SAP Application

The installation of an SAP application generally requires specifying:

- SAP SID

- SAP Instance Number

- A password for the SAP application

Depending on the SAP application you are installing, more parameters may be necessary, such as T-Shirt Sizing or parameters for virtual networking.

SAP HANA Database

The installation of SAP HANA requires specifying:

- SAP SID

- SAP Instance Number

- Whether to enable Multitenant Database Containers (MDC). The multi-tenant support of SAP HANA allows having multiple databases that run as one SAP HANA installation. (To use SAP HANA MDC, you need SAP HANA Life Cycle Manager.)
  For a single-tenant installation, choose *No*.

For a multi-tenant instance administrated by one `SIDadm` user, choose *Yes with low isolation*.

For a multi-tenant instance administrated in which each database has its own `SIDadm` user, choose *Yes with high isolation*.

- A password for the SAP HANA database

For more information about installing SAP software, see the SAP documentation at https://help.s-ap.com ↗ and https://support.sap.com ↗.

2.7 Partitioning

SUSE Linux Enterprise Server for SAP Applications creates the partitioning table in two stages:

1. *Partitioning for the Operating System (stage 1)* (during the installation of the operating system)

2. *Partitioning for the SAP System (stage 2)* (during the installation of the SAP product)

2.7.1 Partitioning for the Operating System (stage 1)

During the installation of the operating system, partitions for the operating system are created.

A logical volume group (LVG) named `/dev/system` will be created. This LVG contains two logical volumes (LVs):

- `/dev/system/root`: by default 60 GB to account for the operating system and SAP media

- `/dev/system/swap`: by default 2 GB, avoid setting a smaller size. See also *SAP Note 1984787: SUSE Linux Enterprise Server 12: Installation notes* (https://launchpad.support.s-ap.com/#/notes/1984787 ↗).

Additionally, a `boot` or UEFI partition will be created as necessary.

2.7.2 Partitioning for the SAP System (stage 2)

The partitioning for the SAP system can be created by:

- The SAP Installation Wizard (see *Section 4.3, "Using the SAP Installation Wizard"*).

- Using YaST on the command line (see *Section 4.5, "Partitioning for an SAP Application Without the SAP Installation Wizard"*).

This part of the partitioning can only be created after the operating system has been installed. That means the partitions are created either in the installation workflow after the reboot or in the running system.

Depending on the product you are installing and your particular use case, the amount of hard disk space necessary can vary.

For information on partitioning for the SAP system using AutoYaST, see *Appendix B, Partitioning for the SAP System Using AutoYaST*.

3 Installing the Operating System

The following section provides instructions for installing the base operating system. Using the installation workflow, you can install either using a local installation medium or over the network. Alternatively, you can install using AutoYaST.

3.1 Using the Installation Workflow

The installation workflow is a guided installation of the operating system with optimized settings for SAP applications. During the installation workflow, you can choose whether you want to install an SAP application. If so, you will be asked to provide SAP installation media when the SUSE Linux Enterprise Server installation is finished. You can also choose whether to install third-party extensions.

This section assumes that you are starting the installation from a local medium. To learn how to start the installation from a remote medium, see *Section 3.2, "Using SLES-SAP Media from the Network"*.

For more information, see *Section 2.5, "Overview of the Installation Workflow"*.

This section guides you through the installation of the SUSE Linux Enterprise Server for SAP Applications operating system.

PROCEDURE 3.1: STARTING THE OS INSTALLATION

1. • On Intel 64/AMD64, boot from the DVD. From the DVD boot menu, select *Installation*.

   • On POWER, follow the instructions in the SUSE Linux Enterprise Server documentation, see *Deployment Guide, Part "Installation Preparation", Chapter "Installation on IBM POWER"* (https://documentation.suse.com/sles-12 ↗).
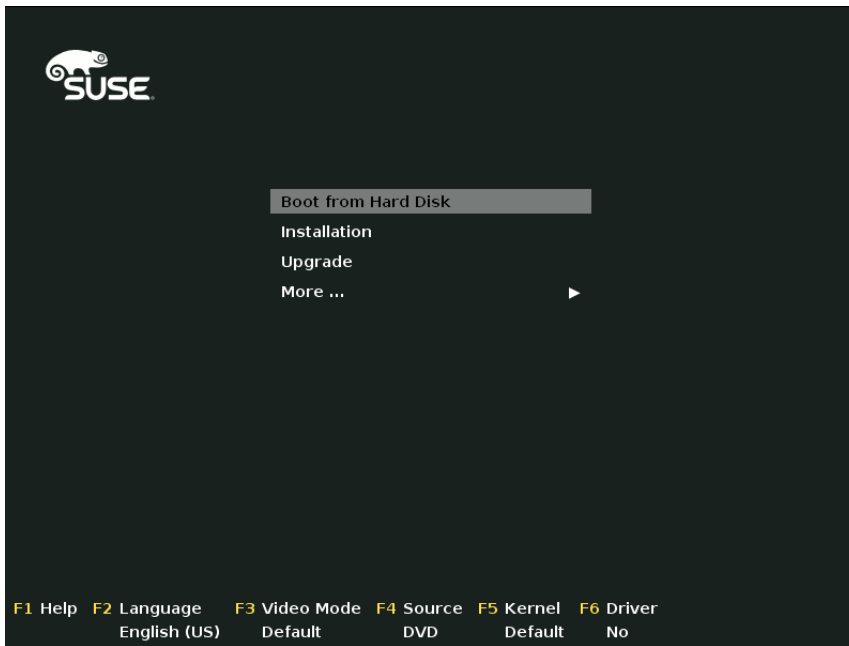
FIGURE 3.1: DVD BOOT MENU

While the initial operating system is starting, you can view boot messages by pressing `Esc` . When this process has completed, the graphical installation workflow will start.

2. Select the default system language under *Language*.
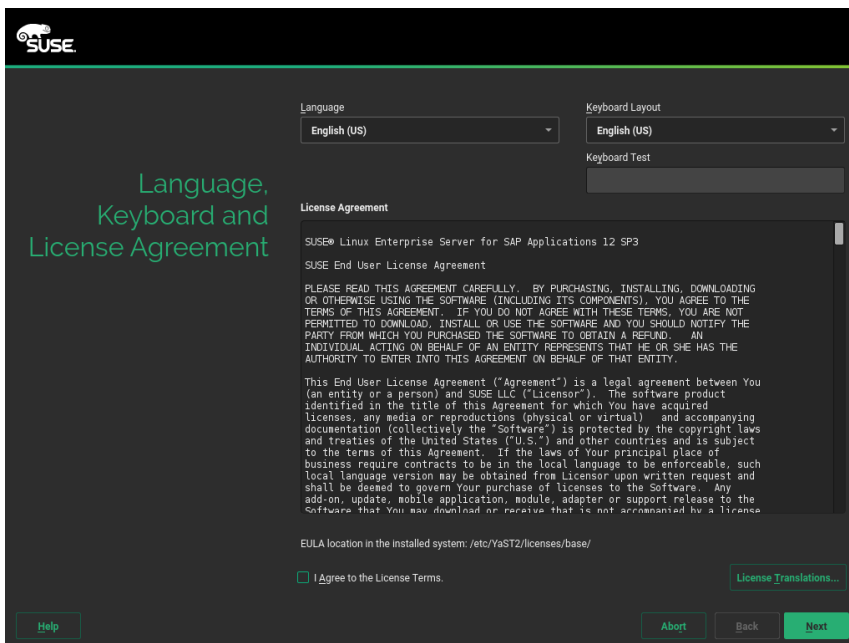
FIGURE 3.2: LANGUAGE, KEYBOARD AND LICENSE AGREEMENT

3. Select the appropriate keyboard layout under *Keyboard Layout*. To test whether the selected layout matches your physical keyboard, use the text box *Keyboard Test*.

4. Read the license agreement. If you agree, select *I Agree to the License Terms*. Proceed with *Next*.
   Otherwise, cancel the installation with *Abort › Abort Installation.*

5. *(Optional)* If automatic network configuration via DHCP fails, the screen *Network Settings* will open.
   If instead the screen *Registration* appears, your network connection works. To change network settings anyway, click *Network Configuration.*
   When you are finished configuring networking, proceed with *Next*.

   ! Important: Configure Networking as Recommended by SAP
   Make sure to configure the network connection as recommended in the documentation provided to you by SAP.

   For information about configuring networking, see *Administration Guide, Chapter "Basic Networking", Section "Configuring a Network Connection with YaST"* (https://documentation.suse.com/sles-12 ↗).

6. On the screen *Registration*, enter your *E-mail Address* and *Registration Code*. Successful registration is a prerequisite for receiving product updates and the entitlement to technical support.
   Proceed with *Next.*

   ! Important: Register at This Step
   Make sure to register your system at this step in the installation. Otherwise, you will not receive package updates immediately.
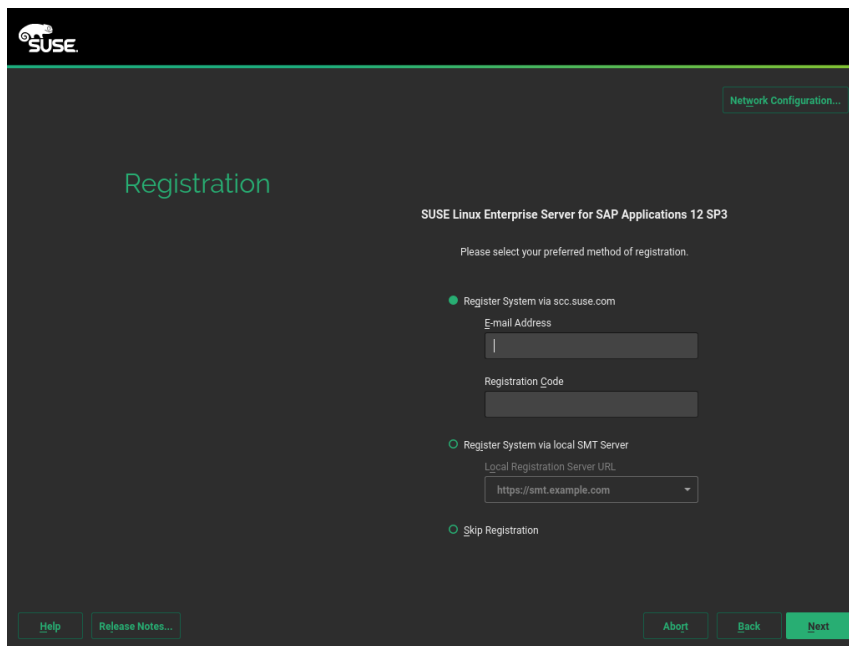
FIGURE 3.3: **REGISTRATION**

7. When asked whether to enable update repositories, choose *Yes*.

8. After the system is successfully registered, YaST lists additional software that is available for your product from the SUSE Customer Center. The list contains modules, which are free, and extensions, which require a registration key that is liable for costs. To enable a module or an extension, activate its entry.
   Proceed with *Next*.

9. The following screen allows you to choose the *Product Installation Mode*. You can now choose between:

   - **A SUSE Linux Enterprise Server Installation.** To install a SLES system without SAP-specific customization, choose *Proceed with standard SLES installation*. For details, see *Installation Quick Start, Section "Installing SUSE Linux Enterprise Server"* (https://documentation.suse.com/sles-12 ↗).

     > **❗ Important: Installing Oracle Databases**
     >
     > To be able to install an Oracle database later, choose *Proceed with standard SLES installation* and later convert your installation to SLES-SAP.

This is necessary because the installer for Oracle databases queries for the existence of certain files, not all of which are included in a SLES-SAP installation.

For more information about converting, see *Section 3.4, "Converting a SLES Installation to a SLES-SAP Installation"*.

- **A SUSE Linux Enterprise Server for SAP Applications Installation.** To install a SLES system with SAP-specific customization, choose *Proceed with standard SLES for SAP Applications installation.*

  - To install an SAP Application together with the system, activate *Launch the SAP Installation Wizard right after the operating system is installed.*

  - To enable RDP access (Remote Desktop Protocol) to this machine, activate *Enable RDP service and open port in firewall.*
    For more information about connecting via RDP, see *Chapter 11, Connecting via RDP*.

  Proceed with *Next*.
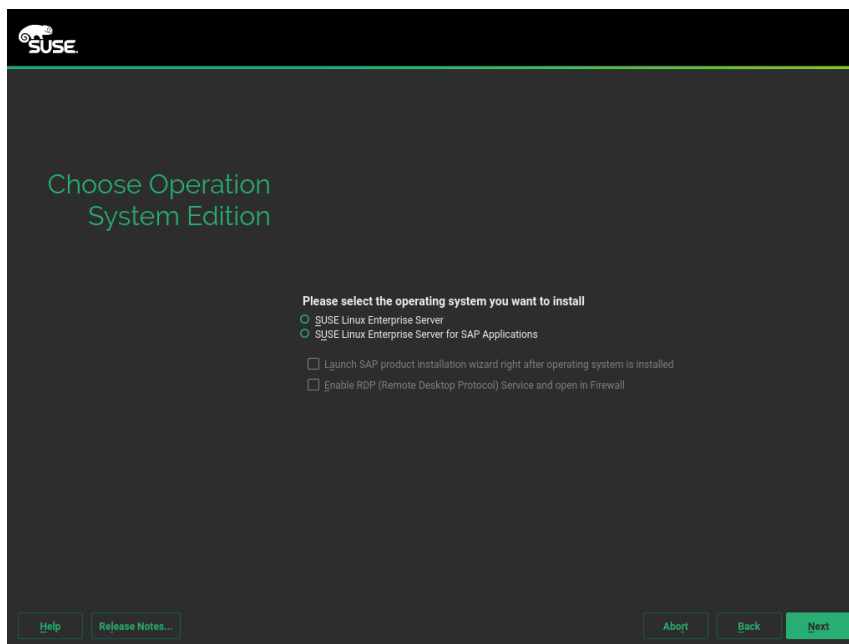
FIGURE 3.4: **INSTALLATION TYPE**

PROCEDURE 3.2: **FINISHING THE OS INSTALLATION**

1. You can now choose whether to install an *Add On Product*. Proceed with *Next*.

2. Review the proposed partition setup for the volumes `/dev/system/root` and `/dev/sys-` `tem/swap`. The volume `/dev/system/data` will be created later, as described in *Section 2.7, "Partitioning"*.

Suitable values are preselected. However, if necessary, change the partition layout. You have the following options:

Edit Proposal Settings

> Allows you to change the options for the proposed settings, but not the suggested partition layout itself.

Create Partition Setup

> Select a disk to which to apply the proposal.

Expert Partitioner

> Open the *Expert Partitioner* described in *Deployment Guide, Chapter "Advanced Disk Setup", Section "Using the YaST Partitioner"* (https://documentation.suse.com/sles-12 ↗). For partitioning advice specific to SUSE Linux Enterprise Server for SAP Applications, see *Section 2.7, "Partitioning"*.

To accept the proposed setup without changes, proceed with *Next*.
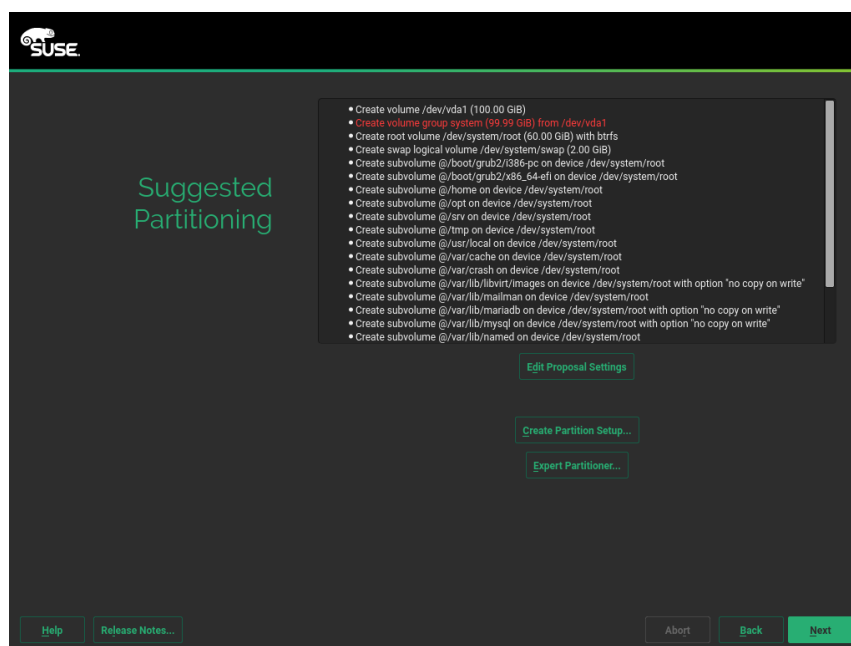
FIGURE 3.5: **SUGGESTED PARTITIONING**

Note: Release Notes

From this point on, the Release Notes can be viewed from any screen during the installation process by selecting *Release Notes*.

3. Select the clock and time zone to use on your system. To manually adjust the time or to configure an NTP server for time synchronization, choose *Other Settings*. For detailed information, see *Deployment Guide, Chapter "Installation with YaST", Section "Clock and Time Zone"* (https://documentation.suse.com/sles-12 ↗).
   Proceed with *Next*.

