SUSE

SUSE Linux Enterprise Server for SAP Applications 15 SP4

Guide

Guide

SUSE Linux Enterprise Server for SAP Applications 15 SP4

Publication Date: December 12, 2024

https://documentation.suse.com ↗

For SUSE trademarks, see https://www.suse.com/company/legal/↗. All third-party trademarks are the property of their respective owners. Trademark symbols (®, ™ etc.) denote trademarks of SUSE and its affiliates. Asterisks (*) denote third-party trademarks.

All information found in this book has been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. Neither SUSE LLC, its affiliates, the authors nor the translators shall be held liable for possible errors or the consequences thereof.

Contents

Preface

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

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

1 Available documentation

This manual contains links to additional documentation resources that are either available on the system or online.

Online documentation

Visit https://documentation.suse.com/#sles-sap ↗ for the latest version of this guide in different formats. You can find whitepapers and other resources in the SUSE Linux Enterprise Server for SAP Applications resource library: https://www.suse.com/products/sles-for-sap/resource-library/ ↗.

Find the online documentation for other products at https://documentation.suse.com/ ↗.

> Ⓢ **Note: Latest updates**
>
> The latest documentation updates are normally available in the English version of the documentation.

SUSE Knowledgebase

If you have run into an issue, also check out the Technical Information Documents (TIDs) that are available online at https://www.suse.com/support/kb/ ↗. Search the SUSE Knowledgebase for known solutions driven by customer need.

Release notes

For release notes, see https://www.suse.com/releasenotes/ ↗.

In your system

For offline use, the release notes are also available under `/usr/share/doc/release-notes` on your system. The documentation for individual packages is available at `/usr/share/doc/packages`.

Many commands are also described in their *manual pages*. To view them, run `man`, followed by a specific command name. If the `man` command is not installed on your system, install it with `sudo zypper install man`.

2 Improving the documentation

Your feedback and contributions to this documentation are welcome. The following channels for giving feedback are available:

Service requests and support

For services and support options available for your product, see https://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/ ↗.

To simplify this process, click the *Report an issue* icon next to a headline in the HTML version of this document. This preselects the right product and category in Bugzilla and adds a link to the current section. You can start typing your bug report right away.

A Bugzilla account is required.

Contributions

To contribute to this documentation, click the *Edit source document* icon next to a headline in the HTML version of this document. This will take you to the source code on GitHub, where you can open a pull request.

A GitHub account is required.

Note: *Edit source document* only available for English

The *Edit source document* icons are only available for the English version of each document. For all other languages, use the *Report an issue* icons instead.

For more information about the documentation environment used for this documentation, see the repository's README.

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.

3   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 a 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

- `AMD/Intel`⟩ This paragraph is only relevant for the AMD64/Intel 64 architectures. The arrows mark the beginning and the end of the text block. ◁
  `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. You can also prefix these commands with the **`sudo`** command to run them as a non-privileged user:

  ```
  # command
  > sudo command
  ```

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

  ```
  > command
  ```

- Commands can be split into two or multiple lines by a backslash character (`\`) at the end of a line. The backslash informs the shell that the command invocation will continue after the line's end:

```
> echo a b \
c d
```

- A code block that shows both the command (preceded by a prompt) and the respective output returned by the shell:

```
> command
output
```

- 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.

4 Support

Find the support statement for SUSE Linux Enterprise Server for SAP Applications and general information about technology previews below. For details about the product lifecycle, see https://www.suse.com/lifecycle ↗ .

If you are entitled to support, find details on how to collect information for a support ticket at https://documentation.suse.com/sles-15/html/SLES-all/cha-adm-support.html ↗ .

4.1 Support statement for SUSE Linux Enterprise Server for SAP Applications

To receive support, you need an appropriate subscription with SUSE. To view the specific support offers available to you, go to https://www.suse.com/support/ ↗ and select your product.

The support levels are defined as follows:

L1

Problem determination, which means technical support designed to provide compatibility information, usage support, ongoing maintenance, information gathering and basic troubleshooting using available documentation.

L2

Problem isolation, which means technical support designed to analyze data, reproduce customer problems, isolate a problem area and provide a resolution for problems not resolved by Level 1 or prepare for Level 3.

L3

Problem resolution, which means technical support designed to resolve problems by engaging engineering to resolve product defects which have been identified by Level 2 Support.

For contracted customers and partners, SUSE Linux Enterprise Server for SAP Applications is delivered with L3 support for all packages, except for the following:

- Technology previews.

- Sound, graphics, fonts, and artwork.

- Packages that require an additional customer contract.

- Some packages shipped as part of the module *Workstation Extension* are L2-supported only.

- Packages with names ending in `-devel` (containing header files and similar developer resources) will only be supported together with their main packages.

SUSE will only support the usage of original packages. That is, packages that are unchanged and not recompiled.

4.2 Technology previews

Technology previews are packages, stacks, or features delivered by SUSE to provide glimpses into upcoming innovations. Technology previews are included for your convenience to give you a chance to test new technologies within your environment. We would appreciate your feedback. If you test a technology preview, please contact your SUSE representative and let them know about your experience and use cases. Your input is helpful for future development.

Technology previews have the following limitations:

- Technology previews are still in development. Therefore, they may be functionally incomplete, unstable, or otherwise *not* suitable for production use.

- Technology previews are *not* supported.

- Technology previews may only be available for specific hardware architectures.

- Details and functionality of technology previews are subject to change. As a result, upgrading to subsequent releases of a technology preview may be impossible and require a fresh installation.

- SUSE may discover that a preview does not meet customer or market needs, or does not comply with enterprise standards. Technology previews can be removed from a product at any time. SUSE does not commit to providing a supported version of such technologies in the future.

For an overview of technology previews shipped with your product, see the release notes at https://www.suse.com/releasenotes ↗ .

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 an overview of which software components and services are also available with other products from SUSE (green) and which are exclusively available with SUSE Linux Enterprise Server for SAP Applications (blue).

```
·········· 24×7 Priority Support for SAP ··········

    Extended Service Pack Overlap Support
              18 Month Grace Period

  HA Cluster      Antivirus      SAP HANA      Simplified
  Connector       ClamSAP        Security      Operations
                                               Management

  High Availability for       SAP HANA       Faster
  SAP NetWeaver & SAP HANA    HA Resource    Installation
  SUSE Linux Enterprise       Agents         Installation
  High Availability                          Workflow

    Reliable, Scalable and Secure Operating System
              SUSE Linux Enterprise Server

·········· 24×7 Priority Support for SAP ··········
```

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 but contains several additional software components such as SUSE Linux Enterprise High Availability, 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 15 SP4. 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

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 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, see the Administration Guide (https://documentation.suse.com/sles-15 ↗) and the White Papers and Best Practice Guides 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 using components of SUSE Linux Enterprise High Availability 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 parent/child resource: The parent assumes responsibility for the SAP HANA databases running in primary mode, whereas the child is responsible for instances that are operated in synchronous (secondary) status. In case of a takeover, the secondary (child resource instance) can automatically be promoted to become the new primary (parent resource instance).

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

- **Performance-optimized scenario.** Two servers (A and B) in the same SUSE Linux Enterprise High Availability 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 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), as 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 cluster and 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 7, 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-15 ↗.

- 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 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 11, Protecting against malware with ClamSAP.*

1.1.6 SAP HANA security

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

1.1.6.1 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` which provides preconfigured rules and integrates with `firewalld`.

For more information, see *Section 10.2, "Configuring HANA-Firewall"*.

1.1.6.2 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*.

1.1.7 Simplified operations management

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

1.1.7.1 System tuning with `saptune`

The system tuning application `saptune` 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, `saptune` 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 *Chapter 8, Tuning systems with* `saptune`.

1.1.7.2 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 is 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.

1.1.7.3 `ClusterTools2`

`ClusterTools2` provides tools that help set up and manage a Corosync & Pacemaker cluster. Among them are `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`.

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	`SLE-Module-Basesys-tem15-SP4-Pool`	`SLE-Module-Basesys-tem15-SP4-Updates`
Basic server functionality of SUSE Linux Enterprise Server	`SLE-Module-Server-Appli-cations15-SP4-Pool`	`SLE-Module-Server-Appli-cations15-SP4-Updates`
Packages specific to SUSE Linux Enterprise Server for SAP Applications	`SLE-Module-SAP-Applica-tions15-SP4-Pool`	`SLE-Module-SAP-Applica-tions15-SP4-Updates`
Packages specific to SUSE Linux Enterprise High Availability	`SLE-Product-HA15-SP4-Pool`	`SLE-Product-HA15-SP4-Up-dates`

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-15-SP4-SAP-15.4-0` 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 SLE Modules and SLE Extensions either during the installation or from the running system using YaST or the command `SUSEConnect`.

For information about all modules and extensions available for the SUSE Linux Enterprise product line, see https://documentation.suse.com/sles-15/html/SLES-all/art-modules.html ↗.

For more information about SUSE Package Hub, see *Section A.3, "SUSE Package Hub"*. 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 (ESPOS). 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).

Extended Service Pack Overlap Support allows you to perform service pack migrations within three and a half 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 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 years of general support and three years of ESPOS for each service pack.

The last service pack in each SLE family does not have ESPOS. Instead of ESPOS, it includes a longer general support period. Because of that, LTSS is available only for the last service pack. All other service packs already include three years of ESPOS, which equals LTSS.

For more information, refer to the following resources:

- Product Lifecycle Support Policies: https://www.suse.com/support/policy-prod-ucts/#sap ⬈

- Product Support Lifecycle Lifecycle Dates by Product: https://www.suse.com/lifecy-cle/ ⬈

- Long Term Service Pack Support: https://www.suse.com/products/long-term-ser-vice-pack-support/ ⬈

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 for 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 (SUSE Package Hub and SUSE 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 and 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↗.

Supported CPU

> Intel 64
>
> IBM POWER 8 (with PowerVM)
>
> IBM POWER 9 (with PowerVM)

Hard disk

> SUSE Linux Enterprise Server for SAP Applications requires at least 41 GB (without swap) 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 refer to https://help.sap.com/docs/SAP_HANA_PLAT-FORM ↗ (the section *Implement › SAP HANA Master Guide › SAP HANA Deployment Options › On-Premise Deployments*).

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: Hardware 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 Installation image

Unlike previous SLE products, the entire SLE 15 SP4 product line can be installed from a single installation medium: `SLE 15 SP4 Online media 1`. To install without network access or registration, download the `SLE 15 SP4 Full media 1` image. Both ISO images are available from https://download.suse.com/ ↗.

Burn the image onto a physical DVD or copy it to a removable flash disk. Make sure the size of the disk is sufficient for the desired image. Alternatively, use a virtual DVD-ROM device for installation in a virtual machine.

💡 Tip: Copying the installation media image to a removable flash disk

Use the following command to copy the contents of the installation image to a removable flash disk.

```
> sudo dd if=IMAGE of=FLASH_DISK bs=4M && sync
```

Replace `IMAGE` with the path to the installation media image file and `FLASH_DISK` with the flash device.

2.3 Offline migration

The migration paths for SUSE Linux Enterprise Server for SAP Applications are identical to those for SUSE Linux Enterprise Server with Enhanced Service Pack Overlay Support (ESPOS). Find detailed information in the *SUSE Linux Enterprise Server* Upgrade Guide at https://documenta-tion.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:

- *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, "Instal-lation workflow"*.

2. SAP Installation Wizard, part 1: Copying all required SAP media to the local disk or se-lecting 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 partic-ular *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*.

Most of these steps do not need to be run immediately after each other, which allows for flexi-bility 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 search list (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 2578899: SUSE Linux Enterprise Server 15: Installation notes* (https://launchpad.support.s-ap.com/#/notes/2578899 ↗).

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 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 for 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.

> ❗ **Important: Installing Oracle databases**
>
> To be able to install an Oracle database later, install SUSE Linux Enterprise Server first and then convert your installation to SUSE Linux Enterprise Server for SAP Applications.
>
> This is necessary because the installer for Oracle databases queries for the existence of certain files, not all of which are included in an SLES for SAP installation.
>
> For more information about converting, see *Section 3.4, "Converting an SLES installation to an SLES for SAP installation"*.

1.   - On AMD64/Intel 64, boot from the installation media. From the 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-15 ↗).

   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. As the first step, the installation workflow will check for updates for itself. After that, it will be ready to start the installation.

2. Select the default system language under *Language*.
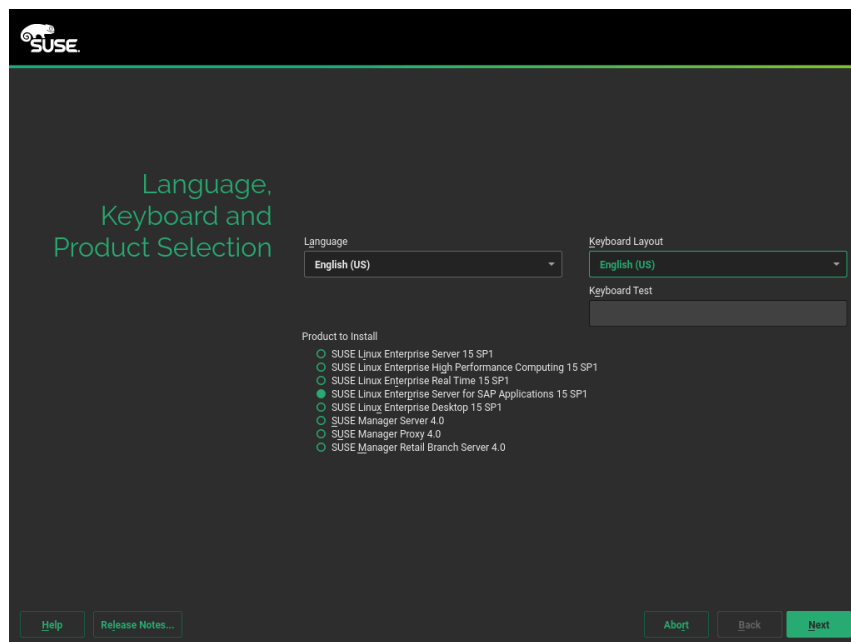
FIGURE 3.1: LANGUAGE, KEYBOARD AND PRODUCT SELECTION

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. SLE 15 SP4 provides a single installation ISO for the entire product line. Therefore, you need to select the product to install on this page.
   Under *Product to install,* choose *SUSE Linux Enterprise Server for SAP Applications 15 SP4*.