4. Type a password for the system administrator account (called `root`) and repeat the password under *Confirm Password*. You can use the text box *Test Keyboard Layout* to make sure that all special characters appear correctly.
   For more information, see *Deployment Guide, Chapter "Installation with YaST", Section "Password for the System Administrator root"* (https://documentation.suse.com/sles-12 ↗).
   Proceed with *Next*.

Important: Do Not Forget the `root` Password

The user `root` has the permission to carry out all administrative tasks. Without this password, you cannot log in to the system as `root`. The password entered here cannot be retrieved later.

5. On the screen *Installation Settings*, you can review and, if necessary, change several proposed installation settings. Each setting is shown alongside its current configuration. To change parts of the configuration, click the appropriate headline or other underlined items.

Important: Firewall Configuration

The software firewall of SLES-SAP is enabled by default. However, often, the ports your SAP product requires to be open are not opened automatically. This means that there may be network issues until you open the required ports manually.

For details, see *Section 8.1, "Configuring SuSEfirewall2"*.
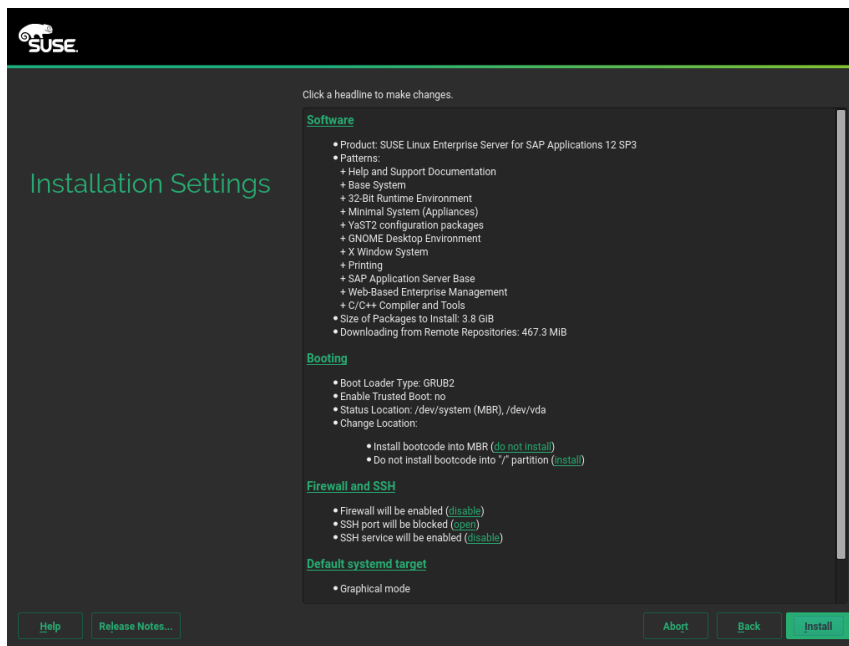
FIGURE 3.6: INSTALLATION SETTINGS

6. When you are satisfied with the system configuration, click *Install*.

   Depending on your software selection, you may need to agree to further license agreements before you are asked to confirm that you want to start the installation process.

   ## ✋ Warning: Deletion of Data

   The installation process fully or partially overwrites existing data on the disk.

   In the installation confirmation box, click *Install*.

   When the installation of the operating system is finished, the system will reboot automatically:

   - If you chose to only prepare the system for installation, the system will boot to a desktop login screen.

   - If you chose to install an SAP application immediately after the operating system, the installation will continue after the reboot.

     In this case, continue with *Chapter 4, Installing SAP Applications*.

3.2   Using SLES-SAP Media from the Network

This section provides a short description of how to install from an installation medium served over the network. This allows, for example, using a regular SLES medium to install SLES-SAP.

1. Copy the content of the SUSE Linux Enterprise Server for SAP Applications DVD to a Web server (for example, `example.com`), to the directory `/srv/www/htdocs/sap_repo`.

2. Boot from a SUSE DVD.

3. Select one of the boot menu options using the keys ⎵↓⎵/⎵↑⎵. Then add to the command line. To do so, specify the parameters listed below:

    *   To allow network usage, add `ifcfg=*=dhcp` (though this should be the default).

    *   Add the parameter `install=SERVER/DIRECTORY`.

4. Follow the instructions in *Section 3.1, "Using the Installation Workflow"*.

For more information, see *Deployment Guide, Chapter "Remote Installation"* (https://documentation.suse.com/sles-12 ↗).

To avoid having to use a DVD to initialize the system, you can boot over the network via PXE. For details, see *AutoYaST Guide, Chapter "Booting via PXE over the Network"* (https://documentation.suse.com/sles-12 ↗).

3.3   Using an External AutoYaST Profile

For more information about installing with AutoYaST, see:

*   *Deployment Guide, Part "Automated Installations", Chapter "Automated Installation"* (https://documentation.suse.com/sles-12 ↗).

*   *AutoYaST Guide* (https://documentation.suse.com/sles-12 ↗).

For more information about partitioning for SAP applications with AutoYaST, see *Section 2.7, "Partitioning"*.

If you plan to deploy SUSE Linux Enterprise Server for SAP Applications from a SUSE Manager server, refer to *SUSE Manager "Reference Manual", "Systems", "Autoinstallation"* and *SUSE Manager "Advanced Topics", Chapter "Minimalist AutoYaST Profile for Automated Installations and Useful Enhancements"* (https://documentation.suse.com/suma ↗).

3.4 Converting a SLES Installation to a SLES-SAP Installation

To convert an installation of SUSE Linux Enterprise Server 12 SP5 or JeOS 12 SP5 to an installation of SLES-SAP, use the script **Migrate_SLES_to_SLES-for-SAP.sh**. The script will register the system correctly and subscribe it to the appropriate repositories.

Make sure that you have an e-mail address for registration and a registration code for SLES-SAP.

1. Install the package `migrate-sles-to-sles4sap`.

2. Execute the following command:

   ```
   root # Migrate_SLES_to_SLES-for-SAP-12.sh
   ```

3. When asked to confirm to continue the migration, press `Y`, then `Enter`.

4. When asked, type the e-mail address to use for registration, then press `Enter`.

5. When asked, type the registration key, then press `Enter`.
   Wait until the script is finished. Afterward, you are subscribed to the SUSE Linux Enterprise Server for SAP Applications software repositories and the package `SLES-release` is removed in favor of `SLES_SAP-release`.

> **Important: Script Does Not Install Default SLES-SAP Packages**
>
> The script does not install all packages that are included with a default SLES-SAP installation. However, you can install these yourself manually. To install the default package selection, use:
>
> ```
> root # zypper in patterns-sles-sap_server
> ```

4 Installing SAP Applications

This section guides you through the installation of SAP media sets you received from SAP.

- If you are installing an SAP application within the installation workflow, continue with *Section 4.2, "First Steps"*.

- If you are installing an SAP application within an installed system, continue with *Section 4.3, "Using the SAP Installation Wizard"*.

4.1 Products That Can Be Installed Using SAP Installation Wizard

Using the SAP Installation Wizard, you can install stand-alone SAP HANA database instances. Additionally, the following SAP products (along with a database) can be installed using the SAP Installation Wizard:

- SAP S/4HANA, on-premise edition 1511

- SAP NetWeaver 7.5

- SAP NetWeaver 7.4 Support Release 2

- SAP NetWeaver 7.4 Support Release 1

- SAP NetWeaver 7.4

- SAP Enhancement Package 1 for SAP NetWeaver 7.3

- SAP NetWeaver 7.3

- SAP NetWeaver Composition Environment (CE) 7.2

- SAP EHP1 for SAP NetWeaver Composition Environment (CE) 7.1

- SAP NetWeaver Composition Environment (CE) 7.1

- SAP EHP1 for SAP NetWeaver Mobile/Banking 7.1

- SAP EHP1 SAP NetWeaver Process Integration 7.1

- SAP EHP1 for SAP NetWeaver Adaptive Computing Controller 7.1

- SAP NetWeaver Mobile/Banking 7.1

- SAP NetWeaver Process Integration 7.1

- SAP NetWeaver Adaptive Computing Controller 7.1

- SAP Business Suite powered by SAP HANA

- SAP Business Suite 7i 2016

- SAP Business Suite 7i 2013 Support Release 2

- SAP Business Suite 7i 2013 Support Release 1

- SAP Business Suite 7i 2011 Java

- SAP Business Suite 7i 2010 Java

- SAP Business Suite 7 Support Release 1 Java

- SAP Solution Manager 7.2 Support Release 1

- SAP Solution Manager 7.1 powered by SAP HANA

- SAP NetWeaver AS ABAP 7.4, OEM version 1.0

❗ Important: Installation of Oracle Databases Not Possible

The SAP Installation Wizard does not allow installing products together with Oracle databases. To install an Oracle database, install the base product SUSE Linux Enterprise Server first, then install the Oracle database and later convert your installation to SLES-SAP. This is necessary because the installer of Oracle databases queries for the existence of certain files, not all of which are included in a SLES-SAP installation.

For more information about converting, see *Section 3.4, "Converting a SLES Installation to a SLES-SAP Installation"*.

4.2 First Steps

These first steps are only relevant during the installation workflow.

1. When the system is booted, it displays the screen *Welcome*. Proceed with *Next*.

2. The screen *Network Settings* will now open. This gives you an opportunity to change the network settings.
   When you are finished configuring networking, proceed with *Next*.

> **！ Important: Configure Networking as Recommended by SAP**
>
> Make sure to configure the network connection according to the documentation of your SAP application.

For information about configuring networking, see *Administration Guide, Chapter "Basic Networking", Section "Configuring a Network Connection with YaST"* (https://documentation.suse.com/sles-12 ↗).

(While the next screen loads, the *Welcome* screen may appear again for a few seconds.)

3. Choose one of the following options:

   Create SAP file systems and start SAP product installation

   Allows installing an SAP application and setting up the system as a server providing SAP installation routines to other systems.
   Continue with *Section 4.3, "Using the SAP Installation Wizard"*.

   Only create SAP HANA file systems, do not install SAP products now

   Create an SAP HANA file system on SAP BusinessOne-certified hardware.

   > **！ Important: Hardware Requirements**
   >
   > Make sure your machine fulfills the hardware requirements for SAP HANA detailed in *Section 2.1, "Hardware Requirements"*. Otherwise, this option will not create a new file system and the installation workflow ends at this point.

   Finish wizard and proceed to OS login

   Do not install an SAP application and continue to the login screen of SUSE Linux Enterprise Server for SAP Applications.

   Proceed with *Next*.

4.3 Using the SAP Installation Wizard

Use the SAP Installation Wizard to install an SAP NetWeaver system (including database) or a simple SAP HANA system (single tenant, same password for all initial users, default settings).

To install other SAP applications or to create a more advanced SAP HANA setup, do not use this wizard. Instead, directly use one of the installation methods provided by SAP.

Tip: Installing an SAP Application in a Fully Installed System

This process is documented as it appears during the installation workflow. However, it also applies to the YaST module *SAP Installation Wizard* which is available in the installed system.

To start the SAP Installer, from the desktop, choose *Applications* › *System* › *YaST*, continue in the YaST control center by choosing *Miscellaneous* › *SAP Installation Wizard*.

Tip: SAP Installation Wizard Configuration

The SAP Installation Wizard configuration is specified and documented in `/etc/sysconfig/sap-installation-wizard`. You can change it according to your needs.

1. In the screen *SAP Installation Wizard*, provide the *Location of the SAP Installation Master* (*Figure 4.1, "Location of SAP Installation Master"*). The location can either be a local, removable, or remote installation source.
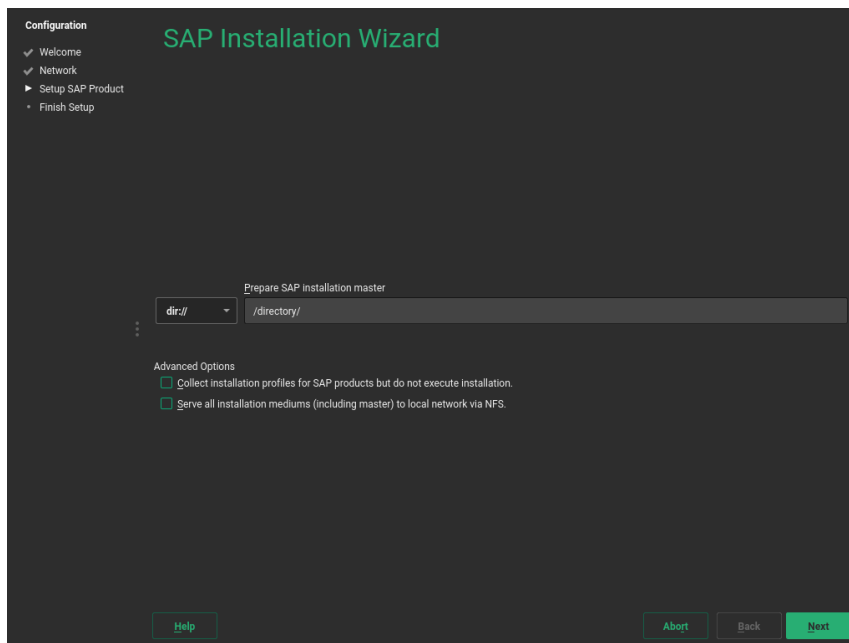
FIGURE 4.1: **LOCATION OF SAP INSTALLATION MASTER**

Select the appropriate option from the drop-down box. In the text box, specify the path to your source according to the format given in the following table.

TABLE 4.1: MEDIA SOURCE PATH

Option	Description	Format of Path
Local Sources		
dir://	a local directory	`/path/to/dir/`
Removable Sources		
device://	a locally connected hard disk	`devicename/path/to/dir/on/device`
usb://	a USB mass storage device	`/path/to/dir/on/USB`
cdrom://	a CD or DVD	`//`
Remote Sources		
nfs://	an NFS share	`server_name/path/to/dir/on/device`
smb://	an SMB share	`[user_name:password@]server_name//path/to/dir/on/server[?workgroup=workgroup_name]`

💡 Tip: Remote Location Specification

To install from an NFS source, specify the name of the server and the complete path to the media data. For information about setting up a remote installation server, see *Chapter 5, Setting Up an Installation Server for SAP Media Sets*.

If you have installed an SAP application from an installation server before (or set up your system to be an installation server), you can also directly choose that server as the provider of the Installation Master. To do so, use the drop-down box below *e an installation master*.

2. Under *Advanced Options,* choose from the following options:

Collect installation profiles for SAP products but do not execute installation

Use this option to set the installation parameters, but not perform the actual installation. With this option, the SAP Installer (SAPinst) will stop without performing the actual SAP product installation. However, the steps that follow fully apply.

For more information, see *Section 4.4, "Continuing an Installation Using an Installation Profile".*

Serve all installation media (including master) to local network via NFS

Set up this system as an installation server for other SUSE Linux Enterprise Server for SAP Applications systems. The media copied to this installation server will be offered through NFS and can be discovered via Service Location Protocol (SLP).

Proceed with *Next.*

The SAP Installation Wizard will now copy the Installation Master to your local disk. Depending on the type of Installation Master you selected, the installation will continue differently:

- If you are installing an SAP HANA database, skip ahead to *Step 8*.

- If you are installing an SAP NetWeaver application, continue with the next step.

3. On the screen *SAP Installation Wizard,* provide the location of additional Installation Media you want to install. This can include an SAP kernel, a database, and database exports.

Copy a medium

Specify a path to additional Installation Media. For more information about specifying the path, see *Table 4.1, "Media Source Path".*

Skip copying of medium

Do not copy additional Installation Media. Choose this option in the following cases: If you do not need additional Installation Media or if you want to install additional Installation Media directly from their source, for example CDs/DVDs or flash disks. When choosing this option despite your SAP product requiring additional Installation Media, you later need to provide the SAP Installer (SAPinst) with the relevant paths.

Proceed with *Next.*

If you chose to copy Installation Media, the SAP Installation Wizard will copy the relevant files to your local hard disk.
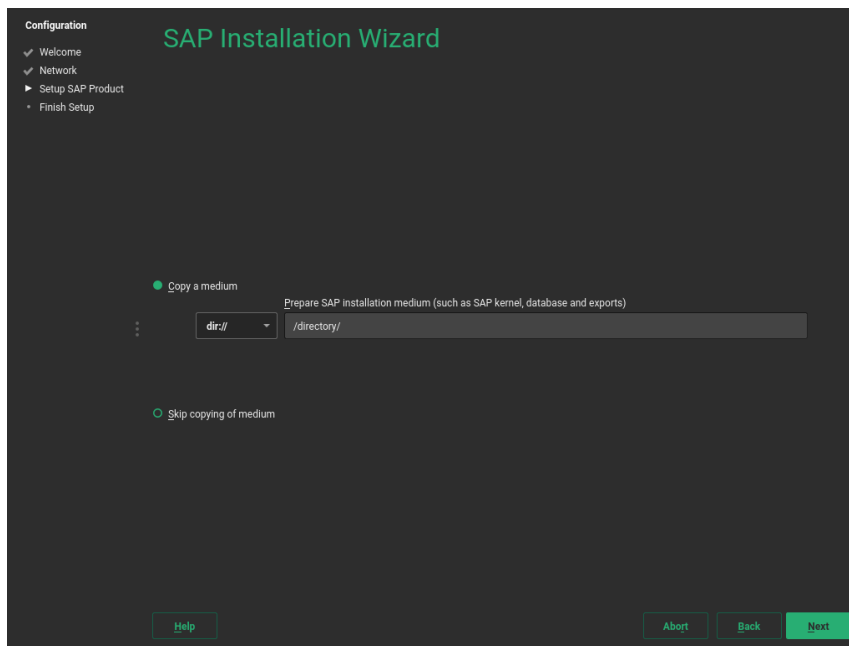
FIGURE 4.2: SAP INSTALLATION WIZARD: ADDITIONAL INSTALLATION MEDIA

4. After copying the Installation Media, you will be asked whether you want to prepare additional Installation Media. To do so, click *Yes*. Then follow the instructions in *Step 3*. Otherwise, click *No*.