5. 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*.

6. *(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-15 ↗).

7. 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 can only install a minimal SLE system and will not receive updates.
   >
   > To install a full (but not updated) SLES for SAP system without network access during the installation, use the SLE 15 SP4 Packages ISO image from https://download.suse.com ↗ . You can then choose *Skip registration* on this page and select the SLE 15 SP4 Packages ISO image as an add-on product on the next page.

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

9. After the system is successfully registered, YaST lists available modules for SUSE Linux Enterprise Server for SAP Applications from the SUSE Customer Center. The default selection covers the most common cases. To enable an additional module, activate its entry.

> ### 🖎 Note: Release notes
>
> From this point on, the release notes can be viewed from any screen during the installation process by selecting *Release Notes*.

Proceed with *Next*.

10. *(Optional)* The *Add On Product* dialog allows you to add additional software sources (so-called "repositories") to SUSE Linux Enterprise Server for SAP Applications, that are not provided by the SUSE Customer Center. Such add-on products may include third-party products, drivers, or additional software for your system.

11. Choose the *System Role*. System roles are predefined use cases which tailor the system for the selected scenario. For SUSE Linux Enterprise Server for SAP Applications, you can choose between:

    - *SLES for SAP Applications*: Default, recommended for most situations. This system role contains the following properties:

        - Supports the installation wizard for SUSE Linux Enterprise Server for SAP Applications.

        - Enables RDP access (*Remote Desktop Protocol*).

        - Provides special partitioning recommendations.

    - *SLES with GNOME*: Can be necessary in specific cases. This installation path is not covered in this document. For more information about this installation path, see *Installation Quick Start, Section "Installing SUSE Linux Enterprise Server"* (https://documentation.suse.com/sles-15 ↗).

    Additional system roles are available for specific use cases (High Availability, text mode, minimal, and KVM/XEN virtualization hosts).
    Proceed with *Next*.
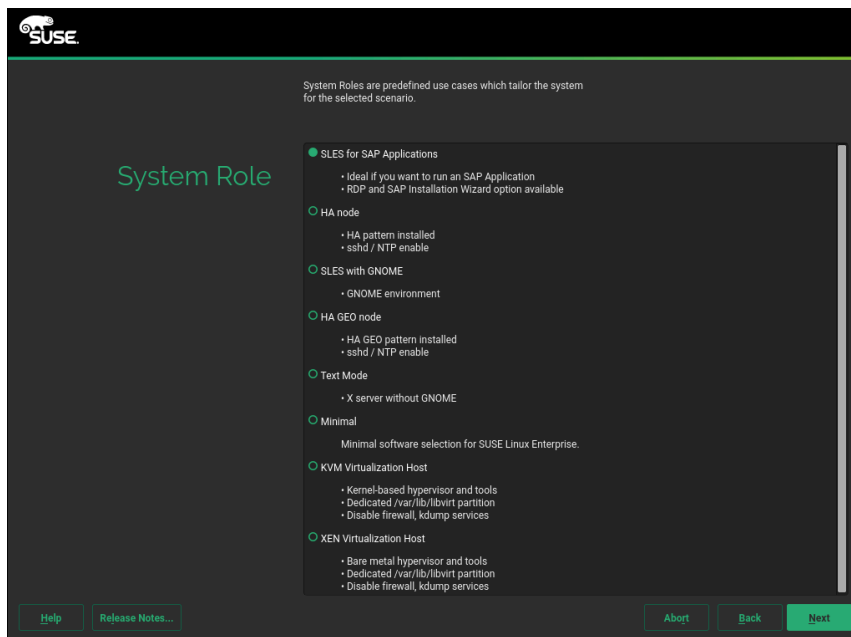
FIGURE 3.2: SYSTEM ROLE

PROCEDURE 3.2: FINISHING THE OS INSTALLATION

1. Choose whether to enable the following options:

    - To install an SAP Application along 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 12, Connecting via RDP.*

2. Review the proposed partition setup for the volumes `/dev/system/root` and `/dev/system/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:

    Guided setup

    Create a new partitioning suggestion based on your input.

    Expert partitioner

    Open the *Expert Partitioner* described in *Deployment Guide, Chapter "Advanced Disk Setup", Section "Using the YaST Partitioner"* (https://documentation.suse.com/sles-15 ↗).

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*.

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-15 ↗).

   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.

   In case you want to enable password-less authentication via SSH login, you can import a key via *Import Public SSH Key*. If you want to completely disable `root` login via password, upload a key only and do not provide a root password. A login as system administrator will only be possible via SSH using the respective key in this case.

   For more information, see *Deployment Guide, Chapter "Installation with YaST", Section "Password for the System Administrator root"* (https://documentation.suse.com/sles-15 ↗).

   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 for 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 10.1, "Configuring `firewalld`"*.

Installation workflow   SLES for SAP 15 SP4
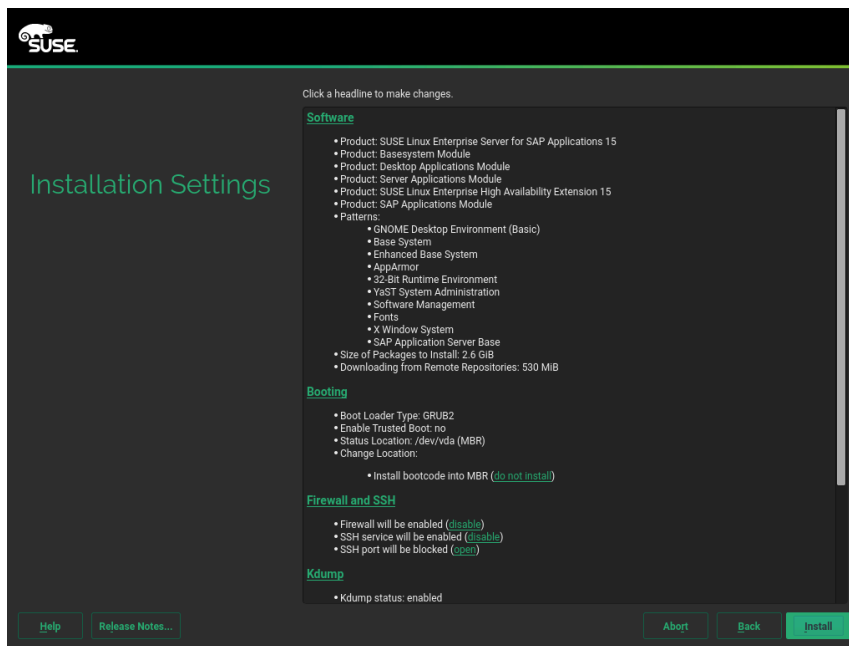
FIGURE 3.3: INSTALLATION SETTINGS

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 now, the installation will continue after a reboot. Continue with *Chapter 4, Installing SAP applications*.

3.2   Using SLES for 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 for SAP.

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

2. Boot from an SLES installation medium.

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, "Installation workflow"*.

For more information, see *Deployment Guide, Chapter "Remote Installation"* (https://documentation.suse.com/sles-15 ↗).

To avoid having to use an SLES installation medium 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-15 ↗).

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-15 ↗).

- *AutoYaST Guide* (https://documentation.suse.com/sles-15 ↗).

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 an SLES installation to an SLES for SAP installation

To convert an installation of SUSE Linux Enterprise Server 15 SP4 or JeOS 15 SP4 to an installation of SLES for 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 for SAP.