5. In the screen *What Would You Like to Install*, under *The SAP product is*, choose how you want to install the product:

SAP Standard System

   Install an SAP application including its database.

SAP Standalone Engines

   Engines that add functionality to a standard product: SAP TREX, SAP Gateway, and Web Dispatcher.

Distributed System

   An SAP application that is separated onto multiple servers.

SAP High-Availability System

   Installation of SAP NetWeaver in a high-availability setup.

System Rename

Allows changing system properties such as the SAP system ID, database ID, instance number or host name. This can be used to install the same product in a very similar configuration on different systems.
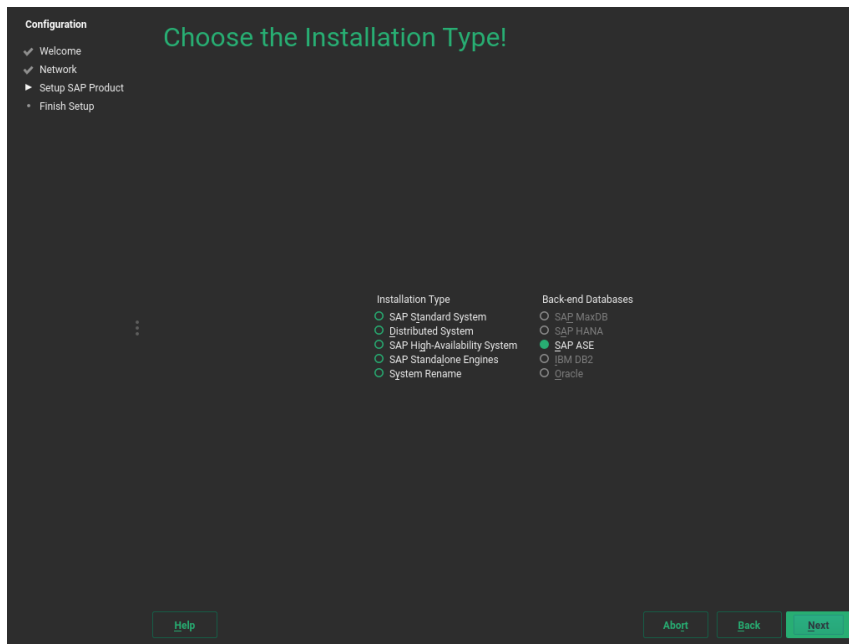
FIGURE 4.3: SAP INSTALLATION WIZARD: INSTALLATION TYPE AND DATABASE

6. If you selected *SAP Standard System, Distributed System,* or *SAP High-Availability System,* additionally choose a back-end database under *The back-end database system is.*
Proceed with *Next.*

7. You will now see the screen *Choose a Product.* The products shown depend on the Media Set and Installation Master you received from SAP. From the list, select the product you want to install.
Proceed with *Next.*
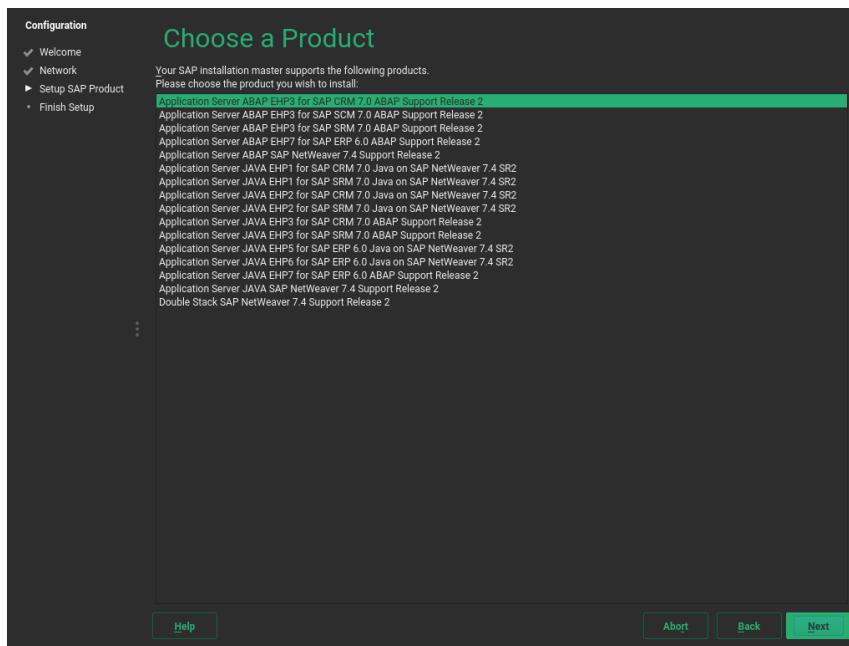
FIGURE 4.4: **SAP INSTALLATION WIZARD: CHOOSE A PRODUCT**

8. You will be asked whether to copy Supplementary Media or Third-Party Media. To do so, click *Yes* and then follow the instructions in *Step 3*.

    Otherwise, click *No*.

> ✏️ **Note: Difference Between Supplementary Media/Third-Party Media and Additional Software Repositories**
>
> Both types of delivery mechanisms allow installing software that is neither part of the SUSE Linux Enterprise Server for SAP Applications media nor part of your Media Set from SAP. However, the delivery mechanism is different:
>
> - Supplementary Media/Third-Party Media is installed using an AutoYaST file which allows creating an installation wizard and custom installation scripts.
>
> - Additional software repositories are RPM package repositories that you will remain subscribed to. This means you receive updates for Third-Party Media along with your regular system updates.

For information on creating Supplementary Media, see *Appendix C, Supplementary Media*.

9. On the screen *Additional software repositories for your SAP installation*, you can add further software repositories. For example, for add-ons that are packaged as RPM. To do so, click *Add new software repositories*. For more information about adding repositories, see *Deployment Guide, Chapter "Installing and Removing Software", Section "Adding Software Repositories"* (https://documentation.suse.com/sles-12 ↗).

Proceed with *Next*.

> ### ⊘ Note: Location of Copied SAP Media
>
> At this point, all data required for the SAP installation has been copied to `/data/SAP_CDs` (unless you chose to skip the process of copying). Each Installation Medium is copied to a separate directory. You might find the following directory structure, for example:
>
> ```
> tux > ls /data/SAP_CDs
> 742-KERNEL-SAP-Kernel-742
> 742-UKERNEL-SAP-Unicode-Kernel-742
> RDBMS-MAX-DB-LINUX_X86_64
> SAP-NetWeaver-740-SR2-Installation-Export-CD-1-3
> SAP-NetWeaver-740-SR2-Installation-Export-CD-2-3
> SAP-NetWeaver-740-SR2-Installation-Export-CD-3-3
> ```
>
> `/data/SAP_CDs` is the default directory as specified in the `/etc/sysconfig/sap-installation-wizard` configuration file.

10. Depending on the product you are installing, one or more dialogs will prompt you to supply values for configuration parameters for the SAP application you are installing.

    Supply the values as described in the documentation provided to you by SAP. Help for the configuration parameters is also available on the left side of the dialog. For more information, see *Section 2.6, "Required Data for Installing"*.

    Fill out the form (or forms), then proceed with *OK*.
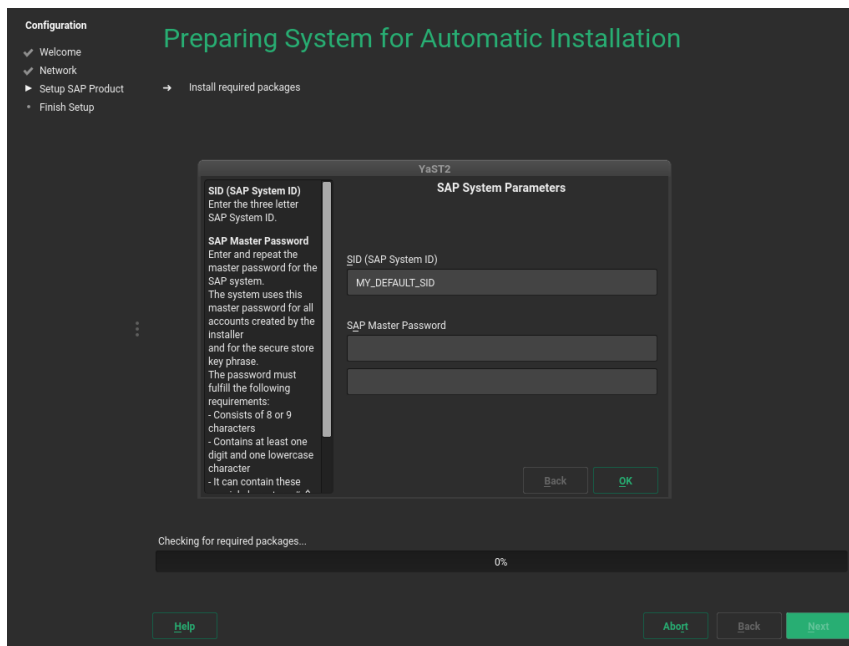
FIGURE 4.5: PRODUCT PARAMETERS

When you are done, the SAP Installation Wizard will download additional software packages.

11. You will be asked whether to continue the installation or prepare another SAP product for installation. If you choose to prepare another SAP product, start from the beginning of this procedure.

12. *(Optional)* When installing SAP HANA on a system that is not certified for SAP HANA and does not meet the minimum hardware requirements for SAP HANA TDI (Tailored Datacenter Integration), you will be asked whether to continue. If you receive this message unexpectedly, check *Section 2.1, "Hardware Requirements"* and the sizing guidelines from SAP at https://service.sap.com/sizing ↗.

    Otherwise, continue with *Yes*.

13. The following steps differ depending on the type of SAP application you are installing:

    - When installing an SAP HANA database, SAP HANA will now be installed without further question.

    - When installing an SAP NetWeaver application, the actual installation will be performed using the SAP Installer (SAPinst). After a few seconds, SAP Installer will open automatically.

Follow the SAP Installer as described in the documentation provided by SAP. Most configuration parameters are correctly filled already.
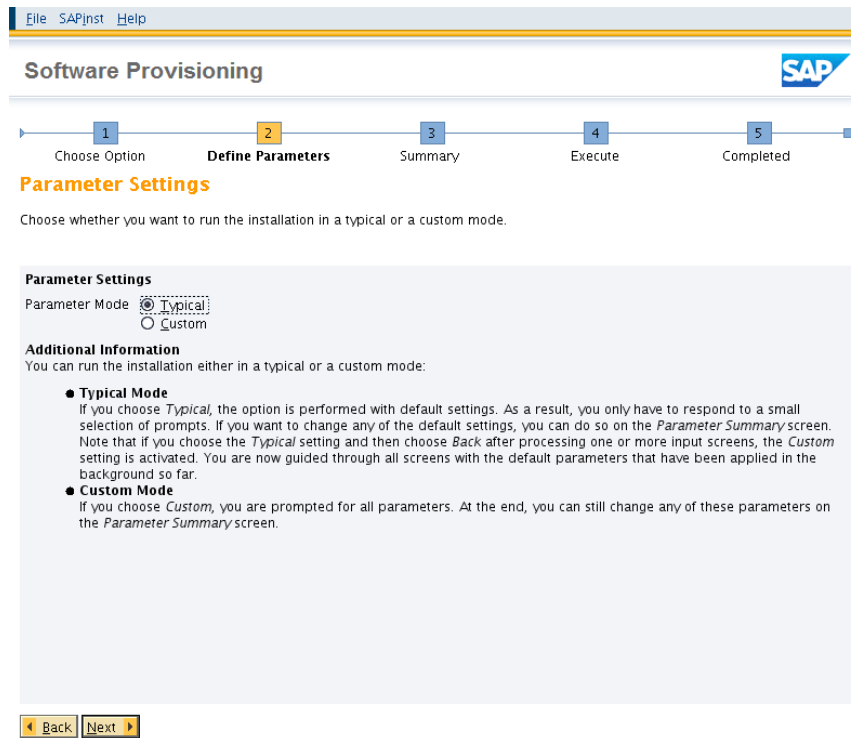
FIGURE 4.6: SAP INSTALLER: DEFINING PARAMETERS

💡 Tip: Installation Log Files

If the installation of the SAP application fails, refer to the installation log files. They are located in `/var/adm/autoinstall`. Failed installations are recorded in files with names ending in `.err`.

For more information about log files, see *Chapter 13, Important Log Files*.

14. The final screen is *Installation Completed*.
    To create an AutoYaST file for this installation, activate *Clone This System for AutoYaST*. The AutoYaST file will be placed in `/root/autoinst.xml`.
    Click *Finish*.

4.4 Continuing an Installation Using an Installation Profile

If you chose *Collect installation profiles but do not execute installation* during the previous run of the SAP Installation Wizard, this section shows you how to continue the installation of the selected SAP applications.

When collecting an installation profile, the SAP Installation Wizard copies product images to `/data/SAP_CDs`. It also prepares an installation environment for every product under the path `/data/SAP_INST`:

```
/data/SAP_INST/0/Instmaster
/data/SAP_INST/1/Instmaster
/data/SAP_INST/2/Instmaster
[...]
```

These files are re-used in the following. To continue the installation, follow these steps:

1. In `/etc/sysconfig/sap-installation-wizard`, set the following:

   ```
   SAP_AUTO_INSTALL="yes"
   ```

2. In the case of an SAP HANA/SAP BusinessOne installation, the SAP Installation Wizard will later use the parameters documented in the AutoYaST files in `/data/SAP_INST/number`. If you need to change any parameters, make sure to adapt the AutoYaST files at this point.

3. Open the YaST control center and start *SAP Installation Wizard*.

4. You will be asked whether to continue the pending installation. Select *Install*.

5. All further interactions happen within the SAP Installer. Follow the steps of SAP Installer as described in the documentation provided to you by SAP.

   - In the case of an SAP NetWeaver installation, all parameters of the SAP Installer will be offered again for fine-tuning.

   - In the case of an SAP HANA/SAP BusinessOne installation, the installer will not be offering to make any changes to parameters.

4.5 Partitioning for an SAP Application Without the SAP Installation Wizard

If you do not want to use the SAP Installation Wizard, you can also create the partitioning for SAP applications directly from the command line. First, find the correct partitioning file in the directory `/usr/share/YaST2/include/sap-installation-wizard/` or create an own partitioning file. For more information, see *Section 2.7.2, "Partitioning for the SAP System (stage 2)"*.

When you have determined the correct partitioning XML file, run:

```
root # yast2 sap_create_storage ABSOLUTE_PATH_TO_PARTITIONING_FILE
```

4.6 Automated Installation of SAP Applications with AutoYaST

The SAP Installation Wizard can be used from AutoYaST to automate the installation of SAP Applications.

4.6.1 SAP HANA Installation

The following AutoYaST snippet shows how an SAP HANA or SAP TREX installation can be automated:

```
<sap-inst>
  <products config:type="list">
    <product>
      <media config:type="list">
        <medium>
          <url>nfs://server/path1</url>
          <type>sap</type>
        </medium>
        <medium>
          <url>nfs://server/path3</url>
          <type>supplement</type>
        </medium>
      </media>
      <sapMasterPW>PASSWORD</sapMasterPW>
      <sid>SID</sid>
      <sapInstNr>INSTANCE_NUMBER</sapInstNr>
      <sapMDC>no</sapMDC>
```

```
      </product>
  </products>
</sap-inst>
```

- The `sapMDC` element is only applicable to SAP HANA.

- The `sapVirtHostname` element must be specified for distributed or highly available installations.

For a full SAP HANA example, including partitioning, see `/usr/share/doc/packages/sap-installation-wizard/hana-autoyast.xml`.

4.6.2   SAP NetWeaver Installation

For SAP NetWeaver, the following example shows how the installation can be automated. Specifically, this example is tailored to installing ASCS Instance of an SAP NetWeaver 7.5 ABAP Server distributed system with MaxDB (product ID `NW_ABAP_ASCS:NW750.ADA.ABAP`). When installing other products based on SAP NetWeaver, not all of the following variables may be necessary (or these variables might need to be replaced by others):

- The master password for the SAP NetWeaver instance: *MASTER_PASSWORD*

- The SAP Identifier (SID): *SID*

- The SAP kernel: *KERNEL*

- The SAP instance number: *INSTANCE_NUMBER*

- The ASCS virtual host name: *ASCS_VIRTUAL_HOSTNAME*

- The SCS virtual host name: *SCS_VIRTUAL_HOSTNAME*

```
<sap-inst>
  <products config:type="list">
    <product>
      <media config:type="list">
        <medium>
          <url>nfs://SERVER/PATH1</url>
          <type>sap</type>
        </medium>
        <medium>
          <url>nfs://SERVER/PATH2</url>
          <type>sap</type>
        </medium>
        <medium>
```

```
          <url>nfs://SERVER/PATH3</url>
          <type>supplement</type>
        </medium>
      </media>
      <productID>NW_ABAP_ASCS:NW750.ADA.ABAP</productID>
      <iniFile>
        <![CDATA[
# Password for the Diagnostics Agent specific <dasid>adm user. Provided value
# may be encoded.
DiagnosticsAgent.dasidAdmPassword =

# Windows domain in which the Diagnostics Agent users must be created.
# The property is Microsoft Windows only. This is an optional property.
DiagnosticsAgent.domain =

# Password for the Diagnostics Agent specific SAPService<DASID> user.
# Provided value may be encoded.
# The property is Microsoft Windows only.
DiagnosticsAgent.sapServiceDASIDPassword =

NW_GetMasterPassword.masterPwd = MASTER_PASSWORD

# Human readable form of the Default Login language - valid names are stored
# in a table of the subcomponent NW_languagesInLoadChecks. Used when freshly
# installing an ABAP stack for the machine that performs an ABAP load (in the
# case of a distributed system, that is the database, otherwise it is used by
# the normal installer). The available languages must be declared in the
# LANGUAGES_IN_LOAD parameter of the product.xml . In this file, the one
# character representation of the languages is used. Check the same table in
# the subcomponent mentioned above.
NW_GetSidNoProfiles.SAP_GUI_DEFAULT_LANGUAGE =

# The drive to use (Windows only)
NW_GetSidNoProfiles.sapdrive =

# The /sapmnt path (Unix only)
NW_GetSidNoProfiles.sapmnt = /sapmnt

# The SAP System ID of the system to install
NW_GetSidNoProfiles.sid = SID

# Will this system be unicode system?
NW_GetSidNoProfiles.unicode = true

NW_SAPCrypto.SAPCryptoFile = /data/SAP_CDs/745-UKERNEL-SAP-Unicode-Kernel-745/DBINDEP/
SAPEXE.SAR
```

```
NW_SCS_Instance.ascsInstanceNumber =

NW_SCS_Instance.ascsVirtualHostname = ASCS_VIRTUAL_HOSTNAME

NW_SCS_Instance.instanceNumber = INSTANCE_NUMBER

NW_SCS_Instance.scsInstanceNumber =

NW_SCS_Instance.scsMSPort =

NW_SCS_Instance.scsVirtualHostname = SCS_VIRTUAL_HOSTNAME

NW_System.installSAPHostAgent = true

NW_Unpack.igsExeSar =

NW_Unpack.igsHelperSar =

NW_Unpack.sapExeDbSar =

NW_Unpack.sapExeSar =

NW_Unpack.sapJvmSar =

NW_Unpack.xs2Sar =

NW_adaptProfile.templateFiles =

# The FQDN of the system.
NW_getFQDN.FQDN =

# Do we want to set the FQDN for the system?
NW_getFQDN.setFQDN = false

# The path to the JCE policy archive to install into the Java home directory
# if it is not already installed.
NW_getJavaHome.jcePolicyArchive =

hostAgent.domain =

# Password for the SAP Host Agent specific sapadm user. Provided value may be
# encoded.
hostAgent.sapAdmPassword = MASTER_PASSWORD

nwUsers.sapDomain =

nwUsers.sapServiceSIDPassword =
```

```
nwUsers.sidadmPassword =
            ]]>
      </iniFile>
    </product>
  </products>
</sap-inst>
```

SAP NetWeaver Installation   SLES-SAP 12 SP5

5 Setting Up an Installation Server for SAP Media Sets

Using the SAP Installation Wizard, it is possible to copy the SAP media sets from a remote server (for example, via NFS or SMB). However, using the option provided there means that you need to install the product at the same time. Additionally, it does not allow for copying all SAP media used in your organization to a single server.

However, you can easily create such a server on your own. For example, to put the SAP media sets on an NFS Server, proceed as follows:

PROCEDURE 5.1: ADDING SAP PRODUCT INSTALLATION FILES TO AN NFS SERVER

1. On your installation server, create the directory `/srv/www/htdocs/sap_repo`.

2. Open the file `/etc/exports` and add the following:

   ```
   /srv/www/htdocs/sap_repo *(ro,no_root_squash,sync,no_subtree_check,insecure)
   ```

   > ⚠ **Important: Executable Rights Must Be Visible**
   >
   > Clients must be able to see which files are executable. Otherwise, SUSE's SAP Installation Wizard cannot execute the SAP Installer.

3. In `/srv/www/htdocs/sap_repo`, create a directory for every SAP medium you have. Give these directories speaking names, so you can identify them later on. For example, you could use names like `kernel`, `java`, or `hana`.

4. Copy the contents of each SAP medium to the corresponding directory with **`cp -a`**.

   > ⚠ **Important: Avoid Using Windows* Operating Systems for Copying**
   >
   > Using Windows operating system for copying or copying from/to Windows file systems like NTFS can break permission settings and capitalization of files and directories.

You can now install from the NFS server you set up. In the SAP Installation Wizard, specify the path this way: *server_name*`/srv/www/htdocs/sap_repo`. For more information about specifying the path, see *Table 4.1, "Media Source Path"*.

For information about setting up an NFS server from scratch, see *Administration Guide, Part "Services", Chapter "Sharing File Systems with NFS", Section "Installing NFS Server"* (https://documentation.suse.com/sles-12 ↗).

For information about installing SUSE Linux Enterprise Server from an NFS server, see *Deployment Guide, Chapter "Remote Installation", Section "Setting Up an NFS Repository Manually"* (https://documentation.suse.com/sles-12 ↗).

6 Setting Up an SAP HANA Cluster

You can use a YaST wizard to set up SAP HANA or SAP S/4HANA Database Server clusters according to best practices, including SAP HANA system replication. A summary of the setup options is given in *Section 1.1.3, "Simplified SAP HANA System Replication Setup"*.

The following *Best Practices* from the SUSE Linux Enterprise Server for SAP Applications Resource Library (https://www.suse.com/products/sles-for-sap/resource-library/ ↗) contain setup instructions:

- **Performance-optimized scenario and multi-tier/chained scenario:** *Setting up a SAP HANA SR Performance Optimized Infrastructure*

- **Cost-optimized scenario:** *Setting up a SAP HANA SR Cost Optimized Infrastructure*

> **!** ## Important: Wizard Can Only Be Used for Initial Configuration
>
> The YaST wizard described in the following can only be used for the initial cluster configuration.
>
> To reconfigure a cluster, use the separate YaST module *Cluster* (available from package `yast2-cluster`). For more information about its usage, see *Administration Guide, Part "Installation, Setup and Upgrade", Chapter "Using the YaST Cluster Module"* at https://documentation.suse.com/sle-ha-12 ↗.

> **!** ## Important: Not Suitable for Cloud Deployments
>
> The instructions in the following sections are suitable for deployments on physical machines. They are not intended for cloud-based deployments.

6.1 Prerequisites

The following procedure has prerequisites:

- Two machines which both have an SAP HANA installation created by the SAP Installation Wizard or SAP HANA Application Lifecycle Management. Both machines need to be on the same L2 network (subnet).

In the case of a multi-tier/chained scenario, there must also be a third machine elsewhere.

- The machines are not yet set up as a high-availability cluster.

- openSSH is running on both machines and the nodes can reach each other via SSH. However, if that has not already happened, the wizard will perform the SSH key exchange itself. For more information about SSH, see *Security Guide, Part "Network Security", Chapter "SSH: Secure Network Operations"* at https://documentation.suse.com/sles-12 ↗.

- A disk device that is available to both nodes under the same path for SBD. It must not use host-based RAID, cLVM2 or reside on a DRBD instance. The device can have a small size, for example, 100 MB.

- You have created either:

    - A key in the SAP HANA Secure User Store on the primary node

    - An initial SAP HANA backup on the primary node

- The package `yast2-sap-ha` is installed on both the primary and the secondary node.

- HANA-Firewall is set up on both computers with the rules `HANA_HIGH_AVAILABILITY` and `HANA_SYSTEM_REPLICATION` on all relevant network interfaces.
  For information about setting up HANA-Firewall, see *Section 8.2, "Configuring HANA-Firewall"*.

- *Cost-optimized scenario only:* The secondary node has a second SAP HANA installation. The database may be running but will be stopped automatically by the wizard.

- *Cost-optimized scenario only:* For the non-production SAP HANA instance, you have created an SAP HANA Secure User Store key `QASSAPDBCTRL` for monitoring purposes. For more information, see *SAP HANA SR Cost Optimized Scenario, Chapter "Installing the SAP HANA Databases on both cluster nodes", Section "Postinstallation configuration", Section "Install the non-productive SAP HANA database (QAS)"* at https://www.suse.com/products/sles-for-sap/resource-library/ ↗.

6.2   Setup

The following procedure needs to be executed on the primary node (also called the "master"). Before proceeding, make sure the prerequisites listed in *Section 6.1, "Prerequisites"* are fulfilled.

1. Open the YaST control center. In it, click *HA Setup for SAP Products* in the category *High Availability*.

2. If an SAP HANA instance has been detected, you can choose between the scale-up scenarios *Performance-optimized, Cost-optimized,* or *Chained (multi-tier)*. For information about these scale-up scenarios, see *Section 1.1.3, "Simplified SAP HANA System Replication Setup"*.
   Continue with *Next*.
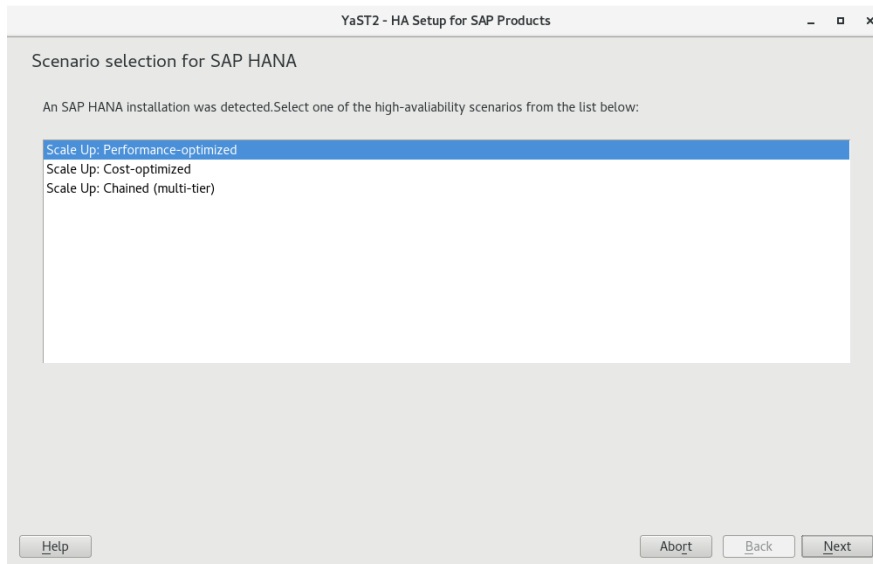
3. This step of the wizard presents a list of prerequisites for the chosen scale-up scenario. These prerequisites are the same as those presented in *Section 6.1, "Prerequisites"*.
   Continue with *Next*.

4. The next step lets you configure the communication layer of your cluster.

   - Provide a name for the cluster.

   - The default transport mode *Unicast* is usually appropriate.

   - Under *Number of rings*, a single communication ring usually suffices.
     For redundancy, it is often better to use network interface bonding instead of multiple communication rings. For more information, see *Administration Guide, Part "Configuration and Administration", Chapter "Network Device Bonding"* at https://documentation.suse.com/sle-ha-12 .

   - From the list of communication rings, configure each enabled ring. To do so, click *Edit selected*, then select a network mask (*IP address*) and a port (*Port number*) to communicate over.

Finish with *OK*.

- Additionally, decide whether to enable the configuration synchronization service Csync2 and Corosync secure authentication using HMAC/SHA1.
  For more information about Csync2, see *Administration Guide Part "Installation, Setup and Upgrade", Chapter "Using the YaST Cluster Module", Section "Transferring the Configuration to All Nodes"* at https://documentation.suse.com/sle-ha-12 ↗.
  For more information about Corosync secure authentication, see *Administration Guide, Part "Installation, Setup and Upgrade", Chapter "Using the YaST Cluster Module", Section "Defining Authentication Settings"* at https://documentation.suse.com/sle-ha-12 ↗.

Proceed with *Next*.

5. The wizard will now check whether it can connect to the secondary machine using SSH. If it can, it will ask for the `root` password to the machine.
   Enter the `root` password.
   The next time, the primary machine needs to connect to the secondary machine, it will connect using an SSH certificate instead of a password.

6. For both machines, set up the host names and IP address (for each ring).
   Host names chosen here are independent from the virtual host names chosen in SAP HANA. However, to avoid issues with SAP HANA, host names must not include hyphen characters (`-`).

If this has not already been done before, host names of all cluster servers must now be added to the file `/etc/hosts`. For this purpose, activate *Append to /etc/hosts*.
Proceed with *Next*.

7. If NTP is not yet set up, do so. This avoids the two machines from running into issues because of time differences.

   a. Click *Reconfigure*.

   b. On the tab *General Settings*, activate *Now and on Boot*.

   c. Add a time server by clicking *Add*. Click *Server* and *Next*. Then specify the IP address of a time server outside of the cluster. Test the connection to the server by clicking *Test*.
   To use a public time server, click *Select › Public server* and select a time server. Finish with *OK*.
   Proceed with *OK*.

   d. On the tab *Security Settings*, activate *Open Port in Firewall*.

   e. Proceed with *Next*.

8. In the next step, choose fencing options. The YaST wizard only supports the fencing mechanism SBD (*STONITH block device*). To avoid split-brain situations, SBD uses a disk device which stores cluster state.
   The chosen disk must be available from all machines in the cluster under the same path. Ideally, use either *by-uuid* or *by-path* for identification.
   The disk must not use host-based RAID, cLVM2 or reside on a DRBD instance. The device can have a small size, for example, 100 MB.

   ## Warning: Data on Device Will Be Lost

   All data on the chosen SBD device or devices will be deleted.

   To define a device to use, click *Add*, then choose an identification method such as *by-uuid* and select the appropriate device. Click *OK*.
   To define additional SBD command line parameters, add them to *SBD options*.
   If your machines reboot particularly fast, activate *Delay SBD start*.

For more information about fencing, see the *Administration Guide* at https://documenta-tion.suse.com/sle-ha-12 ↗ .

Proceed with *Next.*

9. The following page allows configuring watchdogs. They protect against the failure of the SBD daemon itself and force a reboot of the machine in such a case.

It also lists watchdogs already configured using YaST and watchdogs that are currently loaded (as detected by `lsmod`).

To configure a watchdog, use *Add.* Then choose the correct watchdog for your hardware and leave the dialog with *OK.*

For testing, you can use the watchdog `softdog`. However, we highly recommend using a hardware watchdog in production environments instead of `softdog`. For more information about selecting watchdogs, see *Administration Guide, Part "Storage and Data Replication", Chapter "Storage Protection", Section "Conceptual Overview", Section "Setting Up Storage-based Protection", Section "Setting up the Watchdog"* at https://documenta-tion.suse.com/sle-ha-12 ↗ .

Proceed with *Next.*

10. Set up the parameters for your SAP HANA installation or installations. If you have selected the cost-optimized scenario, additionally, fill out details related to the non-production SAP HANA instance.

Production SAP HANA Instance

- Make sure that the *System ID* and *Instance number* match those of your SAP HANA configuration.

- *Replication mode* and *Operation mode* usually do not need to be changed.
  For more information about these parameters, see the *HANA Administration Guide* provided to you by SAP.

- Under *Virtual IP address*, specify a virtual IP address for the primary SAP HANA instance. Under *Virtual IP Mask*, set the length of the subnetwork mask in CIDR format to be applied to the *Virtual IP address.*

- *Prefer site takeover* defines whether the secondary instance should take over the job of the primary instance automatically (*true*). Alternatively, the cluster will restart SAP HANA on the primary machine.

- *Automatic registration* determines whether primary and secondary machine should switch roles after a takeover.

- Specify the site names for the production SAP HANA instance on the two nodes in *Site name 1* and *Site name 2*.

- Having a backup of the database is a precondition for setting up SAP HANA replication.
  If you have not previously created a backup, activate *Create initial backup*. Under *Backup settings*, configure the *File name* and the *Secure store key* for the backup. The key in the SAP HANA Secure User Store on the primary node must have been created before starting the wizard.
  For more information, see the documentation provided to you by SAP.

- *Cost-optimized scenario only:* Within *Production system constraints*, configure how the production instance of SAP HANA should behave while inactive on the secondary node.
  Setting the *Global allocation limit* allows directly limiting memory usage. Activating *Preload column tables* will increase memory usage.
  For information about the necessary global allocation limit, see documentation provided to you by SAP such as *How to Perform System Replication for SAP HANA* at https://archive.sap.com/documents/docs/DOC-47702 ↗.