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

2. Execute the following command:

   ```
   # Migrate_SLES_to_SLES-for-SAP.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 for SAP packages**
>
> The script does not install all packages that are included with a default SLES for SAP installation. However, you can install these yourself manually. To install the default package selection, use:
>
> ```
> # zypper in patterns-server-enterprise-sap_server
> ```

Warning: Migration script on public cloud

On public cloud instances (pay-as-you-go instances in particular), billing depends on internal mechanisms to identify the entitlement and to calculate the actual consumption. This makes the migration script ineffective, as it only performs migration of repositories inside the operating system.

To perform migration, you must follow image migration guidelines by your cloud solution provider.

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

For more information about converting, see *Section 3.4, "Converting an SLES installation to an SLES for 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://documenta-tion.suse.com/sles-15 ↗).

(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 an SAP HANA system.

To install other SAP applications or to create a more advanced SAP HANA setup, directly use one of the installation methods provided by SAP instead of this wizard.

💡 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.
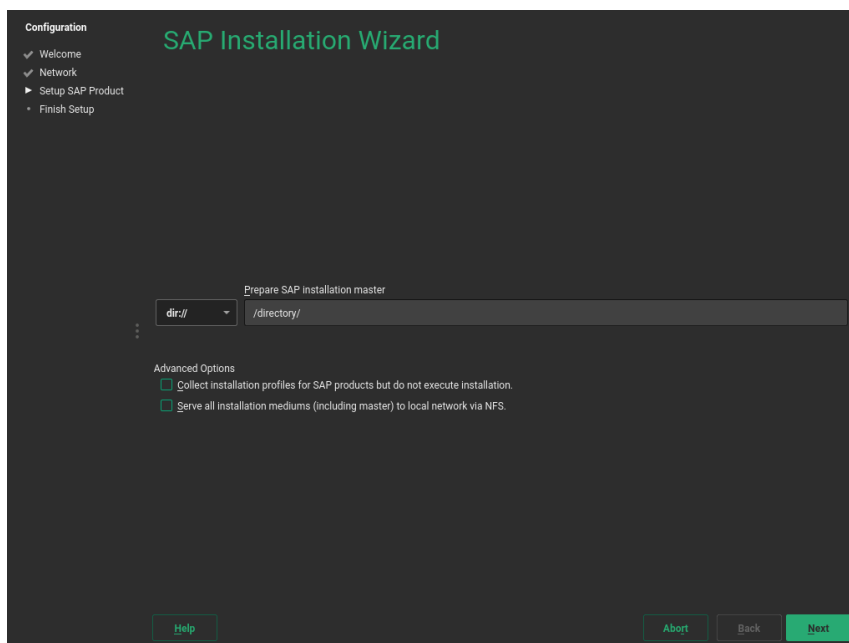
FIGURE 4.1: **LOCATION OF SAP INSTALLATION MASTER**

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 6, 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 *Choose 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 if you do not need additional Installation Media or 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 will 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.
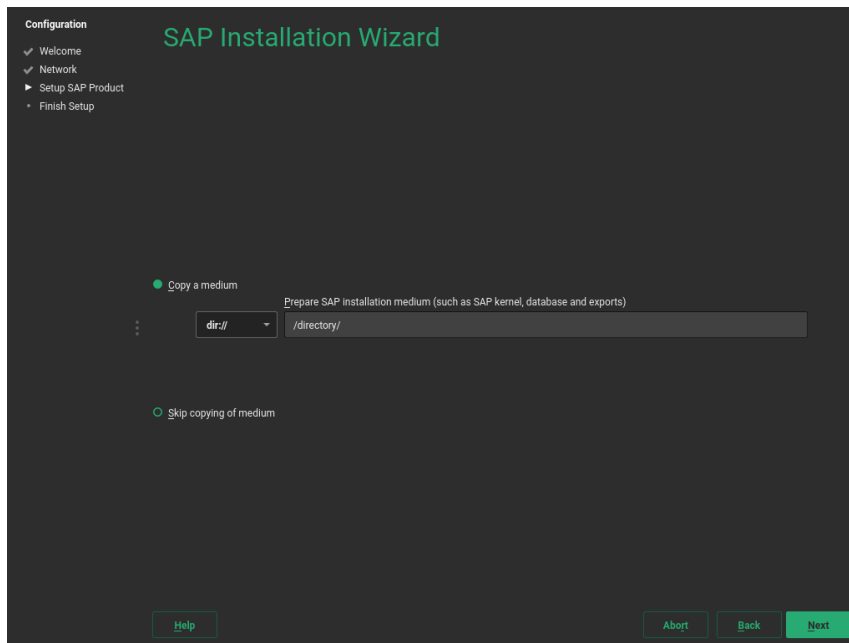
FIGURE 4.2: SAP INSTALLATION WIZARD: ADDITIONAL INSTALLATION MEDIA

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 the various 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.
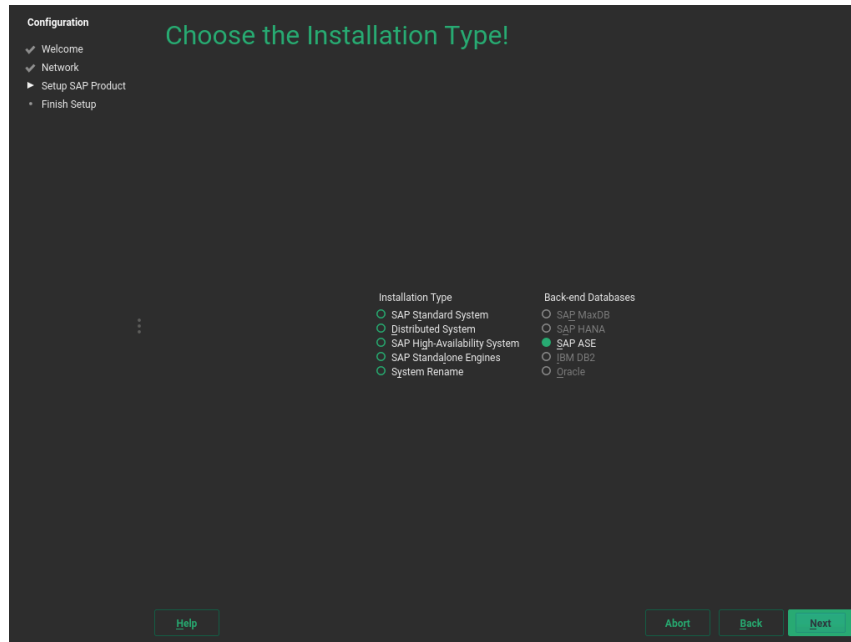
FIGURE 4.3: SAP INSTALLATION WIZARD: INSTALLATION TYPE AND DATABASE

6. If you selected *SAP Standard System*, *Distributed System*, or *SAP High-Availability System*, additionally choose a back-end database under *Back-end Databases*.
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*.
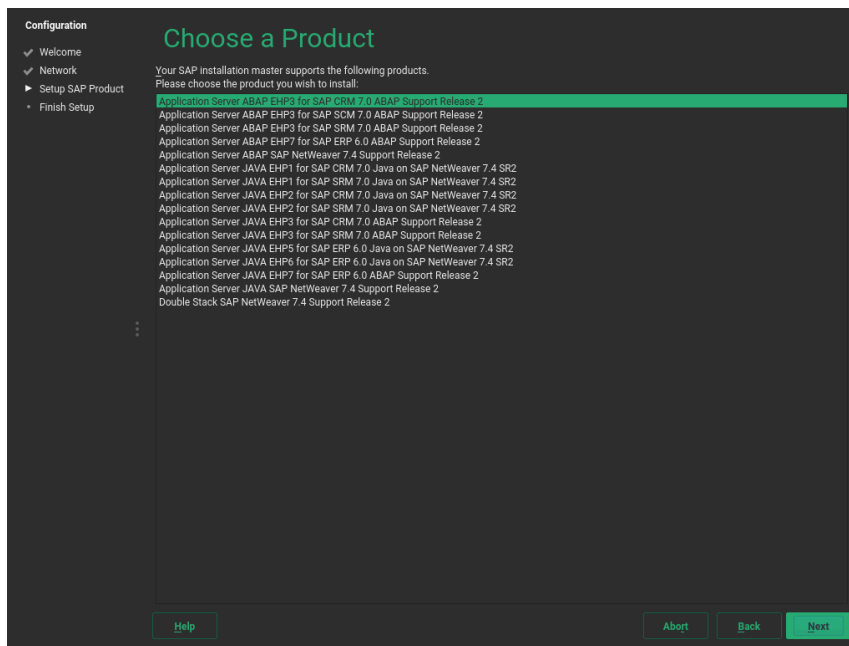
FIGURE 4.4: **SAP INSTALLATION WIZARD: CHOOSE A PRODUCT**

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-15 ↗).

   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:
   >
   > ```
   > > 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 several 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*.
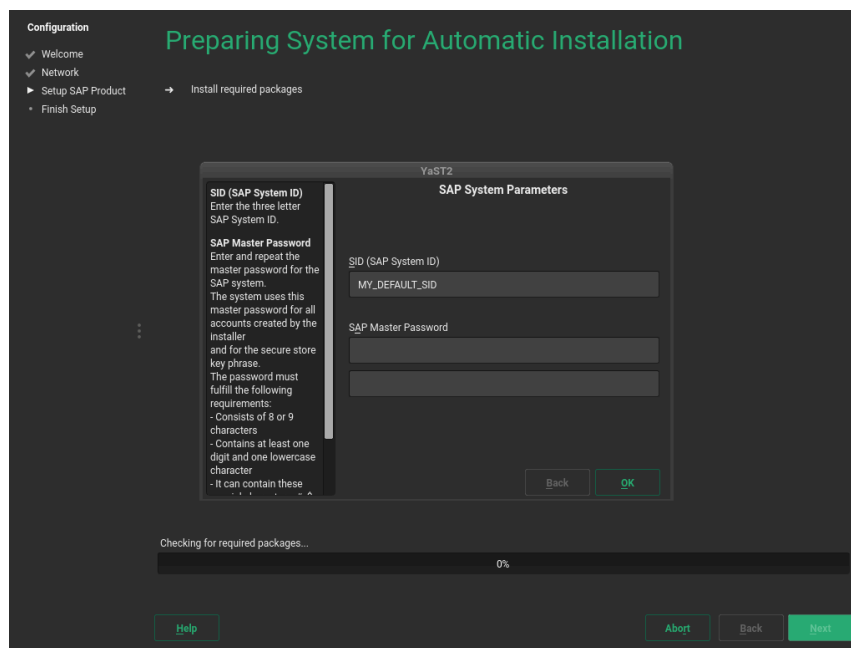
FIGURE 4.5: **PRODUCT PARAMETERS**

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 ↗ (you need your SAP ID to access the information). 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.
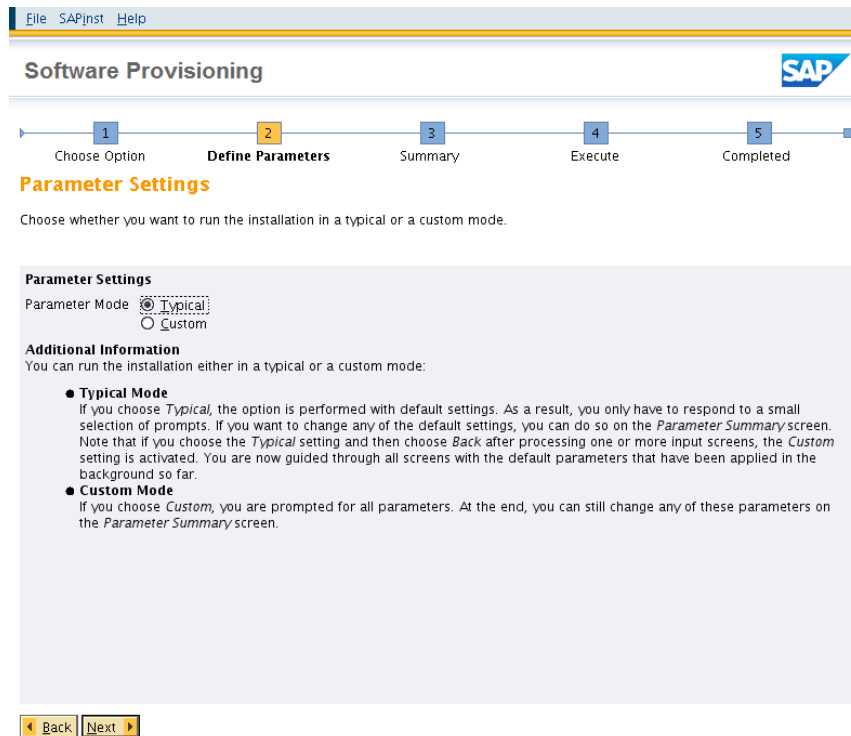
FIGURE 4.6: **SAP INSTALLER: DEFINING PARAMETERS**

> 💡 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 14, 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 will show 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 offer 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 your 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:

```
# yast2 sap_create_storage_ng 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:

```xml
<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>
```

5 Upgrading an SAP HANA cluster

This chapter describes how to upgrade your SAP HANA cluster with the YaST module *SUSE HANA Cluster Update*. This acts as a wizard and guides you through all the SAP HANA cluster maintenance procedures.

The official SAP HANA documentation describes the so-called *Near Zero Downtime Upgrade Process*. The YaST module is based on this process and handles the part of the procedure related to the SUSE cluster. Not all steps can be done automatically. Some steps need to be performed manually by the SAP HANA administrator. The YaST module will inform you during the process.

This YaST module is available in the `yast2-sap-ha` package for SUSE Linux Enterprise Server for SAP Applications 12 SP3 and higher. Currently, the wizard is only prepared to handle the *SAP HANA Scale-up Performance Optimized* scenario.

The upgrade covers the following tasks:

1. *Section 5.1, "Preparing the upgrade"*

2. *Section 5.2, "Upgrading your SAP HANA cluster"*

3. *Section 5.3, "Finishing the upgrade task"*

5.1 Preparing the upgrade

Ensure passwordless SSH access between the two nodes (primary and secondary) for `root`. Keep in mind, some cloud service providers might not have set up SSH access for the `root` by default.

1. Install the `yast2-hana-update` package on both nodes:

   ```
   # zypper install yast2-hana-update
   ```

   After the installation, you can find the module *SUSE HANA Cluster Update* in the *YaST Control Center*.

2. On the secondary node, start the *YaST Control Center* and open the *SUSE HANA Cluster Update* module.

3. In the YaST module, review the prerequisites. Make sure to fulfill all of them before continuing with the next step. Keep in mind that the wizard supports only the HANA Scaleup Performance Optimized scenario.

4. To upgrade the SAP HANA system, select the secondary node.

5. Select the location of the installation medium.
   Point to the location where the SAP medium is located. If wanted, check *Mount an update medium on all hosts* and provide the NFS share and path.

   > **❗ Important: Differences between SAP HANA version 1.0 and 2.0**
   >
   > If you are upgrading from SAP HANA version 1.0 to version 2.0, make sure to check *This is a HANA 1.0 to HANA 2.0 upgrade.*
   >
   > The YaST module will copy the *PKI SSFS keys* from the former secondary node to the former primary node. More information is available through the *Help* button.

Continue with *Section 5.2, "Upgrading your SAP HANA cluster"*.

5.2 Upgrading your SAP HANA cluster

1. Review the update plan generated by the wizard.
   The wizard shows you two steps: automatic and manual. In this automatic step, the wizard puts cluster resources into maintenance mode before it starts with the automatic steps. The manual steps are SAP HANA specific and need to be executed by an SAP HANA administrator. For more information, see the official SAP HANA documentation.

2. Update the SAP HANA software.
   The wizard executes the automatic actions and waits until the SAP HANA administrator performs the SAP HANA upgrade.

3. Perform the SAP HANA upgrade.

4. Review the plan for the primary (remote) node.
   After the SAP HANA upgrade is done, the wizard shows the update plan. When you continue with this step, the wizard turns the primary node into a secondary node to make it ready for the upgrade.

> Keep in mind that this step can take some time.

Continue with *Section 5.3, "Finishing the upgrade task"*.

5.3 Finishing the upgrade task

1. Update the former primary node.

   Pay special attention to the `--hdbupd_server_nostart` option in this step.

2. Restore the previous state of the cluster.

   By default, the wizard registers the former master as now being secondary on the SAP HANA system replication. If you want to revert the system replication to its original state, click the *Reverse* button.

3. Review the update summary.

   You can review the original and current SAP HANA versions and the cluster state.

   ### Note: Dealing with intermediate cluster state

   If the wizard is faster than the status update of the cluster resources, the summary shows an intermediate cluster state. The cluster state is `UNDEFINED` or `DEMOTED`.

   To overcome this, check the cluster status again with the command **SAPHanaSR-showAttr** and make sure the former secondary node is now in the state `PROMOTED`.

Refer to the SUSE blog post https://www.suse.com/c/how-to-upgrade-your-suse-sap-hana-cluster-in-an-easy-way/ ↗ for further information.

6 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 6.1: 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 a Windows operating system for 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-15 ⬈).

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-15 ⬈).

7 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"*.

Administrators can now use the SAP HANA-SR Wizard to run the module unattended, usually for on-premises deployments. Additionally, it is possible to configure the SAP HANA cluster on Azure now. The YaST module identifies automatically when running on Azure and configures an extra resource needed on Pacemaker.

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 an SAP HANA SR Performance Optimized Infrastructure*

- **Cost-optimized scenario:** *Setting up an 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/sles-15 ↗.

7.1 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 and Hardening Guide, Part "Network Security", Chapter "SSH: Secure Network Operations"* at https://documentation.suse.com/sles-15 ↗ .

- 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 10.2, "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, refer to the *SAP HANA System Replication Scale-Up - Cost Optimized Scenario* document at https://documentation.suse.com/sles-sap/ ↗ .