Cost-optimized Scenario Only: Non-production SAP HANA Instance

- Make sure that the *System ID* and *Instance number* match those of your non-production SAP HANA instance.
  These parameters are needed to allow monitoring the status of the non-production SAP HANA instance using the SAPInstance resource agent.

- Generate a hook script for stopping the non-production instance and starting the production instance and removing the constraints on the production system. The script is written in Python 2 and can be modified as necessary later.
  Click *Hook script* and then set up the correct user name and password for the database. Then click *OK*.
  You can now manually verify and change the details of the generated hook script. When you are done, click *OK* to save the hook script at `/hana/shared/SID/srHook`.

⛔ Warning: Passwords Stored in Plain Text

By default, the hook script stores all credentials in plain text. To improve security, modify the script yourself.

Proceed with *Next*.

FIGURE 6.1: **SAP HANA OPTIONS (COST-OPTIMIZED SCENARIO)**

11. On the page *High-Availability Configuration Overview*, check that the setup is correct.
    To change any of the configuration details, return to the appropriate wizard page by clicking one of the underlined headlines.
    Proceed with *Install*.

12. When asked whether to install additional software, confirm with *Install*.

13. After the setup is done, there is a screen showing a log of the cluster setup.
    To be able to reuse the configuration file for an unattended installation, click *Save configuration*. (For more details about unattended cluster installations, see *Section 6.3, "Unattended Installation of Cluster Servers".*)
    To close the dialog, click *Finish*.

14. *Multi-tier/chain scenario only:* Using the administrative user account for the production SAP HANA instance, register the out-of-cluster node for system replication:

```
SIDadm > hdbnsutil -sr_register --remoteHost=SECONDARY_HOST_NAME \
```

```
--remoteInstance=INSTANCE_NUMBER --replicationMode=async \
--name=SITE_NAME
```

6.3   Unattended Installation of Cluster Servers

On bare-metal servers, you can run the YaST module *HA Setup for SAP Products* in an unattended mode. This speeds up the time it takes to configure additional servers with the same configuration.

As a prerequisite, you need a configuration file from a previous run of the YaST wizard *HA Setup for SAP Products*.

1. Copy the configuration file to the target machine.

2. Validate the configuration file:

```
root # yast2 sap_ha readconfig CONFIGURATION_FILE_PATH
```

3. Import, validate, and install the cluster unattendedly:

```
root # yast2 sap_ha readconfig CONFIGURATION_FILE_PATH unattended
```

6.4   Using Hawk

After you have set up the cluster using the wizard, you can open Hawk. directly from the last screen of the *HA Setup for SAP Products* wizard.

To revisit Hawk, open a browser. As the URL, enter the IP address or host name of any cluster node running the Hawk Web service. Alternatively, enter the virtual IP address you configured in *Section 6.2, "Setup"*.

```
https://HAWKSERVER:7630/
```

On the Hawk login screen, use the following login credentials:

- *Username*: `hacluster`

- *Password*: `linux`

> ⓘ **Important: Secure Password**
>
> Replace the default password with a secure one as soon as possible:
>
> ```
> root # passwd hacluster
> ```

For more information about Hawk, see *Administration Guide, Part "Configuration and Administration", Chapter "Configuring and Managing Cluster Resources with Hawk"* (https://documentation.suse.com/sle-ha-12 ↗).

7 Tuning

This chapter presents information about tuning SUSE Linux Enterprise Server for SAP Applications to work optimally with SAP applications.

On SUSE Linux Enterprise Server for SAP Applications you have the choice between `sapconf` and `saptune`. However, `saptune` is the more elaborate tool that offers more features.

> 🖊 **Note: The `sapconf` Command Has Been Removed**
>
> In SUSE Linux Enterprise Server and SUSE Linux Enterprise Server for SAP Applications 11 and 12, the `sapconf` command was included in the package with the same name.
>
> For SUSE Linux Enterprise Server and SUSE Linux Enterprise Server for SAP Applications 15 this has been changed: the command `sapconf` have been removed from the `sapconf` package. The package contains a `systemd` service only. There is no `sapconf` command line tool anymore, no `sapconf`/`tuned` profiles, and no `tuned`.

7.1 Kernel: Page-Cache Limit

Problem

The kernel swaps out rarely accessed memory pages to use freed memory pages as cache to speed up file system operations, for example during backup operations.

SAP NetWeaver and SAP HANA use large amounts of memory for accelerated access to business data. Parts of this memory are rarely accessed. When a user request needs to access paged-out memory, the response time is poor. It is even worse when an SAP application running on Java incurs a Java garbage collection: The system starts heavy page-in (disk I/O) activity and has a poor response time for an extended period of time.

Solution

SUSE Linux Enterprise Server for SAP Applications includes a kernel tuning option. It allows the system administrator to limit the amount of page cache that the kernel uses when there is competition between application memory and page cache. This option tells the kernel that when the page cache is filled to the configured limit, application memory is more important and should thus not be paged out. No pages will be paged out if the memory footprint of the workload plus the configured page-cache limit do not exceed the amount of physical RAM in the system.

These kernel options are available for configuration:

- `vm.pagecache_limit_mb` (`/proc/sys/vm/pagecache_limit_mb`)

- `vm.pagecache_limit_ignore_dirty`   (`/proc/sys/vm/pagecache_limit_ig-nore_dirty`)

> ### 💡 Tip: Use `sapconf`/**saptune** to Configure Parameters
>
> The parameters `vm.pagecache_limit_mb` and `vm.pagecache_limit_ig-nore_dirty` are also configured by **saptune** and `sapconf`.
>
> For more information, see *Section 7.2, "Tuning Systems with `sapconf4`"*, *Section 7.3, "Tuning Systems with `sapconf 5`"*, or *Section 7.4, "Tuning Systems with **saptune**"*.

> ### ⚠ Important: The Following Are Example Values
>
> The values reproduced in *Example 7.1, "Permanently Setting the Page-Cache Limit"* are example values only. Do not set the following parameters on a productive system without first trying and calibrating them on a non-productive system.
>
> If your system does not exhibit page-cache limit issues under the workloads it is running, there is no need to adjust these parameters.
>
> For more information, see *SAP Note 1557506: Linux Paging Improvements* (https://launchpad.support.sap.com/#/notes/1557506 ↗).

EXAMPLE 7.1: **PERMANENTLY SETTING THE PAGE-CACHE LIMIT**

For permanent use, add both parameters to `/etc/sysctl.conf`, for example:

```
vm.pagecache_limit_mb = 1024
vm.pagecache_limit_ignore_dirty = 2
```

7.2 Tuning Systems with `sapconf4`

The package `sapconf` contains the `tuned` profiles `sap-netweaver`, `sap-hana`, `sap-ase` and `sap-bobj`. These tuning profiles set recommended parameters for SAP NetWeaver, SAP HANA, and SAP HANA-based applications.

(The profile `sap-netweaver` (default) is currently identical to `sap-hana`. `sapconf` additionally includes the profiles `sap-ase` and `sap-bobj` which are currently not recommended for use.)

`sapconf` consists of two primary parts:

- A `systemd` service that ensures `tuned` and related services are running and the sapconf-based profile is applied.

- The `tuned` profile that applies configured `sapconf` tuning parameters using a script and configuration files.

To use `sapconf`, install the package `sapconf` on your system.

> 📝 **Note: Unified Profiles in SUSE Linux Enterprise Server and SUSE Linux Enterprise Server for SAP Applications 15**
>
> In SUSE Linux Enterprise Server and SUSE Linux Enterprise Server for SAP Applications 15 and up, only a single `tuned` profile, `sapconf`, is shipped. It is equivalent to the profiles `sap-hana`/`sap-netweaver` (both profiles are identical) that are shipped with `sapconf` in SUSE Linux Enterprise Server for SAP Applications 15.

7.2.1 Enabling and Disabling `sapconf` and Viewing Its Status

After the installation of `sapconf`, `tuned` is enabled and the `sap-netweaver` profile is activated. However, if another `tuned` profile is already enabled, `sapconf` will not enable its own `tuned` profile. If you previously made other changes to system tuning, these may be overridden by the profile of `sapconf`.

To make sure `sapconf` applies all tuning parameters, reboot the machine after installation.

You can inspect or change the status of `sapconf` as described in the following:

- To see the status of the service `sapconf`:

  ```
  root # systemctl status sapconf
  ```

  The service should be displayed as *active (exited)*, as it is only responsible for starting `tuned` and will exit afterward.

- To start the service `sapconf` and with it the service `tuned`:

  ```
  root # systemctl start sapconf
  ```

- Should `sapconf` be disabled, enable and start it with:

```
root # systemctl enable --now sapconf
```

- To stop the service `sapconf` and with it the service `tuned`:

```
root # systemctl stop sapconf
```

  As `tuned` is also terminated at that point, the vast majority of optimizations which do not require a system reboot to disable are disabled immediately.

- To disable `sapconf`, use:

```
root # systemctl disable sapconf
```

  If you yourself have not specifically enabled any of the services that `sapconf` depends on, this will also disable most tuning parameters and all services used by `sapconf`.

Similarly, you can inspect and change the status of the underlying service `tuned`:

- To see the status of the service `tuned`:

```
root # systemctl status tuned
```

- To see which `tuned` profile is currently in use:

```
root # tuned-adm active
```

  If this command does not return the name of the required profile, enable that profile with:

```
root # tuned-adm profile sap-netweaver
```

> 💡 **Tip: Additional Services that `sapconf` Relies On**
>
> In addition to the `sapconf` service itself and the `tuned` service, `sapconf` also relies on the following two services:
>
> - `sysstat` collects data on system activity.
>
> - `uuidd` which generates time-based UUIDs that are guaranteed to be unique even in settings where many processor cores are involved. This is necessary for SAP applications.

7.2.2 Configuring `sapconf4`

In general, the default configuration of `sapconf` already uses the parameter values recommended by SAP. However, if you have special needs, you can configure the tool to better suit those.

The configuration of `sapconf` is split in two parts that can be configured in different ways:

`/usr/lib/tuned/`*PROFILE*`/tuned.conf`

Any file that adheres to this pattern can be edited like in *Procedure 7.1, "Configuring* `sapconf4` *Profiles"*. To configure parameters from this file, copy it to the custom profile directory of `tuned` under `/etc/tuned` first and then change values in it. If you change the file in place instead, you will lose the changes you make on the next update of the `sapconf` package. The following procedure shows an example how to adapt the file `/usr/lib/tuned/sap-netweaver/tuned.conf`. However, as written before, this is possible with any profile. Configure the file as described in the following procedure:

PROCEDURE 7.1: CONFIGURING `sapconf4` **PROFILES**

1. Create a new custom `tuned` profile directory and copy the file `tuned.conf`:

   ```
   root # mkdir /etc/tuned/sap-netweaver
   root # cp /usr/lib/tuned/sap-netweaver/tuned.conf /etc/tuned/sap-netweaver/
   ```

2. Within the newly copied `tuned.conf`, fix the reference to `script.sh` to use an absolute path that points to the script from the original profile:

   ```
   script = /usr/lib/tuned/sap-netweaver/script.sh
   ```

   Do not instead copy `script.sh`, as that provokes update compatibility issues for `sapconf`.

3. Edit the parameters in `/etc/tuned/sap-netweaver/tuned.conf`.

After each update to `sapconf`, make sure to compare the contents of the original and the custom `tuned.conf`.
Log messages related to this file are written to `/var/log/tuned/tuned.log`.

`/etc/sysconfig/sapconf`

This file contains most parameters of `sapconf`. The parameters from this file are applied using the aforementioned script `/usr/lib/tuned/`*PROFILE*`/script.sh`.

This file can be edited directly. All parameters in this file are explained by means of comments and references to SAP Notes which can be viewed at https://launchpad.support.sap.com/ ↗.

When `sapconf` is updated, all customized parameters from this file will be preserved as much as possible. However, sometimes parameters cannot be transferred cleanly to the new configuration file. The previous custom configuration is moved to `/etc/sysconfig/sapconf.rpmsave` during the update. After updating check the difference between the previous custom configuration and the new version at `/etc/sysconfig/sapconf`.

Log messages related to this file are written to `/var/log/sapconf.log`.

When editing either of these files, you will find that some values are commented with a `#` character at the beginning of the line. This means that while the parameter is relevant for tuning, there is no suitable default for it.

Conversely, you can add `#` characters to the beginning of the line to comment specific parameters. However, you should avoid this practice, as it can lead to `sapconf` not properly applying the profile.

To apply edited configuration, restart `sapconf`:

```
root # systemctl restart sapconf
```

Confirming that a certain parameter value was applied correctly works differently for different parameters. Hence, the following serves as an example only:

EXAMPLE 7.2: CHECKING PARAMETERS

To confirm that the setting for `TCP_SLOW_START` was applied, do the following:

- View the log file of `sapconf` to see whether it applied the value. Within `/var/log/sapconf.log`, check for a line containing this text:

```
Change net.ipv4.tcp_slow_start_after_idle from 1 to 0
```

Alternatively, the parameter may have already been set correctly before `sapconf` was started. In this case, `sapconf` will not change its value:

```
Leaving net.ipv4.tcp_slow_start_after_idle unchanged at 1
```

- The underlying option behind `TCP_SLOW_START` can be manually configured at `/proc/sys/net.ipv4.tcp_slow_start_after_idle`. To check its actual current value, use:

```
root # sysctl net.ipv4.tcp_slow_start_after_idle
```

7.2.3 Removing `sapconf`

To remove `sapconf` from a system, uninstall its package with:

```
root # zypper rm sapconf
```

Note that when doing this, dependencies of `sapconf` will remain installed. However, the services `sysstat` and `tuned` will go into a disabled state. If either is still relevant to you, make sure to enable it again.

Certain parameters and files are not removed when `sapconf` is uninstalled. For more information, see the man page `man 7 sapconf`, section *PACKAGE REQUIREMENTS*.

7.2.4 For More Information

The following man pages provide additional information about `sapconf`:

- An overview of tuning parameters for the profiles: `man 7 profile-sapconf-hana` `man 7 profile-sapconf-netweaver` `man 7 profile-sapconf-ase` `man 7 profile-sapconf-bobj`

- Detailed description of all tuning parameters set by `sapconf`: `man 5 sapconf`

- Information about configuring and customizing the `sapconf` profile: `man 7 sapconf`

Also see the blog series detailing the updated version of `sapconf` at https://www.suse.com/c/a-new-sapconf-is-available/ .

7.3 Tuning Systems with `sapconf` 5

The package `sapconf` is available in SUSE Linux Enterprise Server and SUSE Linux Enterprise Server for SAP Applications. It sets recommended parameters for the following types of SAP applications: SAP NetWeaver, SAP HANA and SAP HANA-based applications.

OVERVIEW OF `sapconf`5 IN SUSE® LINUX ENTERPRISE SERVER 12

`sapconf` **5 (without** `tuned` **)**

- `sapconf-netweaver` (`sapconf` profile as a replacement for `tuned` profile)

- `sapconf-hana` (`sapconf` profile as a replacement for `tuned` profile)

- `sapconf-bobj` (`sapconf` profile as a replacement for `tuned` profile)

- `sapconf-ase` (`sapconf` profile as a replacement for `tuned` profile)

OVERVIEW OF `sapconf`5 IN SUSE® LINUX ENTERPRISE SERVER 15

`sapconf` **5 (without** `tuned` **)**
    no profiles anymore

The package `sapconf` contains the `sapconf` profiles `netweaver`, `hana`, `ase` and `bobj`. These profiles set recommended parameters for SAP NetWeaver, SAP HANA, and SAP HANA-based applications.

(The profile `netweaver` (default) is currently identical to `hana`. `sapconf` additionally includes the profiles `ase` and `bobj` which are currently not recommended for use.)

To use `sapconf`, install the package `sapconf` on your system.

> **Note: No Profiles in SUSE Linux Enterprise Server and SUSE Linux Enterprise Server for SAP Applications 15**
>
> In SUSE Linux Enterprise Server and SUSE Linux Enterprise Server for SAP Applications 15, `sapconf` no longer supports profiles.

7.3.1 Enabling and Disabling `sapconf` and Viewing It's Status

After the installation of `sapconf`, the `netweaver` profile is activated.

To make sure `sapconf` applies all tuning parameters, reboot the machine after installation.

You can inspect or change the status of `sapconf` as described in the following:

- To see the status of the service `sapconf`:

```
root # systemctl status sapconf
```

  The service should be displayed as *active (exited)*.

- To start the service `sapconf`:

```
root # systemctl start sapconf
```

- Should `sapconf` be disabled, enable and start it with:

```
root # systemctl enable --now sapconf
```

- To stop the service `sapconf`:

```
root # systemctl stop sapconf
```

  The vast majority of optimizations which do not require a system reboot to disable are disabled immediately.

- To disable `sapconf`, use:

```
root # systemctl disable sapconf
```

  If you yourself have not specifically enabled any of the services that `sapconf` depends on, this will also disable most tuning parameters and all services used by `sapconf`.

- To see which profile is currently in use:

```
root # cat /var/lib/sapconf/act_profile
```

- To change the profile use:

```
root # sapconf stop && sapconf PROFILE
```

Tip: Additional Services that `sapconf` Relies On

In addition to the `sapconf` service the following two services are also relied on:

- `sysstat` collects data on system activity.

- `uuidd` which generates time-based UUIDs that are guaranteed to be unique even in settings where many processor cores are involved. This is necessary for SAP applications.

7.3.2 Configuring `sapconf5`

In general, the default configuration of `sapconf` already uses the parameter values recommended by SAP. However, if you have special needs, you can configure the tool to better suit those.

`/etc/sysconfig/sapconf`

This file contains all parameters of `sapconf` and can be edited directly. All parameters in this file are explained by means of comments and references to SAP Notes which can be viewed at https://launchpad.support.sap.com/ ↗ .

When `sapconf` is updated, all customized parameters from this file will be preserved as much as possible. However, sometimes parameters cannot be transferred cleanly to the new configuration file. The previous custom configuration is moved to `/etc/sysconfig/sapconf.rpmsave` during the update. After updating check the difference between the previous custom configuration and the new version at `/etc/sysconfig/sapconf`.

Log messages related to this file are written to `/var/log/sapconf.log`.

When editing either of these files, you will find that some values are commented with a `#` character at the beginning of the line. This means that while the parameter is relevant for tuning, there is no suitable default for it.

Conversely, you can add `#` characters to the beginning of the line to comment specific parameters. However, you should avoid this practice, as it can lead to `sapconf` not properly applying the profile.

To apply edited configuration, restart `sapconf`:

```
root # systemctl restart sapconf
```

Confirming that a certain parameter value was applied correctly works differently for different parameters. Hence, the following serves as an example only:

EXAMPLE 7.3: CHECKING PARAMETERS

> To confirm that the setting for `TCP_SLOW_START` was applied, do the following:
>
> * View the log file of `sapconf` to see whether it applied the value. Within `/var/log/sapconf.log`, check for a line containing this text:
>
>   ```
>   Change net.ipv4.tcp_slow_start_after_idle from 1 to 0
>   ```
>
>   Alternatively, the parameter may have already been set correctly before `sapconf` was started. In this case, `sapconf` will not change its value:
>
>   ```
>   Leaving net.ipv4.tcp_slow_start_after_idle unchanged at 1
>   ```
>
> * The underlying option behind `TCP_SLOW_START` can be manually configured at `/proc/sys/net.ipv4.tcp_slow_start_after_idle`. To check its actual current value, use:
>
>   ```
>   root # sysctl net.ipv4.tcp_slow_start_after_idle
>   ```

7.3.3 Removing `sapconf`

To remove `sapconf` from a system, uninstall its package with:

```
root # zypper rm sapconf
```

Note that when doing this, dependencies of `sapconf` will remain installed. However, the `sysstat` service will go into a disabled state. If it is still relevant to you, make sure to enable it again.

Certain parameters and files are not removed when `sapconf` is uninstalled. For more information, see the man page `man 7 sapconf`, section *PACKAGE REQUIREMENTS*.

7.3.4 For More Information

The following man pages provide additional information about `sapconf`:

- An overview of tuning parameters for the profiles: **`man 7 profile-sapconf-hana`** **`man 7 profile-sapconf-netweaver`** **`man 7 profile-sapconf-ase`** **`man 7 profile-sapconf-bobj`**

- Detailed description of all tuning parameters set by `sapconf`: **`man 5 sapconf`**

- Information about configuring and customizing the `sapconf` profile: **`man 7 sapconf`**

Also see the blog series detailing the updated version of `sapconf` at https://www.suse.com/c/a-new-sapconf-is-available/ ↗.

7.3.5 Using **tuned** Together with `sapconf`

With version 5 sapconf does not rely on tuned anymore. This means both tools can be used independently. `sapconf` will print a warning in it's log if **tuned** service is started.

> ⚠️ **Important: Using tuned and sapconf together**
>
> If you are going to use **`tuned`** and `sapconf` simultaneously, be very careful, that both tools do not configure the same system parameters.

7.4 Tuning Systems with **saptune**

Using **`saptune`**, you can tune a system for SAP NetWeaver, SAP HANA/SAP BusinessObjects, and SAP S/4HANA applications. This method relies on the system tuning service `tuned`.

If you used the SAP Installation Wizard to install an SAP application, `tuned` is usually already active. It is configured with a profile for the application you installed.

If you did not use the SAP Installation Wizard to install an SAP application, make sure that the packages `tuned` and `saptune` are installed on your system.

7.4.1 Enabling **saptune** to Tune for an SAP Application

1. To tune a system, first find a tuning solutions. To find the appropriate solution, use:

```
tux > saptune solution list
```

saptune knows the following tuning solutions (groups of SAP notes):

- `BOBJ`. Solution for running SAP BusinessObjects.

- `HANA`. Solution for running a SAP HANA database.

- `MAXDB`. Solution for running an SAP MaxDB database.

- `NETWEAVER`. Solution for running SAP NetWeaver application servers.

- `S4HANA-APPSERVER`. Solution for running SAP S/4HANA application servers (identical to SAP NetWeaver solution).

- `S4HANA-APP+DB`. Solution for running both SAP S/4HANA application servers and SAP HANA on the same host (identical to SAP NetWeaver + SAP HANA solution).

- `S4HANA-DBSERVER`. Solution for running the SAP HANA database of an SAP S/4HANA installation (identical to SAP HANA solution)

- `SAP-ASE`. Solution for running an SAP Adaptive Server Enterprise database.

Alternatively, you can tune the computer according to recommendations from specific SAP Notes. A list of notes that you can tune for is available via:

```
root # saptune note list
```

2. - To set up **saptune** with a preconfigured solution, use:

  ```
  root # saptune solution apply SOLUTION
  ```

  - To set up **saptune** for the recommendations of a specific SAP Note, use:

  ```
  root # saptune note apply NOTE
  ```

> ⊘ **Note: Combining Optimizations**
>
> You can combine solutions and notes. However, only one solution can be active at a time. In rare cases, notes can have conflicting options or parameters. To avoid conflicts, order your notes, keeping in mind that the last note always overrides conflicting options or parameters of previous notes.

3. To start `saptune` and enable it at boot, make sure to run the following command:

```
root # saptune daemon start
```

In the background, `saptune` applies a `tuned` profile also named `saptune` that is dynamically customized according to selected "solutions" and "notes". Using `tuned-adm list`, you can also see this profile.

7.4.2   Customizing an SAP Note

Every SAP note can be configured freely with:

```
root # saptune note customise
```

The command includes changing a value or disabling a parameter.

7.4.3   Creating a new SAP Note

It is possible to create a new SAP note with:

```
root # saptune note create
```

All features of `saptune` are available.

7.4.4   Deleting a SAP Note

This command allows to delete a created note including the corresponding override after confirmation:

```
root # saptune note delete
```

The note may not be applied at the time. SAP notes shipped by **saptune** cannot be deleted. Instead, the override file is removed when available.

7.4.5   Renaming a SAP Note

This command allows to rename a created note including the corresponding override after confirmation:

```
root # saptune note rename
```

The note may not be applied at the time. SAP notes shipped by **saptune** cannot be renamed.

7.4.6   Showing the Configuration of an SAP Note

The shipped configuration of a note can be listed with:

```
root # saptune note show
```

7.4.7   Verifying an SAP Note or an SAP Solution

The commands **saptune note verify** *NOTE* and **saptune solution verify** *SOLUTION* list the following data for each active or requested note:

- The parameter name

- The expected value (default)

- A configured override (created using **saptune customise**)

- The current system value

- Whether the current state is in accordance with the SAP recommendation

7.4.8   Simulating the Application of an SAP Note or an SAP Solution

To show each parameter of a `note`, use:

```
root # saptune note simulate
```

To show each parameter of a `solution`, use:

```
root # saptune solution simulate
```

It lists the current system value and the expected values (default and override).

7.4.9   Disabling **saptune**

To disable **saptune** and to stop and disable `tuned` run:

```
root # saptune daemon stop
```

7.4.10   Tuning Kernel Parameters Manually Using **sysctl**

In addition to or instead of tuning kernel parameters using **saptune**, you can also use **sysctl** to make manual adjustments to kernel parameters. However, such changes using **sysctl** do not persist across reboots by default. To make them persist across reboots, add them to the file `/etc/sysctl.conf` (or another configuration file read by **sysctl**).

> 💡 **Tip**
>
> If you plan to configure sysctl parameters for your SAP system, consider to use **saptune** as a central instance for managing your system configuration.

For more information about **sysctl**, see the man pages `sysctl(8)`, `sysctl.conf(5)`, and `sysctl.d(5)`.

7.4.11   For More Information

See the following man pages:

- **man 8 saptune**

- **man 8 saptune_v1**

- **man 8 saptune_v2**

- `man 8 saptune-migrate`

- `man 8 saptune-note`

Also see the project home page https://github.com/SUSE/saptune/ ↗.

7.5 Tuning Kernel Parameters Manually Using `sysctl`

In addition to or instead of tuning kernel parameters using `sapconf`/**saptune**, you can also use `sysctl` to make manual adjustments to kernel parameters. However, such changes using `sysctl` do not persist across reboots by default. To make them persist across reboots, add them to one of the configuration files read by `sysctl`).

For more information about `sysctl`, see the man pages `sysctl(8)`, `sysctl.conf(5)`, and `sysctl.d(5)`.

8 Firewalling

This chapter presents information about restricting access to the system using firewalling and encryption and gives information about connecting to the system remotely.

8.1 Configuring SuSEfirewall2

By default, the installation workflow of SUSE Linux Enterprise Server for SAP Applications enables SuSEfirewall2. The firewall needs to be manually configured to allow network access for the following:

- SAP application

- Database (see the documentation of your database vendor; for SAP HANA, see *Section 8.2, "Configuring HANA-Firewall"*)

Additionally, open the ports `1128` (TCP) and `1129` (UDP).

SAP applications require many open ports and port ranges in the firewall. The exact numbers depend on the selected instance. For more information, see the documentation provided to you by SAP.

8.2 Configuring HANA-Firewall

To simplify setting up a firewall for SAP HANA, install the package `HANA-Firewall`. HANA-Firewall adds rule sets to your existing SuSEfirewall2 configuration.

HANA-Firewall consists of the following parts:

- **YaST Module** *SAP HANA Firewall*.   Allows configuring, applying, and reverting firewall rules for SAP HANA from a graphical user interface.

- **Command Line Utility** `hana-firewall`.   Allows applying and reverting the configured firewall rules for SAP HANA.
  If you prefer, you can configure the rule sets using the configuration file at `/etc/sysconfig/hana-firewall` instead of using YaST.

- **Service** `hana-firewall`.   Ensures that configured firewall rules for SAP HANA are kept.

⚠ Important: SAP HANA MDC Databases

For multi-tenant SAP HANA (MDC) databases, determining the port numbers that need to be opened is not yet possible automatically. If you are working with a multi-tenant SAP HANA database system: Before you use YaST, run a script on the command line to create a new service definition:

```
root # cd /etc/hana-firewall.d
root # ./create_new_service
```

You need to switch to the directory `/etc/hana-firewall.d`, otherwise the rule file for the new service will be created in a place where it cannot be used.

The script will ask several questions: Importantly, it will ask for TCP and UDP port ranges that need to be opened.

✎ Note: Install HANA-Firewall Packages

Before continuing, make sure that the packages `HANA-Firewall` and `yast2-hana-firewall` are installed.

PROCEDURE 8.1: USING HANA-FIREWALL

1. Make sure the SAP HANA databases for which you want to configure the firewall are correctly installed.

2. To open the appropriate YaST module, select *Applications › YaST, Security and Users › SAP HANA Firewall*.

3. When you open this YaST module, it will create a configuration proposal based on the number of installed SAP HANA instances.
   Choose whether you want to accept the proposal using *Yes* or *No*.

   # ⚠ Important: Narrow Down Settings from Proposal

   The proposed settings allow all detected SAP HANA instances on all detected network interfaces. Narrow down the proposal to secure the system further.

4. Under *Global Options*, activate *Enable Firewall*. Additionally, decide whether to *Allow Remote Shell Access (SSH)*.

5. Choose a network interface under *Allowed Services on Network Interface*.

6. Allow network services by selecting them in the list box on the left and clicking →. Remove services by selecting them in the list box on the right and clicking ←.
   To add services other than the preconfigured ones, add them using the following notation:

   ```
   SERVICE_NAME:CIDR_NOTATION
   ```

   For more information about the CIDR notation, see https://en.wikipedia.org/wiki/Class-less_Inter-Domain_Routing ↗. To find out which services are available on your system, use **getent services**.
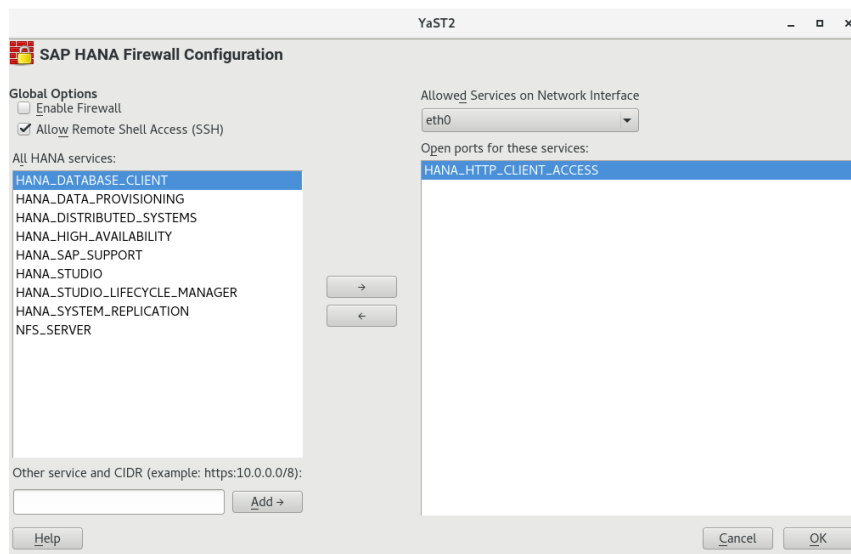
7. Repeat from *Step 5* for all network interfaces.

8. When you are done, click *OK*.
   The firewall rules from HANA-Firewall will now be compiled and applied. Then, the service `hana-firewall` will be restarted.

9. Finally, check whether HANA-Firewall was enabled correctly:

   ```
   root # hana-firewall status
   HANA firewall is active. Everything is OK.
   ```

> **Tip: Checking Which Firewall Rules Are Enabled**
>
> Gaining an overview of which firewall rules are enabled in the current configuration of the script is possible using the command line:
>
> ```
> root # hana-firewall dry-run
> ```

For more information, see the man page of `hana-firewall`.

8.3 SAProuter Integration

The SAProuter software from SAP allows proxying network traffic between different SAP systems or between an SAP system and outside networks. SUSE Linux Enterprise Server for SAP Applications now provides integration for SAProuter into `systemd`. This means, SAProuter will be started and stopped properly with the operating system and can be controlled using `systemctl`.

Before you can use this functionality, make sure the following has been installed, in this order:

- An SAP application that includes SAProuter

- The SAProuter systemd integration, packaged as `systemd-saprouter`

If you got the order of applications to install wrong initially, reinstall `systemd-saprouter`.

To control SAProuter with `systemctl`, use:

- Enabling the SAProuter Service: `systemctl enable saprouter`

- Starting the SAProuter Service: `systemctl start saprouter`

- Showing the Status of SAProuter Service: `systemctl status saprouter`

- Stopping the SAProuter Service: `systemctl stop saprouter`

- Disabling the SAProuter Service: `systemctl disable saprouter`

9 Encrypting Directories Using `cryptctl`

`cryptctl` consists of two components:

- A client is a machine that has one or more encrypted partitions but does not permanently store the necessary key to decrypt those partitions. For example, clients can be cloud or otherwise hosted machines.

- The server holds encryption keys that can be requested by clients to unlock encrypted partitions.

  You can also set up the `cryptctl` server to store encryption keys on a KMIP 1.3-compatible (Key Management Interoperability Protocol) server. In that case, the `cryptctl` server will not store the encryption keys of clients and is dependent upon the KMIP-compatible server to provide these.

> ✋ **Warning: `cryptctl` Server Maintenance**
>
> The `cryptctl` server manages timeouts for the encrypted disks. Depending on the configuration, it can also hold encryption keys. Therefore it should be under your direct control and managed only by trusted personnel.
>
> Additionally, it should be backed up regularly. Losing the server's data means losing access to encrypted partitions on the clients.