7.2 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 7.1, "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*.
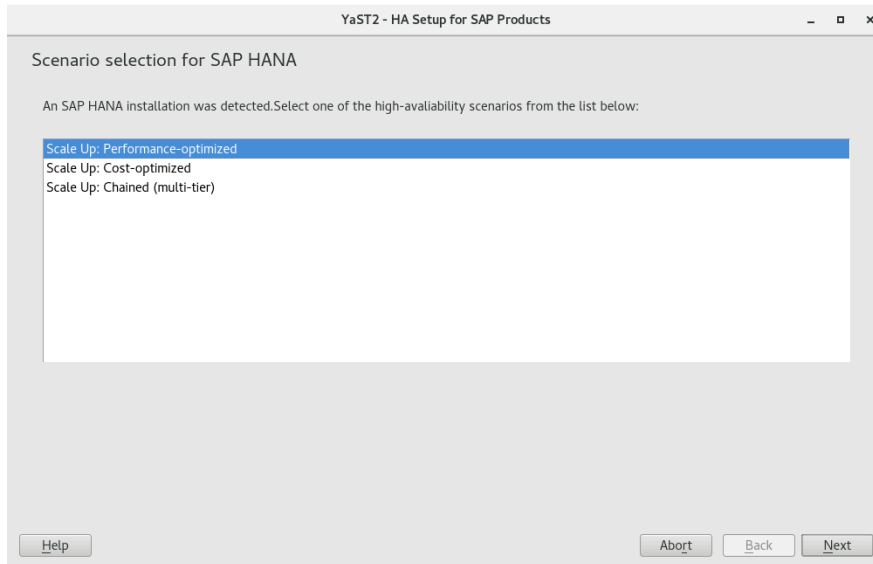
3. This step of the wizard presents a list of prerequisites for the chosen scale-up scenario. These prerequisites are the same as those presented in *Section 7.1, "Prerequisites"*.
   Continue with *Next*.

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/sles-15 ↗ .

   - 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/sles-15 ↗ .

  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/sles-15 ↗ .

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, such as during the initial installation of SAP HANA, 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/sles-15 ↗ .

Proceed with *Next.*

9. The following page allows configuring watchdogs which 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 us-ing a hardware watchdog in production environments instead of `softdog`. For more information about selecting watchdogs, see *Administration Guide, Part "Storage and Da-ta Replication", Chapter "Storage Protection", Section "Conceptual Overview", Section "Set-ting Up Storage-based Protection", Section "Setting up the Watchdog"* at https://documenta-tion.suse.com/sles-15 ↗ .

   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, refer to the documentation provided by SAP.

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 7.1: 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 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
```

7.3 Unattended setup using SAP HANA-SR wizard

An unattended setup requires a manual installation of HANA first. The result is saved into a file containing all configuration options that were chosen. If the administrator needs to reproduce the installation, with this file the installation can be run automatically and unattended.

To use it, perform the following steps on both nodes:

1. On the production machines with SAP HANA installed, create a configuration file by running the `sap_ha` YaST module.

2. On the last screen, click the *Save configuration* button.

3. Decide what you want to do:

   * To review the configuration, upload and validate the configuration on the primary SAP HANA machine and run:

     ```
     # yast2 sap_ha readconfig CONFIGURATION_FILE_PATH
     ```

     It is possible to start the installation on the review screen.

   * To start the installation based on the provided configuration file unattended, run:

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

4. Import, validate, and install the cluster unattended, based on the provided configuration file:

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

7.4 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 and 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 7.2, "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:
>
> ```
> # passwd hacluster
> ```

7.5 For more information

- **Hawk.** *Administration Guide,* Part *Configuration and Administration,* Chapter *Configuring and Managing Cluster Resources with Hawk* (https://documentation.suse.com/sles-15 ↗).

- **Near zero downtime for SAP HANA system replication.** Use SAP HANA System Replication for Near Zero Downtime Upgrades (https://help.sap.com/viewer/2c1988d620e04368aa4103bf26f17727/2.0.03/en-US/ee3fd9a0c2e74733a74e4ad140fde60b.html) ↗ .

- **Implementing the Python hook SAPHanaSR.** https://documentation.suse.com/sbp/all/html/SLES4SAP-hana-sr-guide-PerfOpt-15/ ↗

8 Tuning systems with **saptune**

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

Using `saptune`, you can tune a system for SAP NetWeaver, SAP HANA/SAP BusinessObjects, and SAP S/4HANA applications.

8.1 Installing and updating **saptune**

To install `saptune`, run the `zypper install saptune` command.

When installation is completed, enable and start the `saptune` service (see *Section 8.2, "Enabling and disabling **saptune**"*) and configure the tuning (see *Section 8.3, "Configuring the tuning"*).

To update `saptune`, use the `zypper update saptune` command.

> **⚠ Important**
>
> When installing and updating `saptune`, pay attention to zypper output to ensure that installation and updates are performed correctly. The output is also saved in `/var/log/zypp/history`.

8.2 Enabling and disabling **saptune**

To tune the system after a reboot, enable and start the `saptune` service after installation. In most cases, starting the `saptune` service fails, because `sapconf` already tunes the system. To solve the problem, run the following command:

```
# saptune service takeover
```

This command stops and disables the `sapconf` and `tuned` services, and then starts and enables the `saptune` service.

To disable and stop the `saptune` service, use the command below:

```
# saptune service disablestop
```

8.3 Configuring the tuning

The easiest way to tune a system is to apply an SAP Solution that matches your installed SAP software. SAP Solutions are a group of SAP Notes that perform the actual tuning. To display all available Solutions and their Notes, run the following command:

```
# saptune solution list
```

`saptune` recognizes the following tuning SAP Solutions:

- BOBJ Solution for running SAP BusinessObjects

- HANA Solution for running an 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

- S4HANA-APP+DB Solution for running both SAP S/4HANA Application Servers and SAP HANA on the same host

- S4HANA-DBSERVER Solution for running the SAP HANA database of an SAP S/4HANA installation

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

- NETWEAVER+HANA Solution for running both SAP application servers and SAP HANA on the same host

- NETWEAVER+MAXDB Solution for running both SAP application servers and MAXDB on the same host

To apply a Solution, run the following command:

```
# saptune solution apply SOLUTION
```

Keep in mind that only one Solution can be applied at the time.

To disable a Solution, use the command below:

```
# saptune solution revert SOLUTION
```

To switch to a different Solution, use the following command:

```
# saptune solution change SOLUTION
```

Alternatively, you can tune the computer according to recommendations from specific SAP Notes. Use the `saptune note list` to view a list of notes that you can tune for.

To apply a Note, run the following command:

```
# saptune note apply NOTE
```

Reverting a Note can be done as follows:

```
# saptune note revert NOTE
```

Note: Combining optimizations

It is possible to combine Solutions and Notes by reverting Notes from an applied Solution or applying additional ones. However, only one solution can be active at a time. The `saptune` service restores the combination of Solution and Notes after a service restart or reboot.

In rare cases, Notes can have conflicting options or parameters. Arrange your Notes carefully to avoid conflicts. The last Note always takes priority over conflicting options or parameters of previous notes. In this situation, create your own Solution (see *Section 8.4.2, "Creating a new SAP Note"*) or customize the applied Solution (see *Section 8.4.1, "Customizing an SAP Note"*).

8.4 Managing SAP Notes

The following sections provide information on working with SAP Notes.

An SAP Note configuration contains the OS-specific part of the original SAP Note as complete as possible. A parameter is disabled (it is present in the configuration, but without value) if it does not have a value recommendation, or if `saptune` cannot safely detect the conditions to set the correct value. To set a suitable value, read the corresponding SAP Note and customize the Note (see *Section 8.4.1, "Customizing an SAP Note"*).

8.4.1 Customizing an SAP Note

Any SAP Note can be configured using the following command:

```
# saptune note customise NOTE
```

The command opens the default editor (defined in the environment variable `EDITOR`) with a copy of the Note configuration. Remove everything, except the parameters you want to change or disable, as well as the header of the section the parameter belongs to.

To change or set the parameter value, change or add the value of the parameter. To disable a parameter, remove the value, but leave the parameter and the `=` character. **saptune** lists the parameter, but it does not change it or check it for the compliance status. For more information, refer to the `saptune-note(5)` manpage.

This creates a `/etc/saptune/override/`*NOTE* file. It is possible to create the file elsewhere and place it in `/etc/saptune/override/`.

Configuration sections can be conditional. This is called tagging. Refer to the `saptune-note(5)` for further information.

When you are done customizing a Note, restart the `saptune` service to apply the changes.

8.4.2   Creating a new SAP Note

A new SAP Note can be created using the following command:

```
# saptune note create NOTE
```

The command opens the default editor (defined in the environment variable `EDITOR`) with a Note configuration template. All features of **saptune** are available here. For more information, refer to the `saptune-note(5)` manpage.

This creates a `/etc/saptune/extra/`*NOTE.conf* Note configuration file. It is possible to create the file elsewhere and place it in `/etc/saptune/extra/`.

Configuration sections can be conditional. This is called tagging. Refer to the `saptune-note(5)` for further information.

8.4.3   Editing a custom SAP Note

To edit a custom Note, use the command below:

```
# saptune note edit NOTE
```

The command opens the default editor (`EDITOR`) with the Note configuration. When you are done editing a Note, restart the `saptune` service to apply the changes. Custom Notes can be customized like shipped Notes.

8.4.4 Deleting an SAP Note or a customization

The following command deletes a note, including the corresponding override file, if available:

```
# saptune note delete test
Note to delete is a customer/vendor specific Note.
Do you really want to delete this Note (test2)? [y/n]: y
```

The note may not be applied at the time. Keep in mind the following:

- A confirmation is needed to finish the action.

- Internal SAP Notes shipped by **saptune** cannot be deleted. Instead, the override file is removed when available.

- If the Note is already applied, the command is terminated with the message that the note first needs to be reverted before it can be deleted.

8.4.5 Renaming an SAP Note

This command allows renaming a created Note to a new name. If a corresponding override file is available, this file is renamed as well:

```
# saptune note rename test test2
Note to rename is a customer/vendor specific Note.
Do you really want to rename this Note (test) to the new name 'test2'? [y/n]: y
```

The Note may not be applied at the time. Keep in mind the following points:

- A confirmation is needed to finish the action.

- Internal SAP Notes shipped by **saptune** cannot be renamed.

- If the Note is already applied, the command is terminated with the information that the Note first needs to be reverted before it can be deleted.

8.4.6 Showing the configuration of an SAP Note

The configuration of a Note can be listed using the following command:

```
# saptune note show NOTE
```

8.4.7 Verifying an SAP Note

To verify the tuning of a Note, use the following command:

```
# saptune note verify NOTE
```

For information about the output of the command and verifying the entire tuning instead of a single Note, refer to *Section 8.6, "Verification and troubleshooting"*.

8.4.8 Performing a dry run of an SAP Note

To show each parameter of a Note, use the following command:

```
# saptune note simulate
```

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

> **Note: Deprecation notice**
>
> The **simulate** command is deprecated since version 3.1 and is removed in all **saptune** versions in SUSE Linux Enterprise Server for SAP Applications 16.

8.4.9 Reverting an SAP Note

To revert an SAP Note, run the following command:

```
# saptune note revert NOTE
```

This restores all parameters of the SAP Note to their values at the time of application.

To revert everything, use the following command:

```
# saptune note revert all
```

8.4.10 Listing all enabled or applied SAP Notes

To list all enabled SAP Notes, run the following command:

```
# saptune note enabled
```

To list all applied SAP Notes, run the command below:

```
# saptune note applied
```

Both commands are primarily meant for use in scripts.

8.5 Managing SAP Solutions

This chapter explains how to work with SAP Solutions.

An SAP Solution is a combination of SAP Note configurations grouped logically. It generally represents an SAP product or combination. Applying an SAP Solution effectively applies all SAP Note configurations listed in it. Instructions for listing and setting a solution are provided in *Section 8.3, "Configuring the tuning"*.

8.5.1 Customizing an SAP Solution

An SAP Solution can be customized using the following command:

```
# saptune solution customise SOLUTION
```

The command opens the default editor (defined in the environment variable `EDITOR`) with a copy of the Solution configuration. Change the Note list for the architecture to your liking. For more information, refer to the `saptune-note(5)` manpage.

This creates an override file `/etc/saptune/override/SOLUTION.sol`. It is possible to create the file elsewhere and place it in `/etc/saptune/override/`.

When you are done customizing an SAP Solution, restart the `saptune` service to apply the changes.

8.5.2 Creating a new SAP Solution

To create a new SAP Solution, run the following command:

```
# saptune solution create SOLUTION
```

The command opens the default editor (defined in the environment variable `EDITOR`) with a Solution configuration template. Fill in the template.

This creates a Solution configuration file `/etc/saptune/extra/SOLUTION.sol`. It is possible to create the file elsewhere and place it in `/etc/saptune/extra/`.

8.5.3   Editing a custom SAP Solution

To edit a custom SAP Solution, use the following command:

```
# saptune solution edit NOTE
```

The command opens the default editor (defined in the environment variable `EDITOR`) with the Solution configuration.

When you are done editing an SAP Solution, restart the saptune service to apply the changes.

Custom Solutions can be customized like shipped Solutions.

8.5.4   Deleting SAP Solution

The following command deletes a created Solution (in this example, myHANA), including the corresponding override file or the override file of a shipped Solution, if available:

```
# saptune solution delete myHANA

Solution to delete is a customer/vendor specific Solution.
Do you really want to delete this Solution 'myHANA'? [y/n]: y
```

The SAP Solution may not be applied at the time. Keep in mind the following:

- A confirmation is required to finish the action.

- SAP Solutions shipped by **saptune** cannot be deleted. Only the override file is removed, if available.

- If the SAP Solution is already applied, the command is terminated with the information that the SAP Solution first needs to be reverted before it can be deleted.

8.5.5   Renaming an SAP Solution

To rename an SAP Solution, run the following command:

```
# saptune solution rename myHANA myHANA2
```

```
Solution to rename is a customer/vendor specific Solution.
Do you really want to rename this Solution 'myHANA' to the new name 'myHANA2'? [y/n]:
```

The SAP Solution may not be applied at the time. Keep in mind the following points:

- A confirmation is needed to finish the action.

- SAP Solutions shipped by `saptune` cannot be renamed.

- If the SAP Solution is already applied, the command will be terminated with the information that the SAP Solution first needs to be reverted before it can be renamed.

8.5.6 Showing the configuration of an SAP Solution

To list the configuration of an SAP Solution, run the following command:

```
# saptune solution show SOLUTION
```

8.5.7 Switching to another SAP Solution

Starting with `saptune` version 3.1, it is easier to switch to a different solution using the `saptune solution change SOLUTION` command.

Keep in mind that internally the current solution is reverted first, and then the new solution is applied. If you have additional notes configured, the order is not preserved.

If the same solution is already applied, no action is taken. Otherwise the current solution gets reverted and the new one applied. The command prompts for confirmation before making the change. This can be disabled by adding the `--force` option.

8.5.8 Verifying an SAP Solution

To verify the tuning of a Solution, use the following command:

```
# saptune solution verify SOLUTION
```

For information about the output of the `verify` command and how to verify the entire tuning instead of a single Solution, refer to *Section 8.6, "Verification and troubleshooting"*.

8.5.9 Performing a dry run of an SAP Solution

To show all parameters of all Notes belonging to a Solution, use the following command:

```
# saptune solution simulate SOLUTION
```

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

> 📝 **Note: Deprecation notice**
>
> The **simulate** command is deprecated since 3.1, and it is removed in all **saptune** versions in SUSE Linux Enterprise Server for SAP Applications 16.

8.5.10 Reverting an SAP Solution

To revert an SAP Solution, run the following command:

```
# saptune solution revert SOLUTION
```

The SAP Solution must be applied. This reverts all SAP Notes parts of the SAP Solution that are still applied.

8.5.11 Editing a custom SAP Solution

To edit a custom SAP Solution, run:

```
# saptune solution edit SOLUTION
```

8.5.12 Listing enabled/applied SAP Solution

To list an enabled SAP Solution, run:

```
# saptune solution enabled
```

To list an applied SAP Solution, run:

```
# saptune solution applied
```

If SAP Notes from an applied SAP Solution have been reverted, the string `(partial)` has been added to the solution name.

Both commands are primarily meant for use in scripts.

8.6 Verification and troubleshooting

To see the current status of `saptune`, run the following command:

```
# saptune status
```

The output contains the following:

- status of the `saptune`, `sapconf`, and `tuned` service

- version of package and running `saptune`

- details about configured SAP Solution and SAP Notes

- details about staging

- status of systemd system state

- virtualization environment (new in `saptune` version 3.1)

- tuning compliance (new in `saptune` version 3.1)

To analyze your `saptune` installation, run:

```
# saptune check
```

This command performs the following checks:

- check for mandatory or obsolete configuration files

- check for RPM leftovers

- check if the systemd system state is degraded and list failed units

- check the status of the sapconf, saptune and tuned services

Note

If `saptune check` warns about a degraded systemd system status, in most cases it has no impact on `saptune`. However, failed services require troubleshooting.

The command does not check the tuning itself. To check the tuning, use the command below:

```
# saptune note verify
```

> **Note**
>
> If `saptune note verify` is called without specifying a Note, it verifies all currently applied Notes. This allows you to verify your entire current tuning. As an alternative, use the `saptune solution verify` command that can also verify all currently applied Notes.

The `saptune note verify` command prints a table with all applied Notes, including the following:

- SAP Note and version

- the parameter

- the expected value of the parameter

- the value from an Override if one exists

- the current system value

- the compliance status of the parameter

The last line contains the overall compliance status of the entire tuning.

> **Note**
>
> Some parameters of shipped Notes are disabled, meaning they have empty values in the "Expected" column. In such cases, the SAP Note does not contain a concrete recommendation or `saptune` cannot detect the conditions for a recommendation. Read the SAP Note and set the value manually by customizing the Note (see *Section 8.4.1, "Customizing an SAP Note"*).

If parameters are not compliant, read the footnote if it exists. For some tunings, equivalent parameters exist, for example:

- `grub:intel_idle.max_cstate` covered by `force_latency`

- `grub:processor.max_cstate` covered by `force_latency`

- `grub:numa_balancing` covered by `kernel.numa_balancing`

- `grub:transparent_hugepage` covered by `THP`

A restart of the `saptune` service fixes the problems, except for non-compliant packages (parameter starts with `rpm:`) or GRUB entries (parameter starts with `grub:`). **saptune** does not install, uninstall or upgrade packages, and it never changes the boot loader.

A typical problem is the sysctl parameters that are handled by **saptune** and sysctl. A footnote in the parameter's compliance column indicates if it is also present in one of the sysctl configuration files. Remove the parameter from the sysctl configuration or disable the parameter in **saptune** (see *Section 8.4.1, "Customizing an SAP Note"*) to fix the problem.

Always investigate the cause for the changed tuning and fix it. If **saptune** shall not tune certain parameters, you can revert the Note or just disable parameters via an Override (see *Section 8.4.1, "Customizing an SAP Note"*).

8.7 Machine-readable output

Starting with version 3.1, **saptune** supports machine-readable output (JSON) for the following commands:

- **saptune [daemon|service] status**

- **saptune note list|verify|enabled|applied**

- **saptune solution list|verify|enabled|applied**

- **saptune status**

- **saptune version**

The machine-readable output makes it possible to integrate **saptune** into scripts and configuration management solutions.

To generate JSON output, add `--format json` as the first option, for example:

```
> saptune --format json note applied | jq
{
"$schema": "file:///usr/share/saptune/schemas/1.0/saptune_note_applied.schema.json",
"publish time": "2023-08-29 17:05:45.627",
"argv": "saptune --format json note applied",
"pid": 1538,
```

```
"command": "note applied",
"exit code": 0,
"result": {
    "Notes applied": [
    "941735",
    "1771258",
    "1980196",
    "2578899",
    "2684254",
    "2382421",
    "2534844",
    "2993054",
    "1656250"
    ]
},
"messages": []
}
```

If a command does not yet support JSON output, the command fails with the `result` block set to `"implemented": false`:

```
[+]
> saptune --format json staging status | jq
{
"$schema": "file:///usr/share/saptune/schemas/1.0/saptune_staging_status.schema.json",
"publish time": "2023-08-29 17:08:16.708",
"argv": "saptune --format json staging status",
"pid": 1653,
"command": "staging status",
"exit code": 1,
"result": {
    "implemented": false
},
"messages": []
}
```

8.8 Staging

It is possible that a new **saptune** package can contain both binary changes (for example, bug fixes) and new or altered SAP Notes and SAP Solutions. In certain situations, it is preferable to deploy bug fixes and new features while leaving modifications to the system configuration out.

With staging enabled, SAP Note and SAP Solution changes in a package update are *not* activated immediately. They are placed in a staging area, which can be reviewed and released later.

> **! Important**
>
> With the current implementation, a package update overwrites the staging if staging is enabled.

Staging is disabled by default, and it can be enabled with the following command:

```
# saptune staging enable
```

From that point, SAP Note and SAP Solution changes shipped by a `saptune` package are put in the staging area. To view the staging area, run:

```
# saptune staging list
```

You can print a tabular overview of the differences of the SAP Note and SAP Solution in the staging and working area with the following command:

```
# saptune staging diff [NOTE...|SOLUTION...|all]
```

After reviewing the differences, you can perform an analysis to see if a release has potential issues or requires additional steps. To do this, run the following command:

```
# saptune staging analysis [NOTE...|SOLUTION...|all]
```

To release an SAP Note or an SAP Solution from the staging area, use the command as follows:

```
# saptune staging [--force|--dry-run] [NOTE..|SOLUTION...|all]
```

The command presents an analysis (see `saptune staging analysis`) and carries out the release after asking for confirmation.

8.9 For more information

See the following man pages:

- `man 8 saptune`

- `man 8 saptune-migrate`

- `man 8 saptune-note`

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

9 Tuning Workload Memory Protection

> ⓘ **Important**
>
> If you have `systemd`-based SAP instances, read *Section 9.9, "Systems running both `systemd`-based and regular SAP instances"* and *Section 9.10, "Systems running only `systemd`-based instances"* before setting up Workload Memory Protection.

Keeping SAP applications in physical memory is essential for their performance. In older product versions, the Page Cache Limit prevented a swap out to disk by a growing page cache (in SUSE Linux Enterprise Server for SAP Applications 11 SP1 onwards and in SUSE Linux Enterprise Server for SAP Applications 12). In SUSE Linux Enterprise Server for SAP Applications 15, the Page Cache Limit has been replaced by the more advanced Workload Memory Protection.

Workload Memory Protection puts SAP instances into a dedicated cgroup (v2) and tells the kernel, by the `memory.low` parameter, the amount of memory to keep in physical memory. This protects the processes in this cgroup against any form of memory pressure outside that cgroup, including a growing page cache. Workload Memory Protection cannot protect against memory pressure inside this cgroup. It covers the memory of *all* instances together on one host.

The value for `memory.low` depends on the kind of SAP instance and the workload and needs to be configured manually. If the system is under extreme pressure, the Linux kernel will ignore the `memory.low` value and try to stabilize the whole system, even by swapping or invoking the OOM killer.

For more information about cgroups, see https://documentation.suse.com/sles-15/html/SLES-all/cha-tuning-cgroups.html↗.

9.1 Architecture

WMP relies on three components:

cgroup2 memory controller (Linux kernel)

The cgroup2 memory controller parameter `memory.low` allows defining an amount of memory, which the Linux kernel will keep in physical memory. This amount of memory will be excluded from the reclaiming process unless the entire system is in a critical memory situation.

WMP uses `memory.low` to prevent memory from SAP processes from being paged or swapped out to disk. Apart from the memory controller, cgroup1 controllers are still available, but are not mounted any more.

`systemd`

`systemd` provides the infrastructure to create and maintain the cgroup hierarchy and allows the configuration of cgroup parameters. WMP ships `systemd` configuration files to allow easy configuration of `memory.low` via `systemd` methods.

SAP start service

The SAP start service manages the starting and stopping of SAP instances. An important feature for WMP is the configurable execution of programs before the instance itself gets started in the instance profile. WMP uses this method to call a program to move the `sap-start` process into a designated cgroup, so the SAP instance will be started inside that cgroup.

9.2 Support for Workload Memory Protection

WMP is supported for SUSE Linux Enterprise Server for SAP Applications 15 SP4 on AMD64/ Intel 64 and POWER for one or multiple SAP systems on one host, such as:

- App Server (SAP NetWeaver, SAP S/4HANA) or

- SAP HANA 1.0/2.0

Workload Memory Protection does not cover databases other than SAP HANA. Depending on their start method, the processes might run inside or outside the dedicated cgroup. If they run inside, the memory consumption needs to be taken into account when determining `memory.low`.

> **!** Important: Restrictions of WMP

Using WMP comes with benefits, but you should be aware of some restrictions:

- WMP cannot protect against memory pressure inside the dedicated cgroup.

- WMP cannot protect SAP systems or their instances from each other. All SAP processes share the same memory limit. If you have multiple SAP systems (for example, SAP NetWeaver and SAP S/4HANA), WMP cannot shield one SAP application from the other.

- Support for SUSE's HA cluster solution is not yet available.

9.3   Setting up Workload Memory Protection

9.3.1   Preparing for Workload Memory Protection

1. Check if your SAP software (SAP HANA, SAP NetWeaver etc) is installed. The group `sapsys` is needed during the package installation of `sapwmp` later. If you skip that part, you will get a warning message (see *Important: Watch out for order of packages*).

2. Stop the SAP system:

   ```
   # systemctl stop sapinit
   ```

   The service can be enabled, but all SAP processes need to be terminated.

3. Install the package `sapwmp`:

   ```
   > sudo zypper install sapwmp
   ```

   > **!** Important: Watch out for order of packages
   >
   > The following message should only appear if no SAP software has been installed on the system:
   >
   > ```
   > Warning: sapsys group not found warning: group sapsys does not exist - using
   >   root
   > ```

Remove the package `sapwmp` and install the SAP software first before installing `sapwmp` again.

As an alternative, you can fix ownership and permission *after* installing the SAP software with:

```
> sudo chgrp sapsys /usr/lib/sapwmp/sapwmp-capture && \
chmod +s /usr/lib/sapwmp/sapwmp-capture
```

The following message can be ignored:

```
Warning: Found memory controller on v1 hierarchy. Make sure unified hierarchy only
 is used.
```

Switching to unified hierarchy is done in the next step.

4. Add `systemd.unified_cgroup_hierarchy=true` to the kernel command line by adding it to `GRUB_CMDLINE_LINUX_DEFAULT` in `/etc/default/grub` like:

```
GRUB_CMDLINE_LINUX_DEFAULT="... systemd.unified_cgroup_hierarchy=true swapaccount=1"
```

With this change, only cgroup2 controllers will be mounted on `/sys/fs/cgroup`. Cgroup1 controllers, except the memory controller, are still available and can be used though. Tools using cgroup1 might not work out of the box any more and might need reconfiguration. Also, the required mount structure for cgroup1 needs to be provided.

The parameter `swapaccount=1` is not needed for WMP to work, but it aids the analysis in support cases to show the amount of swapped out memory for each cgroup.

5. Rewrite the GRUB2 configuration:

```
> sudo grub2-mkconfig -o /boot/grub2/grub.cfg
```

After reboot (will be done later), the cgroup hierarchy is switched to v2 (unified hierarchy) only.

6. Configure `MemoryLow` for the `SAP.slice`:

```
> sudo systemctl set-property SAP.slice MemoryLow=...
```

This command creates a drop-in in `/etc/systemd/system.control/SAP.slice.d/` to set `MemoryLow`.

The `sapwmp` package includes the `systemd` configuration `SAP.slice` which creates the cgroup of the same name for the SAP instances. `MemoryLow` is the `systemd` equivalent of the cgroup parameter `memory.low` mentioned in the introduction. The value for `MemoryLow` depends on the type of the SAP application and the workload.

For SAP HANA

Since SAP HANA has a Global Allocation Limit, its value can be used directly.

SAP Application Server (SAP NetWeaver, SAP S/4HANA)

For the Application Server, the sizing for the workload should indicate the value for `MemoryLow`. The `sapwmp` package contains a monitoring part which might be useful to determine `MemoryLow`. See *Section 9.6, "Monitoring memory usage"*.

Keep in mind:

- All SAP instances on one host are inside the `SAP.slice`. `MemoryLow` must cover the amount of memory of *all* instances together on that host. You cannot protect SAP systems or their instances from each other.

- If you are using a database other than SAP HANA, some database processes might be part of `SAP.slice`. Their memory consumption needs to be taken into account when determining the `MemoryLow` value.

- Never choose a value for `MemoryLow` very close to or larger than your physical memory. System services and additional installed software require memory too. If they are forced to use swap too extensively, at the expense of the SAP application, your system can become unresponsive.

> ### Note: Correctly calculate `MemoryLow` value
>
> `MemoryLow` takes the memory size in bytes. If the value is suffixed with K, M, G, or T, the specified memory size is parsed as Kibibytes, Mebibytes, Gibibytes, or Tebibytes (with the base 1024 instead of 1000, see https://en.wikipedia.org/wiki/Binary_prefix ⬈), respectively. Alternatively, a percentage value may be specified, which is taken relative to the installed physical memory on the system.
>
> The underlying cgroup memory controller will round up the value to a multiple of the page size. To avoid confusion, set the value for `MemoryLow` to a multiple of the page size.

7. Create a backup of each SAP instance profile. Errors in a profile can prevent a SAP system from starting.

8. For each SAP instance, add the following line to the instance profile (usually located in `/usr/sap/`*SID*`/SYS/profile/`) after the last `Execute_` line:

```
Execute_20 = local /usr/lib/sapwmp/sapwmp-capture -a
```

If necessary, increase the number of the Execute statement so that it is the highest one, which means that that line is executed last.

> ❗ **Important: Editing instance profiles**
>
> Edit the instance profiles directly **only** if you do not have the profiles imported into the database to manage them via the SAP GUI (transaction RZ11). If you have imported them, use the SAP GUI to add the lines. Profile files located in the file system are overwritten and any manual changes would be lost!

Now the system is ready for a reboot.

9.3.2   Reboot and verification

1. Reboot the system.

2. After rebooting, verify that cgroups v2 has indeed been used:

```
# grep  cgroup /proc/mounts
cgroup /sys/fs/cgroup cgroup2 rw,nosuid,nodev,noexec,relatime 0 0
```

3. Verify that the cgroup was created successfully and the low memory value has been set:

```
> systemctl show -p MemoryLow SAP.slice
MemoryLow=18487889920    <- Should be your chosen value (always in bytes)!


# cat /sys/fs/cgroup/SAP.slice/memory.low
18487889920    <- Should be your chosen value!
```

The variable `MemoryLow` can be set to any value, but the content of the variable is always a multiple of the page size. Keep this in mind if you notice a slight difference between the values.

4. Check that all SAP instance processes are in the correct system slices/cgroup.

   If you have not enabled `sapinit.service` start the service now. If autostart is not enabled in the instance profiles, start the instances before you check.

   Example:

```
# systemd-cgls -a /sys/fs/cgroup/SAP.slice
Directory /sys/fs/cgroup/SAP.slice:
|-wmp-rd91fd6b3ca0d4c1183659ef4f9a092fa.scope
| |-3349 sapstart pf=/usr/sap/HA0/ERS10/profile/HA0_ERS10_sapha0er
| `-3375 er.sapHA0_ERS10 pf=/usr/sap/HA0/ERS10/profile/HA0_ERS10_sapha0er N...
|-wmp-r360ebfe09bcd4df4873ef69898576199.scope
| |-3572 sapstart pf=/usr/sap/HA0/SYS/profile/HA0_D01_sapha0ci
| |-3624 dw.sapHA0_D01 pf=/usr/sap/HA0/SYS/profile/HA0_D01_sapha0ci
...
```

   The `sapstartsrv` process of an instance always remains in the user slice of `SID`adm. Only the `sapstart` process and its children will be moved to the target cgroup.

   For each instance, a directory `wmp-rSCOPEID.scope` exists with all processes of this instance. The `SCOPEID` is a random 128-bit value in hexadecimal.

   The SAP HostAgent is not covered by WMP and remains partly in `sapinit.slice` and partly in the user slice of `sapadm`.

5. If the processes are not in the cgroup, check if the `Execute` lines in the instance profiles are correct. Also each instance start should now be logged in the system log `/var/log/messages`:

```
...
2020-06-16T18:41:28.317233+02:00 server-03 sapwmp-capture: Found PIDs:
2020-06-16T18:41:28.317624+02:00 server-03 sapwmp-capture:      17001
2020-06-16T18:41:28.317813+02:00 server-03 sapwmp-capture:      16994
2020-06-16T18:41:28.317959+02:00 server-03 sapwmp-capture:      16551
2020-06-16T18:41:28.319423+02:00 server-03 sapwmp-capture: Successful capture into
 SAP.slice/wmp-r07a27e12d7f2491f8ccb9aeb0e080aaa.scope
2020-06-16T18:41:28.319672+02:00 server-03 systemd[1]: Started wmp-
r07a27e12d7f2491f8ccb9aeb0e080aaa.scope.
...
```

To verify the correct setup, run **wmp-check**. The script checks the setup of Workload Memory Protection:

- Correct setup of cgroup2.

- Ownership and permission of the capture program.

- WMP entries of SAP instance profiles.

- Correct cgrop of running SAP instance processes.

- Correct setup of `SAP.slice`.

- Sane configuration of MemoryLow. However, it cannot determine if the MemoryLow value has been chosen wisely.

- Setup of the optional memory sampler.

- Setup of optional swap accounting.

It assumes SAP instances profiles can be found beneath `/usr/sap/SID/SYS/profile/`.

9.4 Configuring Workload Memory Protection

To configure WMP, edit `/etc/sapwmp.conf`:

```
# NOTE: Local changes may be reverted after update of WMP package. Check for
#       .rpmsave file to restore & merge changes.

## Description: Slice unit name where workload is put into
## Type:        string
## Default:     "SAP.slice"
DEFAULT_SLICE="SAP.slice"

## Description: Comma-separated list of command names to which capture is
##              applied (matching against /proc/$PID/stat)
## Type:        string
## Default:     sapstart
PARENT_COMMANDS=sapstart
```

After any change, restart all SAP instances.

✋ Warning

Altering `/etc/sapwmp.conf` should not be necessary. Do not do it until you know exactly what you are doing!

9.5 Changing the value of MemoryLow

To change the value of `MemoryLow` run:

```
# systemctl set-property SAP.slice MemoryLow=...
```

The changes will take effect immediately.

The underlying cgroup memory controller will round up the value to a multiple of the page size. To avoid confusion, set the value of `MemoryLow` to a multiple of the page size.

> **!** **Important: Value of `MemoryLow`**
>
> Never set `MemoryLow` to a value lower than the memory already accounted in `SAP.slice`. To check, run:
>
> ```
> # systemctl show -p MemoryCurrent SAP.slice
> ```

9.6 Monitoring memory usage

Logging the memory usage can be necessary not only to determine the value for `memory.low`, but also for monitoring the correct operation of WMP.

To enable monitoring, activate the shipped timer unit:

```
# systemctl enable --now wmp-sample-memory.timer
```

Now the timer should be listed by **systemctl list-timers**:

```
# systemctl list-timers
NEXT     LEFT       LAST     PASSED  UNIT                     ACTIVATES
...
Tue...  9min left  Tue...  4s ago  wmp-sample-memory.timer  wmp-sample-memory.service
...
```

If you check the current configuration, you can see that memory data is collected every 10 minutes with a randomized delay of three minutes:

```
# systemctl cat wmp-sample-memory.timer
# /usr/lib/systemd/system/wmp-sample-memory.timer
[Unit]
```

```
Description=WMP periodic log of memory consumption

[Timer]
OnCalendar=*:0/10
RandomizedDelaySec=180
AccuracySec=60

[Install]
WantedBy=timers.target
```

To change this, create a drop-in file and reload `systemd` (for example, by increasing the interval to 30 minutes):

```
# mkdir /etc/systemd/system/wmp-sample-memory.timer.d

# cat <<EOF >/etc/systemd/system/wmp-sample-memory.timer.d/override.conf
[Timer]
OnCalendar=
OnCalendar=*:0/30
EOF

# systemctl daemon-reload
```

(The first OnCalendar= line is important for deleting previously defined OnCalendar= settings.)

To see the memory consumption, check the system log for lines written by `wmp_memory_current`:

```
# grep wmp_memory_current /var/log/messages
...


2020-09-14T12:02:40.337266+02:00 server-03 wmp_memory_current: SAP.slice :
 memory.low=21474836480 memory.current=2294059008 memory.swap.current=0 , user.slice :
 memory.low=0 memory.current=5499219968 memory.swap.current=0 , init.scope :
 memory.low=0 memory.current=8364032 memory.swap.current=0 , system.slice : memory.low=0
 memory.current=1863335936 memory.swap.current=0
2020-09-14T12:03:00.767838+02:00 server-03 wmp_memory_current: SAP.slice :
 memory.low=21474836480 memory.current=2294022144 memory.swap.current=0 , user.slice :
 memory.low=0 memory.current=5499473920 memory.swap.current=0 , init.scope :
 memory.low=0 memory.current=8364032 memory.swap.current=0 , system.slice : memory.low=0
 memory.current=1862586368 memory.swap.current=0
2020-09-14T12:04:00.337315+02:00 server-03 wmp_memory_current: SAP.slice :
 memory.low=21474836480 memory.current=2294022144 memory.swap.current=0 , user.slice :
 memory.low=0 memory.current=5499207680 memory.swap.current=0 , init.scope :
 memory.low=0 memory.current=8355840 memory.swap.current=0 , system.slice : memory.low=0
 memory.current=1862746112 memory.swap.current=0
```

```
...
```

Here is a reformatted log line to get a better impression:

```
2020-09-14T12:02:40.337266+02:00 server-03 wmp_memory_current:
SAP.slice    : memory.low=21474836480 memory.current=2294059008 memory.swap.current=0 ,
user.slice   : memory.low=0            memory.current=5499219968 memory.swap.current=0 ,
init.scope   : memory.low=0            memory.current=8364032     memory.swap.current=0 ,
system.slice : memory.low=0            memory.current=1863335936 memory.swap.current=0
```

For each cgroup directly below `/sys/fs/cgroup/` one comma-separated block exists. On a normal system, you should find at least `user.slice`, `system.slice`, and `init.scope`. WMP adds `SAP.slice`.

Each block contains the information about the current value of `memory.low` and `memory.current`, and the currently allocated amount of physical memory of processes in this cgroup.

If you enabled swap accounting (`swapaccount=1`) during setup, you also have `memory.swap.current`, the amount of swapped-out memory of the cgroup.

All values are in bytes. See *Step 6* in *Section 9.3.1, "Preparing for Workload Memory Protection"*.

> ### 💡 Tip: Script for printing
> You can find a script to print the information as table or CSV here: https://github.com/scmschmidt/wmp_log_extract ↗

9.7 Verifying correct operation

Besides monitoring memory consumption and swapping (see *Section 9.6, "Monitoring memory usage"*), you should also regularly check that all SAP instance processes are in their scopes below `SAP.slice`.

To do so, run **systemd-cgls** and check each instance process.

Example:

```
# systemd-cgls -a /sys/fs/cgroup/SAP.slice
Directory /sys/fs/cgroup/SAP.slice:
|-wmp-rd91fd6b3ca0d4c1183659ef4f9a092fa.scope
| |-3349 sapstart pf=/usr/sap/HA0/ERS10/profile/HA0_ERS10_sapha0er
| `-3375 er.sapHA0_ERS10 pf=/usr/sap/HA0/ERS10/profile/HA0_ERS10_sapha0er N...
|-wmp-r360ebfe09bcd4df4873ef69898576199.scope
```

```
| |-3572 sapstart pf=/usr/sap/HA0/SYS/profile/HA0_D01_sapha0ci
| |-3624 dw.sapHA0_D01 pf=/usr/sap/HA0/SYS/profile/HA0_D01_sapha0ci
...
```

A simpler test would be to list all processes, including cgroups, for all `SID`s used on the system. Example:

```
> ps -eo user,pid,cgroup:60,args | grep -e [h]a0adm
ha0adm    2062 0::/user.slice/user-1001.slice/user@1001.service/init.scope  /usr/lib/
systemd/systemd --user
ha0adm    2065 0::/user.slice/user-1001.slice/user@1001.service/init.scope  (sd-pam)
ha0adm    3081 0::/SAP.slice/wmp-r73c594e050904c9c922a312dd9a28fd4.scope    sapstart pf=/
usr/sap/HA0/SYS/profile/HA0_ASCS00_sapha0as
ha0adm    3133 0::/SAP.slice/wmp-r73c594e050904c9c922a312dd9a28fd4.scope
 ms.sapHA0_ASCS00 pf=/usr/sap/HA0/SYS/profile/HA0_ASCS00_sapha0as
ha0adm    3134 0::/SAP.slice/wmp-r73c594e050904c9c922a312dd9a28fd4.scope
 en.sapHA0_ASCS00 pf=/usr/sap/HA0/SYS/profile/HA0_ASCS00_sapha0as
ha0adm    3327 0::/SAP.slice/wmp-ra42489517eb846c282c57681e627a496.scope    sapstart pf=/
usr/sap/HA0/ERS10/profile/HA0_ERS10_sapha0er
...
```

All instance processes except `sapstartsrv` need to be in a scope below `0::/SAP.slice/`.

To verify the correct setup, use the **wmp-check** tool. See *Section 9.3.2, "Reboot and verification"* for more details.

9.8 Uninstalling Workload Memory Protection

1. Stop the SAP system completely. The `sapinit.service` has to be stopped, but can stay enabled. All SAP processes have to be terminated.

2. Remove any changes made to `SAP.slice`, such as setting `MemoryLow`:

   ```
   # systemctl revert SAP.slice
   ```

3. *(Optional)* Remove the package `sapwmp`:

   ```
   # zypper remove sapwmp
   ```

   This step is optional. The package can stay on the system without having an influence.

4. *(Optional)*   Remove   `systemd.unified_cgroup_hierarchy=true`   from `GRUB_CMDLINE_LINUX_DEFAULT` in `/etc/default/grub`.

This step is optional. You can keep cgroup2 without using WMP.

5. Rewrite the GRUB2 configuration:

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

After the next boot, the system is switched back to the hybrid cgroup hierarchy.

6. Remove the line to call sapwmp-capture from each SAP instance profile (usually located in `/usr/sap/SID/SYS/profile/`):

```
Execute_20 = local /usr/lib/sapwmp/sapwmp-capture -a
```

> **❗ Important: Backup is necessary**
>
> Before editing an instance profile, create a backup! Errors in a profile can prevent an SAP system from starting!

> **❗ Important: About editing profiles directly**
>
> Edit the instance profiles directly **only** if you have not imported the profiles into the database to manage them via the SAP GUI (transaction RZ11). If you have imported them, use the SAP GUI to add the lines. Profile files located in the file system are overwritten, and any manual changes will be lost!

7. Reboot the system and verify that your SAP system has been started successfully.

9.9 Systems running both `systemd`-based and regular SAP instances

With SAP kernel version 788 and certain specified patch levels of older kernel releases (722, 753, 777, 781, 785), the SAP Start Service supports `systemd`. Such a `systemd`-based instance comes with its own `systemd` service that the SAP Start Service places into its own cgroup (`SAPSID_NR.service`) under `SAP.slice`. This affects the Workload Memory Protection setup. For `systemd`-based SAP instances, the following instructions apply:

- Skip the step that adds the capture program to the instance profile. If you are migrating an instance to support `systemd`, remove this entry from the profile. Keep in mind that `systemd`-based instances are always put into a cgroup under `SAP.slice` and therefore become part of Workload Memory Protection protection.

- Set `MemoryLow=infinity` for the `SAPSID_NR.service` to make the protection work correctly. Example for `SAPNW1_01.service`:

  ```
  sudo systemctl set-property SAPNW1_01.service MemoryLow=infinity
  ```

- Set `MemoryLow=infinity` for the `saphostagent.service` to make the protection work correctly

  ```
  sudo systemctl set-property saphostagent.service MemoryLow=infinity
  ```

The following example demonstrates a mixed environment. The SAP Host Agent and instance 01 are `systemd`-based, instance 00 is not. Both instances are under `SAP.slice` either in the cgroup managed by the SAP Start Service (`SAPNW1_01.service`) or Workload Memory Protection (`wmp-rece5b7fa372e4619a9623e120aa23a23.scope`). `MemoryLow=` for `SAP.slice` has been set as well as `MemoryLow=infinity` for all the cgroups below.

```
>  systemd-cgls
Control group /:
-.slice
├─SAP.slice
│ ├─SAPNW1_01.service
│ │ ├─ 887 /usr/sap/NW1/ASCS01/exe/sapstartsrv pf=/usr/sap/NW1/SYS/profile/
NW1_ASCS01_systemdproj
│ │ ├─2838 sapstart pf=/usr/sap/NW1/SYS/profile/NW1_ASCS01_systemdproj
│ │ ├─2895 ms.sapNW1_ASCS01 pf=/usr/sap/NW1/SYS/profile/NW1_ASCS01_systemdproj
│ │ └─2896 en.sapNW1_ASCS01 pf=/usr/sap/NW1/SYS/profile/NW1_ASCS01_systemdproj
│ ├─saphostagent.service
```

```
| | ├─ 900 /usr/sap/hostctrl/exe/saphostexec pf=/usr/sap/hostctrl/exe/host_profile -
nodaemon -trace
| | ├─ 984 /usr/sap/hostctrl/exe/sapstartsrv pf=/usr/sap/hostctrl/exe/host_profile
| | └─2428 /usr/sap/hostctrl/exe/saposcol -l -w60 pf=/usr/sap/hostctrl/exe/host_profile
| └─wmp-rece5b7fa372e4619a9623e120aa23a23.scope
|   ├─5522 sapstart pf=/usr/sap/NW1/SYS/profile/NW1_D00_systemdproj
|   ├─7824 dw.sapNW1_D00 pf=/usr/sap/NW1/SYS/profile/NW1_D00_systemdproj
|   ├─7825 ig.sapNW1_D00 -mode=profile pf=/usr/sap/NW1/SYS/profile/NW1_D00_systemdproj
...

> systemctl show -p MemoryLow SAP.slice
MemoryLow=24584065024

> systemctl show -p MemoryLow wmp-rece5b7fa372e4619a9623e120aa23a23.scope
MemoryLow=infinity

> systemctl show -p MemoryLow SAPNW1_01.service
MemoryLow=infinity

> systemctl show -p MemoryLow saphostagent.service
MemoryLow=infinity
```

Further information about the `systemd` integration is available at SAP Note 139184 - Linux: `systemd` integration for sapstartsrv and SAP Host Agent (https://launchpad.support.sap.com/#/notes/3139184)↗.

9.10 Systems running only `systemd`-based instances

> **Note**
>
> SUSE's HA cluster solutions are supported for fully `systemd`-based SAP systems.

An SAP system consisting of `systemd`-based instances can only rely on cgroup2 for memory protection.

The instances are placed into a dedicated cgroup (default is `SAP.slice`) by the SAP Start Service. Therefore it is only necessary to switch to the unified cgroup2 hierarchy and set `MemoryLow=` correctly.

If Workload Memory Protection is already configured, follow the instructions in *Section 9.9, "Systems running both* `systemd`*-based and regular SAP instances"*. Afterwards, you can remove the package `sapwmp` from the system, unless you want to use the `wmp-sample-memory.timer` and `wmp-sample-memory.service` to monitor memory usage.

To configure the memory protection for a system consisting of only `systemd`-based SAP instances, perform the following steps:

- Stop all SAP instances and the SAP Host Agent.

- Switch to the unified cgroup hierarchy, rewrite the boot loader (step 4 and 5 in *Section 9.3.1, "Preparing for Workload Memory Protection"*).

- Set `MemoryLow=` for `SAP.slice` (step 6 in *Section 9.3.1, "Preparing for Workload Memory Protection"*).

- Set `MemoryLow=infinity` for `saphostagent.service` and for the services of all SAP instances.

- Reboot the system.

The following example demonstrates a `systemd`-based environment. The SAP Host Agent and all instances are in their cgroups under `SAP.slice`. `MemoryLow=` for `SAP.slice` has been set as well as `MemoryLow=infinity` for all the cgroups below.

```
>  systemd-cgls
Control group /:
-.slice
├─SAP.slice
| ├─SAPNW1_00.service
| | ├─5522 sapstart pf=/usr/sap/NW1/SYS/profile/NW1_D00_systemdproj
| | ├─7824 dw.sapNW1_D00 pf=/usr/sap/NW1/SYS/profile/NW1_D00_systemdproj
| | ├─7825 ig.sapNW1_D00 -mode=profile pf=/usr/sap/NW1/SYS/profile/NW1_D00_systemdproj
...
| ├─SAPNW1_01.service
| | ├─ 887 /usr/sap/NW1/ASCS01/exe/sapstartsrv pf=/usr/sap/NW1/SYS/profile/
NW1_ASCS01_systemdproj
| | ├─2838 sapstart pf=/usr/sap/NW1/SYS/profile/NW1_ASCS01_systemdproj
| | ├─2895 ms.sapNW1_ASCS01 pf=/usr/sap/NW1/SYS/profile/NW1_ASCS01_systemdproj
| | └─2896 en.sapNW1_ASCS01 pf=/usr/sap/NW1/SYS/profile/NW1_ASCS01_systemdproj
| └─saphostagent.service
|    ├─ 900 /usr/sap/hostctrl/exe/saphostexec pf=/usr/sap/hostctrl/exe/host_profile -
nodaemon -trace
|    ├─ 984 /usr/sap/hostctrl/exe/sapstartsrv pf=/usr/sap/hostctrl/exe/host_profile
|    └─2428 /usr/sap/hostctrl/exe/saposcol -l -w60 pf=/usr/sap/hostctrl/exe/host_profile
```

```
...

> systemctl show -p MemoryLow SAP.slice
MemoryLow=24584065024

> systemctl show -p MemoryLow wmp-rece5b7fa372e4619a9623e120aa23a23.scope
MemoryLow=infinity

> systemctl show -p MemoryLow SAPNW1_01.service
MemoryLow=infinity

> systemctl show -p MemoryLow saphostagent.service
MemoryLow=infinity
```

10 Configuring a firewall

This chapter provides information about restricting access to the system using a firewall and encryption and gives information about connecting to the system remotely.

10.1 Configuring `firewalld`

By default, the installation workflow of SUSE Linux Enterprise Server for SAP Applications enables `firewalld`.

> ### Note: `firewalld` replaces SuSEfirewall2
>
> SUSE Linux Enterprise Server for SAP Applications 15 introduces `firewalld` as the new default software firewall, replacing SuSEfirewall2. SuSEfirewall2 has not been removed from SUSE Linux Enterprise Server for SAP Applications 15 and is still part of the main repository, but it is not installed by default. If you are upgrading from a release older than SUSE Linux Enterprise Server for SAP Applications 15, SuSEfirewall2 will be unchanged and you must manually upgrade to `firewalld` (see Security and Hardening Guide).

The firewall must be manually configured to allow network access for the following components:

- SAP application

- Database (see the documentation of your database vendor; for SAP HANA, see *Section 10.2, "Configuring HANA-Firewall"*)

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

SAP applications require multiple 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.

10.2 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`. Creates XML files containing firewall rules for SAP HANA.

  Instead of using YaST, you can configure firewall rules using the configuration file at `/etc/sysconfig/hana-firewall`.

> **!** **Important: SAP HANA MDC databases**
>
> For multi-tenant SAP HANA (MDC) databases, determining automatically the port numbers that need to be opened is not yet possible. If you are working with a multi-tenant SAP HANA database system, run a script to create a new service definition before using YaST:
>
> ```
> # cd /etc/hana-firewall.d
> # hana-firewall define-new-hana-service
> ```
>
> The script prompts you to answer a series of questions, including 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 10.1: 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* › *Configure system firewall for SAP HANA*.

3. Under *Global Options*, activate *Enable and reload firewalld*.

4. Select the desired zone from the *Zone* drop-down list, and add the required services using the right arrow button.

To add services other than the preconfigured ones, use the following notation:

```
SERVICE_NAME:CIDR_NOTATION
```

For more information about the CIDR notation, see https://en.wikipedia.org/wiki/Classless_Inter-Domain_Routing ↗. To find out which services are available on your system, use `getent services`.

5. 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.

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

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

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

10.3   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 that 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 `saprouter-systemd`

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

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`

11 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 15 SP4 supports NW-VSI version 2.0.

> **❗ Important: Avoid false positive reports for large files exceeding maximum file size**
>
> By default, ClamAV does not scan files exceeding various limits like file sizes, nesting level, or scan time. Such files are reported as "OK". The current default settings for the ClamAV virus scan engine in the `clamscan` commandline tool and the `clamd` scan daemon are set in a way that:
>
> - Files and archives are scanned, but only up to the configured or default limits for size, nesting level, scan time, etc.
>
> - The scan engine reports these files as being "OK".
>
> - This could potentially allow attackers to bypass the virus scanning.
>
> Alerts can be enabled to set the `--alert-exceeds-max=yes` option on the `clamscan` commandline or via `AlertExceedsMax TRUE` in `clamd.conf` for daemon based scans. Settings these options will cause a "FOUND" report of status type `Heuristics.Limits.Exceeded`. You need to handle such files differently in front-ends or processing of reports.
>
> Before enabling the alert, ensure that front-ends will not suddenly quarantine or remove those files.

11.1 Installing ClamSAP

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

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

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