To handle encryption, `cryptctl` uses LUKS with aes-xts-256 encryption and 512-bit keys. Encryption keys are transferred using TLS with certificate verification.
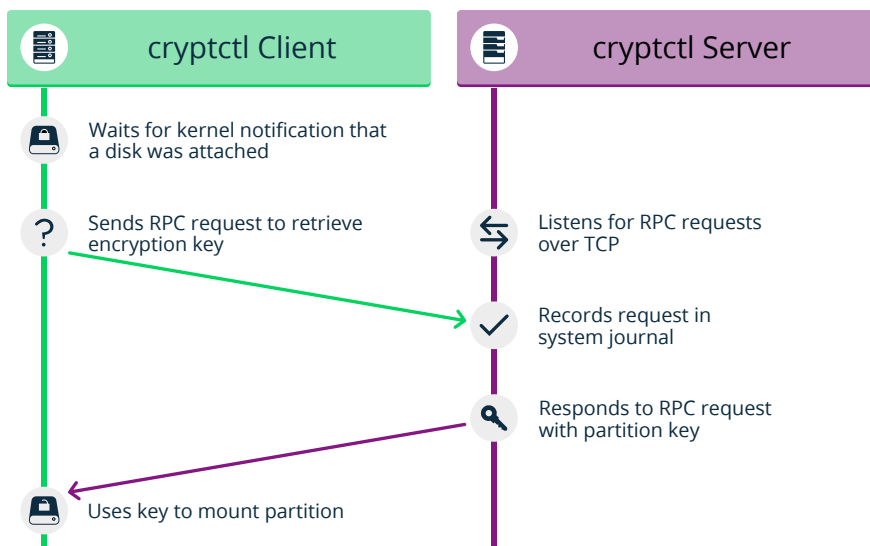
FIGURE 9.1: KEY RETRIEVAL WITH `cryptctl` (MODEL WITHOUT CONNECTION TO KMIP SERVER)

> ## Note: Install `cryptctl`
>
> Before continuing, make sure the package `cryptctl` is installed on all machines you intend to set up as servers or clients.

9.1   Setting Up a **cryptctl** Server

Before you can define machine as a `cryptctl` client, you need to set up a machine as a `cryptctl` server.

Before beginning, choose whether to use a self-signed certificate to secure communication between the server and clients. If not, generate a TLS certificate for the server and have it signed by a certificate authority.

Additionally, you can have clients authenticate to the server using certificates signed by a certificate authority. To use this extra security measure, make sure to have a CA certificate at hand before starting this procedure.

1. As `root`, run:

   ```
   root # cryptctl init-server
   ```

2. Answer each of the following prompts and press `Enter` after every answer. If there is a default answer, it is shown in square brackets at the end of the prompt.

a. Choose a password with at least 10 characters and confirm it. This password assumes the role of a master password, able to unlock all partitions that are registered on the server.

b. Specify the path to a PEM-encoded TLS certificate or certificate chain file or leave the field empty to create a self-signed certificate. If you specify a path, use an absolute path.

c. If you want the server to be identified by a host name other than the default shown, specify a host name. `cryptctl` will then generate certificates which include the host name.

d. Specify the IP address that belongs to the network interface that you want to listen on for decryption requests from the clients, then set a port number (the default is port 3737).
   The default IP address setting, `0.0.0.0` means that `cryptctl` will listen on all network interfaces for client requests using IPv4.

e. Specify a directory on the server that will hold the decryption keys for clients.

f. Specify whether clients need to authenticate to the server using a TLS certificate. If you choose *No*, this means that clients authenticate using disk UUIDs only. (However, communication will be encrypted using the server certificate in any case.)
   If you choose *Yes*, pick a PEM-encoded certificate authority to use for signing client certificates.

g. Specify whether to use a KMIP 1.3-compatible server (or multiple such servers) to store encryption keys of clients. If you choose this option, provide the host names and ports for one or multiple KMIP-compatible servers.
   Additionally, provide a user name, password, a CA certificate for the KMIP server, and a client identity certificate for the `cryptctl` server.

> ⚠️ **Important: No Easy Reconfiguration of KMIP Setting**
>
> The setting to use a KMIP server cannot easily be changed later. To change this setting, both the `cryptctl` server and its clients need to be configured afresh.

h. Finally, configure an SMTP server for e-mail notifications for encryption and decryption requests or leave the prompt empty to skip setting up e-mail notifications.

> **ⓘ Note: Password-Protected Servers**
>
> `cryptctl` currently cannot send e-mail using authentication-protected SMTP servers. If that is necessary, set up a local SMTP proxy.

    i. When asked whether to start the `cryptctl` server, enter `y`.

3. To check the status of the service `cryptctl-server`, use:

```
root # systemctl status cryptctl-server
```

To reconfigure the server later, do either of the following:

- Run the command `cryptctl init-server` again. `cryptctl` will then propose the existing settings as the defaults, so that you only need to the specify values that you want to change.

- Make changes directly in the configuration file `/etc/sysconfig/cryptctl-server`. However, to avoid issues, do not change the settings `AUTH_PASSWORD_HASH` and `AUTH_PASSWORD_SALT` manually. The values of these options need to be calculated correctly.

9.2   Setting Up a **cryptctl** Client

The following interactive setup of `cryptctl` is currently the only setup method.

Make sure the following preconditions are fulfilled:

- A `cryptctl` server is available over the network.

- There is a directory to encrypt.

- The client machine has an empty partition available that is large enough to fit the directory to encrypt.

- When using a self-signed certificate, the certificate (`*.crt` file) generated on the server is available locally on the client. Otherwise, the certificate authority of the server certificate must be trusted by the client.

- If you set up the server to require clients to authenticate using a client certificate, prepare a TLS certificate for the client. It must be signed by the CA certificate you chose for the server.

1. As `root`, run:

```
root # cryptctl encrypt
```

2. Answer each of the following prompts and press Enter after every answer. If there is a default answer, it is shown in square brackets at the end of the prompt.

   a. Specify the host name and port to connect to on the **cryptctl** server.

   b. If you configured the server to have clients authenticate to it using a TLS certificate, specify a certificate and a key file for the client. The client certificate must be signed by the certificate authority chosen when setting up the server.

   c. Specify the absolute path to the server certificate (the `*.crt` file).

   d. Enter the encryption password that you specified when setting up the server.

   e. Specify the path to the directory to encrypt. Specify the path to the empty partition that will contain the encrypted content of the directory.

   f. Specify the number of machines that are allowed to decrypt the partition simultaneously.
   Then specify the timeout in seconds before additional machines are allowed to decrypt the partition after the last vital sign was received from the client or clients. When a machine unexpectedly stops working and then reboots, it needs to be able to unlock its partitions again. That means, this timeout should be set to a time slightly shorter than the reboot time of the client.

   > **! Important: Timeout Length**
   >
   > If the time is set too long, the machine cannot decrypt encrypted partitions on the first try. **cryptctl** will then continue to periodically check whether the encryption key has become available. However, this will introduce a delay.

If the timeout is set too short, machines with a copy of the encrypted partition have an increased chance of unlocking the partition first.

3. To start encryption, enter `yes`.

   `cryptctl` will now encrypt the specified directory to the previously empty partition and then mount the newly encrypted partition. The file system type will be of the same type as the original unencrypted file system.

   Before creating the encrypted partition, `cryptctl` moves the unencrypted content of the original directory to a location prefixed with `cryptctl-moved-`.

4. To check that the directory is indeed mounted correctly, use:

```
tux > lsblk -o NAME,MOUNTPOINT,UUID
NAME                            MOUNTPOINT          UUID
[...]
sdc
└─sdc1                                              PARTITION_UUID
  └─cryptctl-unlocked-sdc1  /secret-partition   UNLOCKED_UUID
```

   `cryptctl` identifies the encrypted partition by its UUID. For the previous example, that is the UUID displayed next to `sdc1`.

   On the server, you can check whether the directory was decrypted using `cryptctl`:

```
root # cryptctl list-keys
2016/10/10 10:00:00 ReloadDB: successfully loaded database of 1 records
Total: 1 records (date and time are in zone EDT)
Used By      When                  UUID  Max.Users  Num.Users  Mount Point
IP_ADDRESS   2016-10-10 10:00:00   UUID  1          1          /secret-partition
```

   Verify that the UUID shown is that of the previously encrypted partition.

5. After verifying that the encrypted partition works, delete the unencrypted content from the client. For example, use `rm`. For more safety, overwrite the content of the files before deleting them, for example, using `shred -u`.

   > ❗ **Important: `shred` Does Not Guarantee That Data Is Completely Erased**
   >
   > Depending on the type of storage media, using `shred` is not a guarantee that all data is completely removed. In particular, SSDs usually employ wear leveling strategies that render `shred` ineffective.

The configuration for the connection from client to server is stored in `/etc/syscon-fig/cryptctl-client` and can be edited manually.

The server stores an encryption key for the client partition in `/var/lib/cryptctl/key-db/PARTITION_UUID`.

9.3 Checking Partition Unlock Status Using Server-side Commands

When a `cryptctl` client is active, it will send a "heartbeat" to the `cryptctl` server every 10 seconds. If the server does not receive a heartbeat from the client for the length of the timeout configured during the client setup, the server will assume that the client is offline. It will then allow another client to connect (or allow the same client to reconnect after a reboot).

To see the usage status of all keys, use:

```
root # cryptctl list-keys
```

The information under `Num. Users` shows whether the key is currently in use. To see more detail on a single key, use:

```
root # cryptctl show-key UUID
```

This command will show information about mount point, mount options, usage options, the last retrieval of the key and the last three heartbeats from clients.

Additionally, you can use `journalctl` to find logs of when keys were retrieved.

9.4 Unlocking Encrypted Partitions Manually

There are two ways of unlocking a partition manually, both of which are run on a client:

- **Online Unlocking.** Online unlocking allows circumventing timeout or user limitations. This method can be used when there is a network connection between client and server but the client could not (yet) unlock the partition automatically. This method will unlock all encrypted partitions on a machine.

To use it, run **`cryptctl online-unlock`**. Be prepared to enter the password specified when setting up the server.

- **Offline Unlocking.** This method can be used when a client cannot or must not be brought online to communicate with its server. The encryption key from the server must still be available. This method is meant as a last resort only and can only unlock a single partition at a time.

    To use it, run **`cryptctl offline-unlock`**. The server's key file for the requisite partition (`/var/lib/cryptctl/keydb/`*`PARTITION_UUID`*) needs to be available on the client.

9.5 Maintenance Downtime Procedure

To ensure that partitions cannot be decrypted during a maintenance downtime, turn off the client and disable the **`cryptctl`** server. You can do so by either:

- Stopping the service `cryptctl-server`:

    ```
    root # systemctl stop cryptctl-server
    ```

- Unplugging the **`cryptctl`** server from the network.

9.6 For More Information

For more information, also see the project home page https://github.com/HouzuoGuo/cryptctl/ ↗.

10 Protecting Against Malware With ClamSAP

ClamSAP integrates the ClamAV anti-malware toolkit into SAP NetWeaver and SAP Mobile Platform applications. ClamSAP is a shared library that links between ClamAV and the SAP NetWeaver Virus Scan Interface (NW-VSI). The version of ClamSAP shipped with SUSE Linux Enterprise Server for SAP Applications 12 SP5 supports NW-VSI version 2.0.

10.1 Installing ClamSAP

1. On the application host, install the packages for ClamAV and ClamSAP. To do so, use the command:

   ```
   tux > sudo zypper install clamav clamsap
   ```

2. Before you can enable the daemon `clamd`, initialize the malware database:

   ```
   tux > sudo freshclam
   ```

3. Start the service `clamd`:

   ```
   tux > sudo systemctl start clamd
   ```

4. Check the status of the service `clamd` with:

   ```
   tux > systemctl status clamd
   ● clamd.service - ClamAV Antivirus Daemon
   Loaded: loaded (/usr/lib/systemd/system/clamd.service; enabled; vendor preset:
    disabled)
   Active: active (running) since Tue 2017-04-11 10:33:03 UTC; 24h ago
   [...]
   ```

10.2 Creating a Virus Scanner Group in SAP NetWeaver

1. Log in to the SAP NetWeaver installation through the GUI. Do not log in as a `DDIC` or `SAP*` user, because the virus scanner needs to be configured cross-client.

2. Create a Virus Scanner Group using the transaction *VSCANGROUP*.
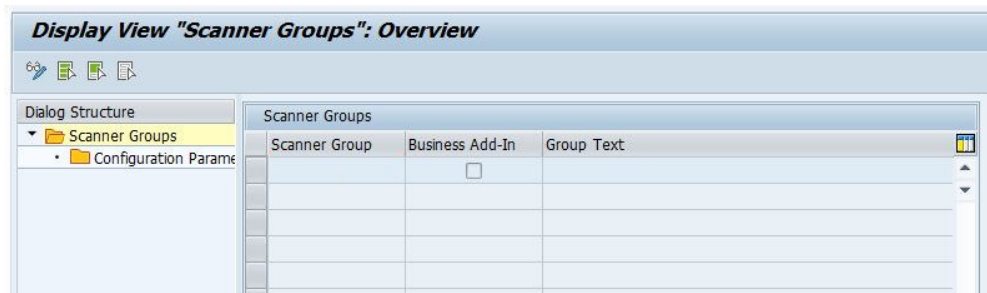
3. To switch from view mode to change mode, click the button *Change View* (🔧).
   Confirm the message *This table is cross-client* by clicking the check mark. The table is now editable.

4. Select the first empty row. In the text box *Scanner Group*, specify `CLAMSAPVSI`. Under *Group Text*, specify `CLAMSAP`.
   Make sure that *Business Add-in* is not checked.
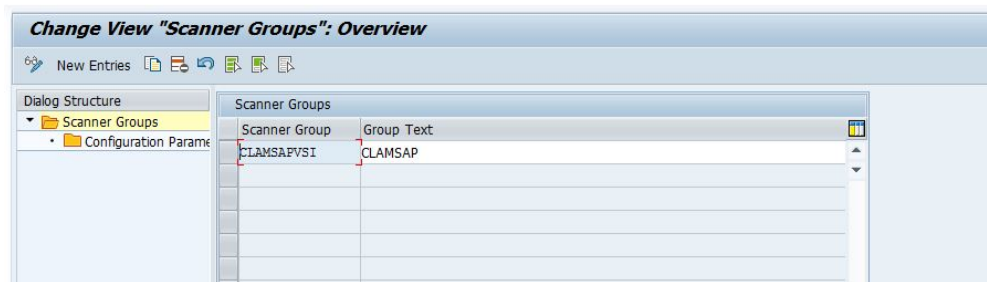
5. To save the form, click the button *Save* (💾).

10.3   Setting Up the ClamSAP Library in SAP NetWeaver

1. In the SAP NetWeaver GUI, call the transaction *VSCAN*.

2. To switch from view mode to change mode, click the button *Change View* (🔧).
   Confirm the message *This table is cross-client* by clicking the check mark. The table is now editable.

3. Click *New entries*.

4. Fill in the form accordingly:

- *Provider Type*: `Adapter (Virus Scan Adapter)`

- *Provider Name*: `VSA_HOSTNAME` (for example: `VSA_SAPSERVER`)

- `Scanner Group`: The name of the scanner group that you set up in *Section 10.2, "Creating a Virus Scanner Group in SAP NetWeaver"* (for example: `CLAMSAPVSI`)

- *Server*: `HOSTNAME_SID_INSTANCE_NUMBER` (for example: `SAPSERVER_P04_00`)

- *Adapter Path*: `libclamsap.so`

5. To save the form, click the button 🖫.

10.4 Engaging ClamSAP

To run ClamSAP, go to the transaction *VSCAN*. Then click *Start*.
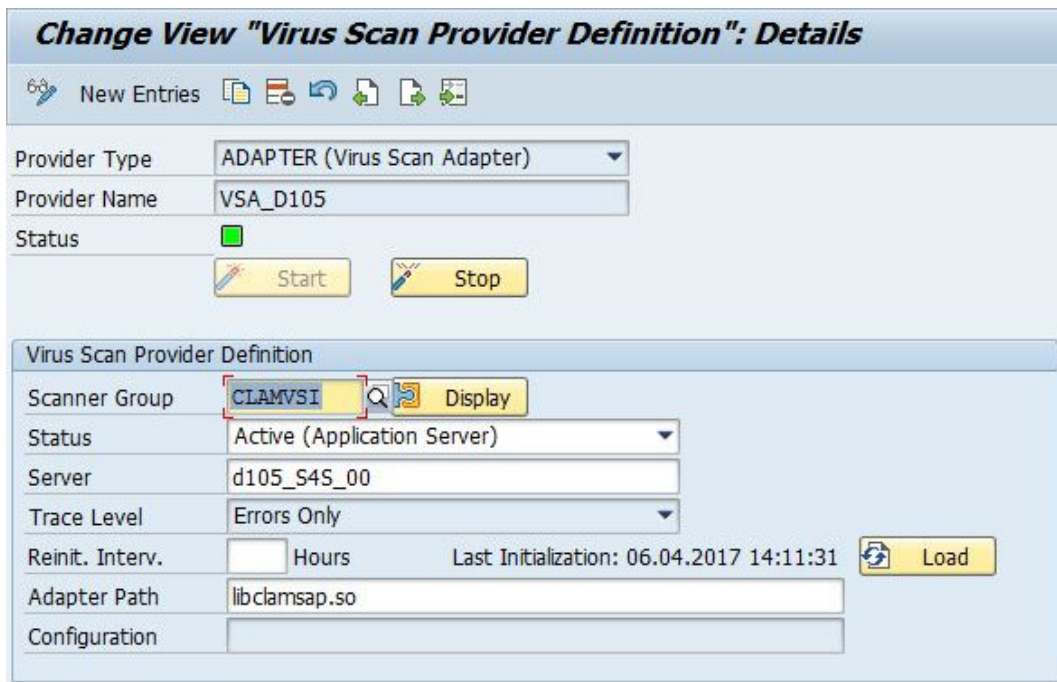
FIGURE 10.1: **CHANGE VIEW "VIRUS SCAN PROVIDER DEFINITION"**

Afterward, a summary will be displayed, including details of the ClamSAP and ClamAV (shown in *Figure 10.2, "Summary of ClamSAP Data"*).