```
> sudo freshclam
```

3. Start the service `clamd`:

```
> sudo systemctl start clamd
```

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

```
> 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
[...]
```

11.2 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*.
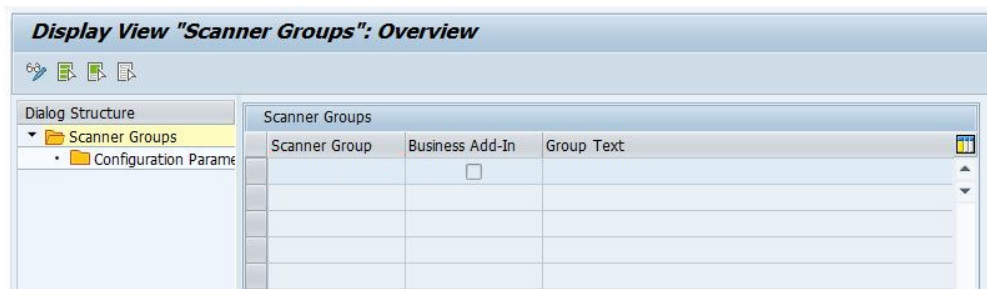
3. To switch from view mode to change mode, click the button *Change View* (🔧).
   Confirm the message *This table is cross-client* by clicking the check mark. The table is now editable.

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.
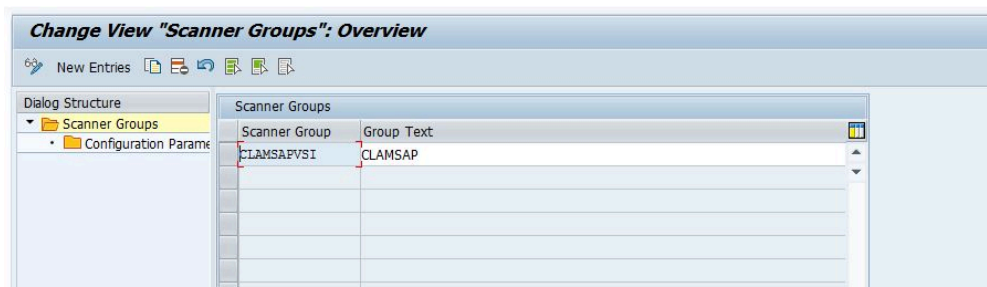
5. To save the form, click the button *Save* (💾).

11.3 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 11.2, "Creating a virus scanner group in SAP NetWeaver"* (for example: `CLAMSAPVSI`)

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

   - *Adapter Path*: `libclamdsap.so`

New Entries: Details of Added Entries

Provider Type	ADAPTER (Virus Scan Adapter) ▼
Provider Name	VSA_<HOSTNAME>
Status		

Start   Stop

Virus Scan Provider Definition

Scanner Group	☑ Display
Status	Active (Application Server) ▼
Server	<hostname>_<SID>_<instance number>
Reinit. Interv.	Hours Load
Adapter Path	libclamdsap.so

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

11.4 Configuring the default location of virus definitions

By default, ClamAV expects the virus definitions to be located in `/var/lib/clamsap`. To change this default location, proceed as follows:

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. Select the `CLAMSAPVSI` group.

3. In the left navigation pane, click *Configuration Parameters*.

4. 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.
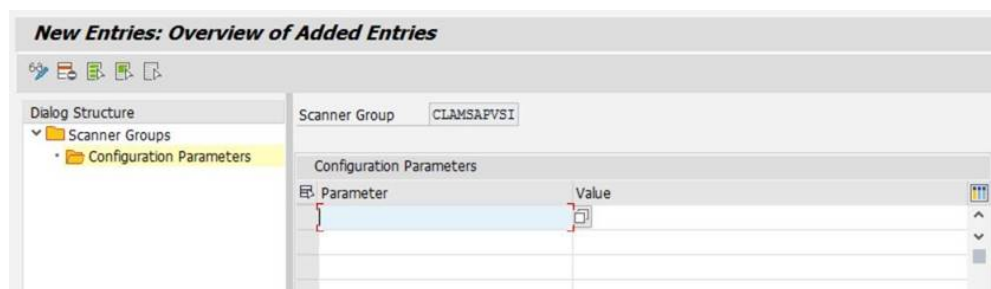
FIGURE 11.1:

5. Click *New Entries* and select `INITDRIVERDIRECTORY`.
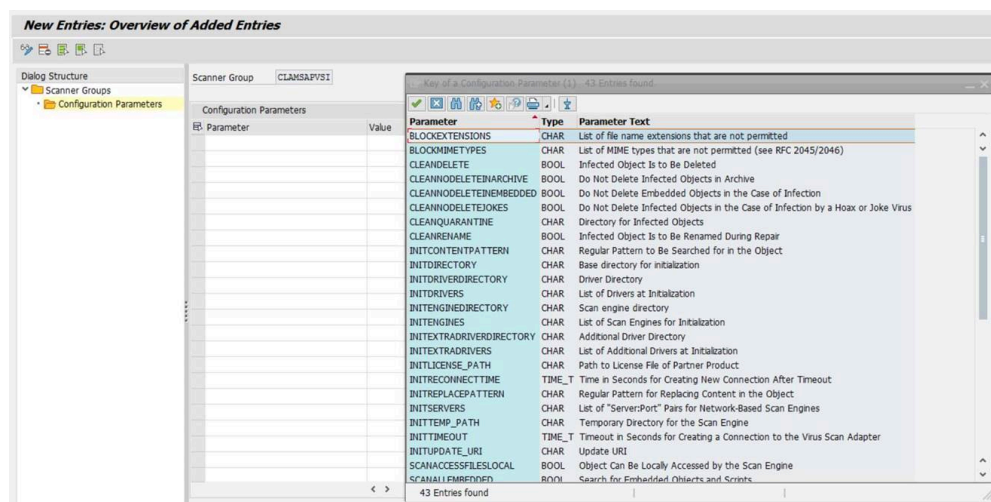
FIGURE 11.2:

6. Enter the path to a different virus scanner location.

7. To save the form, click the button *Save* (💾).

11.5 Engaging ClamSAP

To run ClamSAP, go to the transaction *VSCAN*. Then click *Start*.

FIGURE 11.3: CHANGE VIEW "VIRUS SCAN PROVIDER DEFINITION"

Afterward, a summary will be displayed, including details of the ClamSAP and ClamAV (shown in *Figure 11.4, "Summary of ClamSAP data"*).

FIGURE 11.4: **SUMMARY OF CLAMSAP DATA**

11.6 For more information

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

12 Connecting via RDP

If you installed SLES for SAP with the RDP option activated or if you installed from a KIWI NG 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 WE 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 12.1: 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 the necessary exception for your firewall, opening port TCP 3389 in all relevant zones. For example, if your internal network uses the `internal` zone, use the following command:

   ```
   # firewall-cmd --zone=internal --add-port=3389/tcp
   ```

   This is a temporary assignment for testing the new setting. If you need to change more than one zone, change and test each zone one at a time.

2. Make the new configuration permanent:

   ```
   # firewall-cmd --runtime-to-permanent
   # firewall-cmd --reload
   ```

   For more information on using firewalld, refer to https://docs.suse.com/sles/15/html/SLES-all/cha-security-firewall.html#sec-security-firewall-firewalld ↗.

3. Nxt, set up `xrdp`.

   Install the package `xrdp`:

   ```
   # zypper install xrdp
   ```

4. Enable and start the `xrdp` service:

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

   You can now connect to the machine.

13 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 NG, 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.

13.1 Creating images with KIWI NG

KIWI NG 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 for SAP images with KIWI NG.

SUSE Linux Enterprise Server for SAP Applications now supports creating images with KIWI NG 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`, as 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:

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

2. Build the VMX image:

   ```
   # 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 12, Connecting via RDP*.

For more information about KIWI NG and SLES for SAP:

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

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

13.2 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 NG image from scratch. For example, when your image needs to contain additional software or configuration that cannot be installed using KIWI NG.

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.

13.2.1 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`.

13.2.2 Using `clone-master-clean-up`

To use the script, do:

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

Then follow the instructions.

13.2.3 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`.

14 Important log files

The most important log files for this product can be found as follows:

- 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 for 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-15 ↗.

- *SUSE Connect Program* allows installing packages created and supported by third parties, specifically for SLES for 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 the SUSE Linux Enterprise community without support. See *Section A.3, "SUSE Package Hub"*.

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 may have been changed. If you need further details, refer to the article at https://www.suse.com/support/kb/doc/?id=7023490 ↗.

`/etc/products.d/baseproduct`

A link to an XML file. The `/etc/products.d/` directory contains different `.prod` files. Depending on which products you have purchased and how you installed your system, the link `/etc/products.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/) ↗.

You can extract any information from the file `/etc/products.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 you want to avoid the XML tags in the output, use the **xmlstarlet** command instead.

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

```
> 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:

```
> 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 that `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:

```
>> 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 SUSE Package Hub

SUSE Package Hub 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 SUSE Package Hub, add the repository as described at https://package-hub.suse.com/how-to-use/ ↗.

For more information, see the SUSE Package Hub 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 illustrate 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-15 ↗).

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.

The two XML files (`autoyast.xml` and `product.xml`) will be merged after the media is read and 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-15 ↗).

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-15 ↗).

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 and 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, because 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-15 ↗). 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 the Supplementary Media directory contains only a file called `product.xml`.

D Cheat sheet for Windows administrators

D.1 Managing users

To manage users, launch YaST and switch to *User and Group Management.* To use the ncurses version of YaST, run the `sudo /sbin/yast2 users` command. For more information, refer to https://documentation.suse.com/sles/html/SLES-all/cha-yast-userman.html/ ⬈ .

D.2 Assigning administrator privileges

On Linux, administrator privileges are defined in the `/etc/sudoers` file. You can use YaST to manage administrator privileges. Install the required YaST module using the `sudo zypper in yast2-sudo` command. Launch YaST and switch to the *Sudo* section. To use the ncurses version of YaST, run the `sudo /sbin/yast2 sudo` command. For more information, see https://documentation.suse.com/sles/single-html/SLES-administration/#cha-adm-sudo/ ⬈ .

D.3 Managing system services

Use the *Services Manager* module in YaST to view and manage enabled services. To use the ncurses version of YaST, run the `sudo /sbin/yast2 services-manager` command. For more information, see https://documentation.suse.com/sles/single-html/SLES-administration/#cha-sys-temd/ ⬈ .

D.4 Managing firewall settings

The *Firewall* module in YaST allows you to manage the firewall settings. To use the ncurses version of YaST, run the `sudo /sbin/yast2 firewall` command.

The firewall on SUSE Linux Enterprise allows setting rules for each interface independently. You can also enable masquerading, port forwarding and broadcasting in the firewall settings. For more information, see https://documentation.suse.com/sles/html/SLES-all/cha-security-firewall.html/ ⬈ .

D.5 Joining a Windows domain (Active Directory/SMB file sharing)

Install the `yast2-auth-client` package first. Then use the *User Logon* module to join a Windows domain. To use the ncurses version of YaST, run the **`sudo /sbin/yast2 auth-client`** command. For more information, see https://documentation.suse.com/sles/html/SLES-all/cha-security-auth.html/ ↗ .

D.6 Managing partitions and storage devices

Use the *Partitioner* module in YaST to view or change the partition layout. To use the ncurses version of YaST, run the **`sudo /sbin/yast2 disk`** command.

To avoid data loss, unmount partitions before modifying them! To unmount a partition, open a terminal and run the **`mount`** command. This returns a list of entries structured like this: `DEVICE on MOUNT_POINT type FILE_SYSTEM_TYPE (FILE_SYSTEM_OPTIONS)`. To unmount the desired partition, use the **`sudo umount MOUNT_POINT`** command with the appropriate mount point. For more information, see https://documentation.suse.com/sles/html/SLES-all/book-storage.html/ ↗ .

D.7 Creating a Windows share

On Linux, Samba implements the SMB protocol that makes it possible to create Windows shares. Use the *Samba Server* module in YaST to set up an SMB server. To use the ncurses version of YaST, run the **`sudo /sbin/yast2 samba-server`** command. For more information, see https://documentation.suse.com/sles/15-SP4/html/SLES-all/cha-samba.html/ ↗ .

E GNU licenses

This appendix contains the GNU Free Documentation License version 1.2.

GNU Free Documentation License

Copyright (C) 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondarily, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or non-commercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.

B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.

C. State on the Title page the name of the publisher of the Modified Version, as the publisher.

D. Preserve all the copyright notices of the Document.

E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.

F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.

G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.

H. Include an unaltered copy of this License.

I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.

J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.

K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.

L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.

M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.

N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.

O. Preserve any Warranty Disclaimers.

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

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

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

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

5. COMBINING DOCUMENTS

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

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

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

6. COLLECTIONS OF DOCUMENTS

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

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

7. AGGREGATION WITH INDEPENDENT WORKS

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

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

8. TRANSLATION

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

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

9. TERMINATION

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

10. FUTURE REVISIONS OF THIS LICENSE

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

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

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

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

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

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

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