FIGURE 10.2: **SUMMARY OF CLAMSAP DATA**

10.5 For More Information

For more information, also see the project home page https://sourceforge.net/projects/clam-sap/ .

11 Connecting via RDP

If you installed SLES-SAP with the RDP option activated or if you installed from a KIWI image, RDP is enabled on the machine via the service `xrdp`. Alternatively, you can enable RDP later as described at the end of this section.

You can connect using any software that supports RDP, such as:

- **Linux:** Vinagre (available in SUSE Linux Enterprise Desktop/SLE Workstation Extension and openSUSE) or Remmina (available in openSUSE)

- **Windows:** Remote Desktop Connection

> **❗ Important: Connection Parameters**
>
> Make sure to set up the connection with the following parameters:
>
> - **Port:** 3389
>
> - **Color Depth:** 16-bit or 24-bit only

PROCEDURE 11.1: **SETTING UP RDP**

> If you have not set up an RDP connection during the installation, you can also do so later using the following instructions.
>
> 1. First, create an exception for the firewall. Start by creating a file that sets up the port that needs to be opened for RDP.
>   As `root`, create a new file under `/etc/sysconfig/SuSEfirewall2.d/services/` with the name `xrdp` and the following content:
>
>    ```
>    ## Name: Remote Desktop Protocol
>    TCP="3389"
>    ```
>
> 2. Open the file `/etc/sysconfig/SuSEfirewall2` and change the lines for the settings `FW_CONFIGURATIONS_EXT`, `FW_CONFIGURATIONS_DMZ`, and `FW_CONFIGURATIONS_INT` to include `xrdp`. If there are no other services enabled, the respective lines should read:
>
>    ```
>    FW_CONFIGURATIONS_EXT="xrdp"
>    FW_CONFIGURATIONS_DMZ="xrdp"
>    FW_CONFIGURATIONS_INT="xrdp"
>    ```

If there are other services, separate them within the quotes using a space character.

3. Now set up `xrdp` itself.

   If the package `xrdp` is not installed, install it:

   ```
   root # zypper install xrdp
   ```

4. Enable and start the service:

   ```
   root # systemctl restart SuSEfirewall2
   ```

5. Enable and start the service:

   ```
   root # systemctl enable xrdp
   root # systemctl start xrdp
   ```

   You can now connect to the machine.

12 Creating Operating System Images

There are multiple ways to create custom operating system images from SUSE Linux Enterprise Server for SAP Applications. The preferred way is generally to use KIWI which ingests an XML configuration file and then runs fully automatically.

Alternatively, you can also create an image from an existing installation that is cleaned up before re-use.

12.1 Creating Images with KIWI

KIWI is a tool to create operating system images that can be easily copied to new physical or virtual machines. This section will present information on creating SLES-SAP images with KIWI.

SUSE Linux Enterprise Server for SAP Applications now supports creating images with KIWI using the template from the package `kiwi-template-sap`. However, there are certain restrictions in the current implementation:

- Only building VMX disk images is supported. Building other image types is not supported.

- You must provide an ISO image of SUSE Linux Enterprise Server for SAP Applications at `/tmp/SLES4SAP.iso`. The Open Build Service does not contain all necessary packages.

To build a basic image, use the following two commands:

1. Build the root file system:

   ```
   root # kiwi -p SLES4SAP --root fsroot
   ```

2. Build the VMX image:

   ```
   root # kiwi --create fsroot --type vmx -d build
   ```

To enable running graphical installations using SAPinst, the default settings of the image enable the following:

- Installation of an IceWM desktop

- The service `xrdp` is started automatically, so you can connect to the machine via RDP. For more information, see *Chapter 11, Connecting via RDP*.

For more information about KIWI and SLES-SAP:

- On the KIWI configuration for SLES-SAP, see `/usr/share/kiwi/image/SLES4SAP/README`.

- On KIWI in general, see the *openSUSE-KIWI Image System Cookbook* (https://doc.open-suse.org/projects/kiwi/doc/ ↗).

12.2   Cleaning Up an Instance Before Using It as a Master Image

In some cases, it makes sense to use an image of an already-configured master instance on multiple systems instead of generating a KIWI image from scratch. For example, when your image needs to contain additional software or configuration that cannot be installed using KIWI.

However, normally such an image would contain certain configuration data that should not be copied along with the rest of the system.

To avoid needing to clean up manually, use the script **clone-master-clean-up** (available from the package of the same name).

It deletes the following data automatically:

- Swap device (zero-wiped, then re-enabled)

- SUSE registration information and repositories from SUSE, and the Zypper ID

- User and host SSH keys and domain and host names

- The generated `HANA-Firewall` script (but not the configuration itself)

- Shell history, mails, cron jobs, temporary files (`/tmp`, `/var/tmp`), log files (`/var/log`), random seeds, `systemd` Journal, `collectd` statistics, **postfix** configuration, parts of `/root`

- `/var/cache`, `/var/crash`, `/var/lib/systemd/coredump`

Additionally, the following configuration is restored to defaults:

- Network interfaces that do not use DHCP and network configuration (`/etc/hostname`, `/etc/hosts`, and `/etc/resolv.conf`)

- `sudo` settings

Additionally, you can choose to set up a new `root` password. UUID-based entries in `/etc/fstab` are replaced by device strings. This script also ensures that, if the first-boot section of the installation workflow was used for the original installation, it is run again on the next boot.

12.2.1 Configuring `clone-master-clean-up`

Before running `clone-master-clean-up`, the script can be configured in the following ways:

- To configure the script to not clean up certain data, use the configuration file `/etc/sysconfig/clone-master-clean-up`.
  This file also gives short explanations of the available options.

- To configure the script to clean up additional directories or files, create a list with the absolute paths of such directories and files:

  ```
  /additional/file/to/delete.now
  /additional/directory/to/remove
  ```

  Save this list as `/var/adm/clone-master-clean-up/custom_remove`.

12.2.2 Using `clone-master-clean-up`

To use the script, do:

```
root # clone-master-clean-up
```

Then follow the instructions.

12.2.3 For More Information

The following sources provide additional information about `clone-master-clean-up`:

- For general information, see the man page `clone-master-clean-up`.

- For information on which files and directories might additionally be useful to delete, see `/var/adm/clone-master-clean-up/custom_remove.template`.

13 Important Log Files

The most important files for this product are:

- The SAP Installation Wizard is a YaST module. You can find its log entries in `/var/log/YaST/y2log`.

- All SAP knowledge is bundled in a library. You can find its log entries in `/var/log/SAPmedia.log`.

- You can find log files related to auto-installation in `/var/adm/autoinstall/logs`.

A   Additional Software for SLES-SAP

SUSE Linux Enterprise Server for SAP Applications makes it easy to install software that is not included with your subscription:

- Extensions and modules allow installing additional software created and supported by SUSE. For more information about extensions and modules, see *Deployment Guide, Part "Initial System Configuration", Chapter "Installing Modules, Extensions, and Third Party Add-On Products"* at https://documentation.suse.com/sles-12 ↗.

- *SUSE Connect Program* allows installing packages created and supported by third parties, specifically for SLES-SAP. It also gives easy access to third-party trainings and support. See *Section A.2, "SUSE Connect Program"*.

- SUSE Package Hub allows installation of packages created by SUSE Linux Enterprise community without support. See *Section A.3, "PackageHub"*.

A.1   Identifying a Base Product for SUSE Linux Enterprise Server for SAP Applications

To identify and distinguish SUSE products, use one of the following files:

`/etc/os-release`

> A text file with key-value pairs, similar to shell-compatible variable assignments. Each key is on a separate line.
>
> You can search for the `CPE_NAME` key; however, between different releases and service packs, the value have been changed. If you need further details, refer to the article at https://www.suse.com/support/kb/doc/?id=7023490 ↗.

`/etc/product.d/baseproduct`

> A link to an XML file. The `/etc/product.d/` directory contains different `.prod` files. Depending on which products you have purchased and how you installed your system, the link `/etc/product.d/baseproduct` can point to a different `.prod` file, for example, `sle-module-sap-applications.prod`. The same information as `CPE_NAME` is stored in the tag `<cpeid>`.

Among other information, both files contain the operating system and base product. The base product (key `CPE_NAME` and tag `<cpeid>`) follow the Common Platform Enumeration Specification (http://scap.nist.gov/specifications/cpe/) ↗.

Basically, you can extract any information from the file `/etc/product.d/baseproduct` either with the commands **grep** or **xmlstarlet** (both are available for your products). As XML is also text, use **grep** for "simple searches" when the format of the output does not matter much. However, if your search is more advanced, you need the output in another script, or your would like to avoid the XML tags in the output, use the **xmlstarlet** command instead.

For example, to get your base product, use **grep** like this:

```
tux > grep cpeid /etc/products.d/baseproduct
<cpeid>cpe:/o:suse:sle-module-sap-applications:RELEASE:spSP_NUMBER</cpeid>
```

The `RELEASE` and `SP_NUMBER` are placeholders and describe your product release number and service pack.

The same can be achieved with **xmlstarlet**. You need an XPath (the steps that lead you to your information). With the appropriate options you can avoid the `<cpeid>`/`</cpeid>` tags:

```
tux > xmlstarlet sel -T -t -v "/product/cpeid" /etc/products.d/baseproduct
cpe:/o:suse:sle-module-sap-applications:RELEASE:spSP_NUMBER
```

A more advanced search (which would be difficult for **grep**) would be to list all required dependencies to other products. Assuming, `basename` points to `sle-module-sap-applications.prod`, the following command will output all product dependencies which are required for SUSE Linux Enterprise Server for SAP Applications:

```
>tux > xmlstarlet sel -T -t -v "/product/productdependency[@relationship='requires']/
@name" /etc/products.d/baseproduct
SUSE_SLE
sle-ha
```

A.2 SUSE Connect Program

Start SUSE Connect Program from the YaST control center using *SUSE Connect Program.* Choose from the available options. To enable a software repository, click *Add repository*.

All software enabled by SUSE Connect Program originates from third parties. For support, contact the vendor in question. SUSE does not provide support for these offerings.

> ⬢ **Note: SUSEConnect command-line tool**
>
> The `SUSEConnect` command-line tool is a separate tool with a different purpose: It allows you to register installations of SUSE products.

A.3 PackageHub

PackageHub provides many packages for SLE that were previously only available on openSUSE. Packages from SUSE Package Hub are created by the community and come without support. The selection includes, for example:

- The R programming language

- The Haskell programming language

- The KDE 5 desktop

To enable PackageHub, add the repository as described at https://packagehub.suse.com/how-to-use/ ↗ .

For more information, see the PackageHub Web site at https://packagehub.suse.com ↗ .

B Partitioning for the SAP System Using AutoYaST

Partitioning for the SAP system is controlled by the files from the directory `/usr/share/YaST2/include/sap-installation-wizard/`. The following files can be used:

- **SAP NetWeaver or SAP S/4HANA Application Server Installation.** `base_partitioning.xml`

- **SAP HANA or SAP S/4HANA Database Server Installation.** `hana_partitioning.xml`

- **SAP HANA or SAP S/4HANA Database Server Installation on SAP BusinessOne-Certified Hardware.** hardware-specific partitioning file

The files will be chosen as defined in `/etc/sap-installation-wizard.xml`. Here, the content of the element `partitioning` is decisive.

If the installation is, for example, based on HA or a distributed database, no partitioning is needed. In this case, `partitioning` is set to `NO` and the file `base_partitioning.xml` is used.

> 🖊 Note: `autoinst.xml` Cannot Be Used Here
>
> `autoinst.xml` is only used for the installation of the operating system. It cannot control the partitioning for the SAP system.

The files that control partitioning are AutoYaST control files that contain a `partitioning` section only. However, these files allow using several extensions to the AutoYaST format:

- If the `partitioning_defined` tag is set to `true`, the partitioning will be performed without any user interaction.
  By default, this is only used when creating SAP HANA file systems on systems certified for SAP HANA (such as from Dell, Fujitsu, HP, IBM, or Lenovo).

- For every partition, you can specify the `size_min` tag. The size value can be given as a string in the format of *RAM*N*. This way you can specify how large the partition should minimally be (*N* times the size of the available memory (*RAM*)).

PROCEDURE B.1: **CREATING A CUSTOM SAP PARTITIONING SETUP**

The steps below illustrates how to create a partitioning setup for TREX. However, creating a partitioning setup for other applications works analogously.

1. In `/usr/share/YaST2/include/sap-installation-wizard/`, create a new XML file. Name it `TREX_partitioning.xml`, for example.

2. Copy the content of `base_partitioning.xml` to your new file and adapt the new file to your needs.

3. Finally, adapt `/etc/sap-installation-wizard.xml` to include your custom file. In the `listitem` for `TREX`, insert the following line:

```
<partitioning>TREX_partitioning</partitioning>
```

> ⚠️ **Important: Do Not Edit `base_partitioning.xml`**
>
> Do not edit `base_partitioning.xml` directly. With the next update, this file will be overwritten.

For more information about partitioning with AutoYaST, see *AutoYaST Guide, Chapter "Partitioning"* (https://documentation.suse.com/sles-12 ↗).

C Supplementary Media

Supplementary Media allow partners or customers to add their own tasks or workflows to the Installation Wizard.

This is done by adding an XML file which will be part of an AutoYaST XML file. To be included in the workflow, this file must be called `product.xml`.

This can be used for various types of additions, such as adding your own RPMs, running your own scripts, setting up a cluster file system or creating your own dialogs and scripts.

C.1 `product.xml`

The `product.xml` file looks like a normal AutoYaST XML file, but with some restrictions.

The restrictions exist because only the parts of the XML that are related to the second stage of the installation are run, as the first stage was executed before.

Both XML files (`autoyast.xml` and `product.xml`) will be merged after the media is read. A "new" AutoYaST XML file is generated on the fly for the additional workflow.

The following areas or sections will be merged:

```
<general>
  <ask-list>          ❶
  ...
<software>            ❷
  <post-packages>
  ...
<scripts>
  <chroot-scripts>    ❸
  <post-scripts>      ❹
  <init-scripts>      ❺
  ...
```

❶   see *Section C.2, "Own AutoYaST Ask Dialogs"*

❷   see *Section C.3, "Installing Additional Packages"*

❸   after the package installation, before the first boot

❹   during the first boot of the installed system, no services running

❺   during the first boot of the installed system, all services up and running

All other sections will be replaced.

For more information about customization options, see *AutoYaST Guide, Chapter "Configuration and Installation Options", Section "Custom User Scripts"* (https://documentation.suse.com/sles-12 ↗).

C.2 Own AutoYaST Ask Dialogs

For more information about the "Ask" feature of AutoYaST, see *AutoYaST Guide, Chapter "Configuration and Installation Options", Section "Ask the User for Values During Installation"* (https://documentation.suse.com/sles-12 ↗).

For the Supplementary Media, you can only use dialogs within the `cont` stage (`<stage>cont</stage>`), which means they are executed after the first reboot.

Your file with the dialogs will be merged with the base AutoYaST XML file.

As a best practice, your dialog should have a dialog number and an element number, best with steps of 10. This helps to include later additions and could be used as targets for jumping over a dialog or element dependent on decisions. We also use this in our base dialogs. If you provide the right dialog number and element number, you can place your dialog between our base dialogs.

You can store the answer to a question in a file, to use it in one of your scripts later. Be aware that you *must* use the prefix `/tmp/ay` for this: The Installation Wizard will copy such files from the `/tmp` directory to the directory where your media data also will be copied. This is done because the next Supplementary Media could have the same dialogs or same answer file names and would overwrite the values saved here.

Here is an example with several options:

```xml
<?xml version="1.0"?>
<!DOCTYPE profile>
<profile xmlns="http://www.suse.com/1.0/yast2ns"
         xmlns:config="http://www.suse.com/1.0/configns">
<general>
  <ask-list config:type="list">
     <ask>
         <stage>cont</stage>
         <dialog config:type="integer">20</dialog>
         <element config:type="integer">10</element>
         <question>What is your name?</question>
         <default>Enter your name here</default>
         <help>Please enter your full name within the field</help>
         <file>/tmp/ay_q_my_name</file>
         <script>
```

```
            <filename>my_name.sh</filename>
            <rerun_on_error config:type="boolean">true</rerun_on_error>
            <environment config:type="boolean">true</environment>
            <source><![CDATA[
function check_name() {
        local name=$1
        LC_ALL=POSIX
        [ -z "$name" ] && echo "You need to provide a name." && return 1
        return 0
}
check_name "$VAL"
]]>
            </source>
            <debug config:type="boolean">false</debug>
            <feedback config:type="boolean">true</feedback>
        </script>
      </ask>
   </ask-list>
</general>
</profile>
```

C.3 Installing Additional Packages

You can also install RPM packages within the `product.xml` file. To do this, you can use the `<post-packages>` element for installation in stage 2.

For more information, see *AutoYaST Guide, Chapter "Configuration and Installation Options",
Section "Installing Packages in Stage 2"* (https://documentation.suse.com/sles-12 ↗). An example
looks as follows:

```
...
<software>
 <post-packages config:type="list">
  <package>yast2-cim</package>
 </post-packages>
</software>
...
```

C.4 Example Directory for Supplementary Media

A minimal example for Supplementary Media directory contains only a file called `product.xml`.